

Deliberation and Voting in Approval-Based Multi-Winner Elections

Kanav Mehra , Nanda Kishore Sreenivas and Kate Larson

David R. Cheriton School of Computer Science, University of Waterloo

{kanav.mehra, nksreenivas, kate.larson}@uwaterloo.ca

Abstract

Citizen-focused democratic processes where participants deliberate on alternatives and then vote to make the final decision are increasingly popular today. While the computational social choice literature has extensively investigated voting rules, there is limited work that explicitly looks at the interplay of the deliberative process and voting. In this paper, we build a deliberation model using established models from the opinion-dynamics literature and study the effect of different deliberation mechanisms on voting outcomes achieved when using well-studied voting rules. Our results show that deliberation generally improves welfare and representation guarantees, but the results are sensitive to how the deliberation process is organized. We also show, experimentally, that simple voting rules, such as approval voting, perform as well as more sophisticated rules such as proportional approval voting or method of equal shares if deliberation is properly supported. This has ramifications on the practical use of such voting rules in citizen-focused democratic processes.

1 Introduction

Scenarios, where a committee must be selected to represent the interests of some larger group, are ubiquitous, ranging from political domains such as parliamentary elections and participatory budgeting (PB) [Cabannes, 2004] to technical applications such as designing recommender systems [Skowron *et al.*, 2016] and diversifying search results [Skowron *et al.*, 2017]. *Multi-winner* voting has been well studied within the social choice literature, with a focus on understanding how the “best” committee can be selected. However, even defining what is meant by “best” is no trivial undertaking. In some contexts, such as aggregation of expert judgements, one might want a committee that consists of the highest-rated k alternatives. However, in other tasks such as choosing k locations for constructing a public facility (*e.g.* hospitals, fire stations), it is preferable to ensure as many voters as possible have access to the facility.

In citizen-focused democratic processes such as citizens’ assemblies [Elstub *et al.*, 2022] and participatory budgeting

[Cabannes, 2004], there exists extensive scope for discussion over the multitude of possible alternatives. For example, deliberation is an important phase in most implementations of participatory budgeting as it allows voters to refine their preferences and facilitates the exchange of information, with the objective of reaching consensus [Aziz and Shah, 2021]. While deliberation is a vital component of democratic processes [Fishkin, 2009; Habermas, 1996], it cannot completely replace voting because, in reality, deliberation does not guarantee unanimity. A decision must still be made. Accordingly, we argue that it is essential to understand the relationship between voting and deliberation. To this end, we (1) introduce an agent-based model of deliberation, and (2) study the effect of different deliberation mechanisms on the outcomes obtained when using well-studied voting rules.

We focus on approval-based elections, where voters express preferences by sharing a subset of approved candidates. Approval ballots are used in practice due to their simplicity and flexibility [Brams and Fishburn, 2007; Benadè *et al.*, 2021; Aziz and Shah, 2021]. They also offer scope for deliberation as often voters are left to decide between many different alternatives. We present an agent-based model of deliberation and explore various alternatives for structuring deliberation groups. We evaluate standard multi-winner voting rules, both before and after voters have the opportunity to deliberate, with respect to standard objectives from the literature, including social welfare, representation, and proportionality. We show that deliberation, in almost all scenarios, significantly improves welfare, representation, and proportionality. However, the results are sensitive to the deliberation mechanism; increased exposure to diverse opinions (or agents from different backgrounds) leads to better outcomes. More importantly, our results indicate that in the presence of effective deliberation, *simple*, explainable voting rules such as approval voting perform as well as more sophisticated, *complex* rules. This can serve to guide the design and deployment of voting rules in citizen-focused democratic processes.

Related Work: The social choice literature has extensively studied the quality of approval-based multi-winner voting rules. From the quantitative perspective, a recent paper [Lackner and Skowron, 2020] provides an in-depth theoretical and empirical analysis of different approval-based multi-winner voting rules with respect to (utilitarian) social welfare and representation guarantees. [Fairstein *et al.*, 2022]

extended this work to study the welfare-representation trade-off in the more general PB setting. The traditional axiomatic approach, on the other hand, provides a qualitative evaluation, *i.e.* whether a voting rule satisfies a property or not. For approval-based rules, recent work has focused heavily on proportionality axioms [Aziz *et al.*, 2017; Sánchez-Fernández *et al.*, 2017; Aziz *et al.*, 2018; Brill *et al.*, 2017; Lackner and Skowron, 2018; Skowron, 2021]. We refer the reader to [Faliszewski *et al.*, 2017] for an extensive survey on the properties of multi-winner rules.

Deliberation, specifically within social choice, has been studied through multiple approaches. From the theoretical perspective, a wide variety of mathematical deliberation models have been proposed [Chung and Duggan, 2020; Zvi *et al.*, 2021]. For example, [Goel and Lee, 2016; Fain *et al.*, 2017] introduce iterative small-group deliberation mechanisms for reaching consensus in collective decision-making problems. [Elkind *et al.*, 2021] propose a consensus-reaching deliberation protocol based on coalition formation. A recent experimental study [Rad and Roy, 2021] shows that deliberation leads to meta-agreements and single-peaked preferences under specific conditions. [Perote-Peña and Piggin, 2015] look at deliberation and voting simultaneously, but their work is limited to the *ground-truth* setting with ordinal preferences over three alternatives. They do not study the impact of deliberation on the quantitative and qualitative properties of voting rules.

In this paper, we bridge the gap between deliberation and voting literature. To our knowledge, we are the first to experimentally study the effect of deliberation on voting outcomes across different deliberation strategies.

2 Preliminaries

Let $E = (C, N)$ be an election, where $C = \{c_1, c_2, \dots, c_m\}$ and $N = \{1, \dots, n\}$ are sets of m candidates and n voters, respectively. Each voter $i \in N$, has an *approval ballot* $A_i \subseteq C$, containing the set of its approved candidates. The *approval profile* $A = \{A_1, A_2, \dots, A_n\}$ represents the approval ballots for all voters. For a candidate $c_j \in C$, $N(c_j)$ is the set of voters that approve c_j and its *approval score*, $V(c_j) = |N(c_j)|$. Let $S_k(C)$ denote all k -sized subsets of the candidate set C , representing the set of all possible committees of size k . Given approval profile A and desired committee size $k \in \mathbb{N}$, the objective of a multi-winner election is to select a subset of candidates that form the winning committee $W \in S_k(C)$. An *approval-based committee rule*, $R(A, k)$, is a social choice function that takes as input an approval profile A and committee size k and returns a set of *winning committees*.¹ For any voting rule $R(A, k)$, we will use W_R to denote its selected committee (after tie-breaking).

2.1 Properties

We ideally want our voting rules to exhibit certain desired properties, representing the principles that should govern the selection of winners given individual ballots. In this paper, we compare voting rules across three dimensions: *social welfare*,

¹A tie-breaking method is used to pick one winning committee in cases where multiple winning committees are returned.

representation, and *proportionality*. Intuitively, the *welfare* objective focuses on selecting candidates that garner maximum support from the voters. *Representation* cares about *diversity*; carefully selecting a committee that maximizes the number of voters represented in the winning committee. A voter is represented if the final committee contains at least one of its approved candidates.

Definition 1 (Utilitarian Social Welfare). *For a given approval profile A and committee size k , the utilitarian social welfare of a committee W is:*

$$SW(A, W) = \sum_{i \in N} \sum_{c \in W} u_i(c), \quad (1)$$

$u_i(c) \in \mathbb{R}$ is the utility voter i derives from candidate c .

Definition 2 (Representation Score). *For a given approval profile A and committee size k , the representation score of a committee W is defined as:*

$$RP(A, W) = \sum_{i \in N} \min(1, |A_i \cap W|) \quad (2)$$

It may not be possible to maximize both social welfare and representation, so *proportionality* serves as an important third objective to capture a compromise between welfare and representation. It requires that if a large enough voter group collectively approves a shared candidate set, then the group must be “fairly represented”. Definitions of proportionality differ based on how they interpret “fairly represented”.

Definition 3 (T-Cohesive Groups). *Consider an election $E = (C, N)$ with n voters and committee size k . For any integer $T \geq 1$, a group of voters N' is T -cohesive if it contains at least Tn/k voters and collectively approves at least T common candidates, *i.e.* if $|\bigcap_{i \in N'} A_i| \geq T$ and $|N'| \geq Tn/k$.*

Definition 4 (Proportional Justified Representation (PJR)). *A committee W of size k satisfies PJR if for each integer $T \in \{1, \dots, k\}$ and every T -cohesive group $N' \subseteq N$, it holds that $|\bigcup_{i \in N'} A_i \cap W| \geq T$.*

Definition 5 (Extended Justified Representation (EJR)). *A committee W of size k satisfies EJR if for each integer $T \in \{1, \dots, k\}$, every T -cohesive group $N' \subseteq N$ contains at least one voter that approves at least T candidates in W , *i.e.* for some $i \in N'$, $|A_i \cap W| \geq T$.*

Unlike EJR [Aziz *et al.*, 2017], where the focus is on a single group member, PJR provides a more natural requirement for group representation. However, EJR provides stronger guarantees for average voter satisfaction and implies PJR [Sánchez-Fernández *et al.*, 2017].

2.2 Multi-winner Voting Rules

In this section, we introduce the set of approval-based multi-winner voting rules that form the basis of our analysis.

Approval Voting (AV): For an approval profile A , the AV-score of committee W is $sc_{av}(A, W) = \sum_{c \in W} V(c)$. The AV rule is defined as $R_{AV}(A, k) = \arg \max_{W \in S_k(C)} sc_{av}(A, W)$. This rule selects k candidates with the highest individual approval scores.

Approval Chamberlin-Courant (CC): The CC rule [Chamberlin and Courant, 1983], $R_{CC}(A, k)$, picks committees that

maximize representation score $RP(A, W)$. Given profile A , $R_{CC}(A, k) = \arg \max_{W \in S_k(C)} RP(A, W)$. It maximizes voter coverage by maximizing the number of voters with at least one approved candidate in the winning committee.

Proportional Approval Voting (PAV): [Thiele, 1895] For profile A and committee W , the PAV-score is defined as $sc_{pav}(A, W) = \sum_{i \in N} h(|W \cap A_i|)$, where $h(t) = \sum_{i=1}^t 1/i$. The PAV rule is defined as $R_{PAV}(A, k) = \arg \max_{W \in S_k(C)} sc_{pav}(A, W)$. Based on the idea of diminishing returns, a voter’s utility from having an approved candidate in the elected committee W decreases according to the harmonic function $h(t)$. It is a variation of the AV rule that ensures proportional representation, as it guarantees EJR [Aziz *et al.*, 2017]. PAV is the same as AV when committee size $k = 1$, but computing PAV is NP-hard [Aziz *et al.*, 2014].

Method-of-Equal-Shares (MES): $R_{MES}(A, k)$, also known in the literature as Rule-X [Peters *et al.*, 2021; Peters and Skowron, 2020], is an iterative process that uses the idea of budgets to guarantee proportionality. Each voter starts with a budget of k/n and each candidate is of unit cost. In round t , a candidate c is added to W if it is q -affordable, *i.e.* for some $q \geq 0$, $\sum_{i \in N(c)} \min(q, b_i(t)) \geq 1$, where $b_i(t)$ is the budget of voter i in round t . If a candidate is successfully added then the budget of each supporting voter is reduced accordingly. This process continues until either k candidates are added to the committee or it fails (then another voting rule is used to select the remaining candidates).

We elect to study these rules since they exhibit a wide range of properties, allowing for comparisons to be drawn across several axes. First, AV is known to maximize social welfare under certain conditions on voters’ utility functions [Lackner and Skowron, 2020; Lackner and Skowron, 2018], however, there are no guarantees that AV satisfies proportionality (EJR criterion) [Aziz *et al.*, 2017]. On the other hand, CC maximizes representation, but its welfare properties are less well understood. Both PAV and MES guarantee EJR and maintain a balance between representation and social welfare. Thus, this collection of multi-winner voting rules covers the set of properties we are interested in better understanding.

2.3 Objectives

We compare the above voting rules according to different standard objectives [Lackner and Skowron, 2020]. In particular, we consider objectives across three dimensions: welfare, representation, and proportionality.

Utilitarian Ratio: This ratio compares the (utilitarian) social welfare achieved by $W_R = R(A, k)$ to the maximum social welfare achievable:

$$UR(R) = \frac{SW(A, W_R)}{\max_{W \in S_k(C)} SW(A, W)} \quad (3)$$

Representation Ratio: This ratio measures the diversity of the committee $W_R = R(A, k)$, by comparing the representation score achieved by W_R to the optimal representation score amongst all k -sized committees:

$$RR(R) = \frac{RP(A, W_R)}{\max_{W \in S_k(C)} RP(A, W)}. \quad (4)$$

Note that $RR(R_{CC}) = 1$.

Utility-Representation Aggregate Score: This score captures how well a voting rule, $R(A, k)$ balances both social welfare and representation:

$$URagg(R) = UR(R) * RR(R) \quad (5)$$

Finally, we are interested in experimentally verifying whether or not the generated profile instances **satisfy EJR** or **PJR**. To this end, we count the number of profile instances that satisfy these two properties.

3 Deliberation Models

We describe how we model deliberation processes. We first define our underlying agent population and how we model their initial preferences. We then discuss the deliberation process, through which agents exchange information and update their preferences. Finally, we observe that deliberation is often done, not at the full population level, but instead in smaller subgroups. We discuss different ways these deliberation subgroups can be created.

3.1 Voting Population: Preferences and Utilities

Our agent population N is divided into two sets — a *majority* and *minority*, where the number of agents in the majority is greater than that in the minority. Agents’ initial preferences depend on their population group. Consistent with previous work [Lackner and Skowron, 2020], our preference model is based on the ordinal Mallows model. The rankings are then converted to an approval ballot using the top-ranked candidates. In particular, we assume an agent i ’s initial preference ranking, P_i^0 , is sampled from a Mallows model [Mallows, 1957], with reference rankings, Π_{maj} and Π_{min} , for the majority and minority populations respectively.²

We further assume that agents have underlying cardinal utilities for candidates, consistent with their ordinal preferences. These are represented by a vector $U_i = \langle u_i(c_1), u_i(c_2), \dots, u_i(c_m) \rangle$, where $u_i(c_x) \geq u_i(c_y)$ if and only if $c_x \succeq_i c_y$ in P_i^0 , and $u_i(c_x) \in [0, 1]$. We work in this cardinal space as it allows us to leverage standard deliberation models and measure welfare across voting rules in settings where voters derive some utility from elected candidates who were not on their ballot. Our particular instantiation of utility functions subsumes earlier work (e.g. [Lackner and Skowron, 2020]) and is consistent with utility models used in the social choice literature (e.g. [Procaccia and Rosenschein, 2006; Fain *et al.*, 2018]).

3.2 The Deliberation Process

Deliberation is defined as a “discussion in which individuals are amenable to scrutinizing and changing their preferences in the light of persuasion (but not manipulation, deception or coercion) from other participants” [Dryzek and List, 2003].

²The Mallows model is a standard noise model for preferences. It defines a probability distribution over rankings over alternatives (*i.e.* preferences), defined as $\mathbb{P}(r) = \frac{1}{Z} \phi^{d(r, \Pi)}$ where Π is a reference ranking, $d(r, \Pi)$ is the Kendall-tau distance between r and Π , and Z is a normalizing factor.

Deliberation thus requires a group of peers with whom to deliberate and a methodology for changing preferences. In this section, we describe the process in which agents update their preferences, deferring details about peer groups until later.³

Deliberation is an iterative process, involving, at each step, a speaker and listeners. The speaker makes a report, based on their preferences, and the listeners update their own preferences based on this information. In this work, we use a variation of the Bounded Confidence (BC) model to capture the (abstract) deliberation process [Hegselmann and Krause, 2002]. The BC model is a particularly good match for modelling deliberation in groups because it was intended to “describe formal meetings, where there is an effective interaction involving many people at the same time” [Castellano *et al.*, 2009]. In the BC model, listeners consider the speaker’s report (e.g. utilities for different candidates), and update their opinions/preferences of the candidates independently, only if the speaker’s report is not “too far” from their own. The notion of distance is captured by a confidence parameter for each listener, Δ_i , where agents may have different confidence levels [Lorenz, 2007; Weisbuch *et al.*, 2002]. We make a simplifying assumption that agents’ utilities for all m candidates are independent of each other, and apply the BC model to each dimension (candidate) independently. The interested reader is referred to Appendix A⁴ for more details on the original BC model.

Given time step t , some agent, x , selected as the speaker, makes its report (which reveals x ’s thoughts and utilities for the candidates). Each listener updates its own preferences across candidates $c_j \in C$ according to the following rule

$$u_i^{t+1}(c_j) = \begin{cases} (1 - w_{ix})u_i^t(c_j) + w_{ix}u_x^t(c_j), & \text{if } |u_i^t(c_j) - u_x^t(c_j)| \leq \Delta_i \\ u_i^t(c_j), & \text{otherwise} \end{cases} \quad (6)$$

where $w_{ix} \in [0, 1]$ is the *influence weight* that i places on x ’s perspective. It is known that opinions from sources similar to oneself have a higher influence than opinions from dissimilar sources [Wilder, 1990; Mackie *et al.*, 1990]. To capture this phenomenon, we let w_{ix} take on one of two values, contingent on the relationship between i and x . If i and x are both members of the majority group (N_{maj}) or the minority group (N_{min}) then $w_{ix} = \alpha_i$, otherwise $w_{ix} = \beta_i$ where $\alpha_i \geq \beta_i$.

3.3 Deliberation Groups

In the real world, deliberation typically happens in small discussion or peer groups [Elstub *et al.*, 2022; Gözl, 2022]. To this end, we divide the agent population into g sub-groups of approximately equal size. The deliberation process is conducted within these sub-groups where one *round* of deliberation is complete when all agents in each group have had the opportunity to speak.

We want to explore how group-formation strategies influence the deliberation process and the final decision made through voting. Our strategies are informed by common

³As is common in much of the deliberation literature (e.g [Dryzek and List, 2003; Perote-Peña and Piggins, 2015]), we assume agents are non-strategic and truthfully reveal their utilities.

⁴All appendices can be found in the longer version of the paper [Mehra *et al.*, 2023].

heuristics or rationale used in practice and none rely on private/unknown information such as the agents’ underlying utilities or preferences. We do, however, assume that whether an agent is a member of the majority or minority group is public information and allow group-formation strategies to use such information. Finally, we consider both single-round and iterative group-formation strategies where agents are divided into different groups in each round [Elstub *et al.*, 2022].

Single-Round Group-Formation Strategies

Homogeneous group: Each group contains only agents who are members of N_{maj} or N_{min} . That is, there is no mixing of minority and majority agents.

Heterogeneous group: Each group is selected such that the ratio of the number of majority agents to the number of minority agents within the group is approximately equal to the majority:minority ratio in the overall population. Each group created through this strategy is *diverse* and representative of the overall agent population. This strategy is already popular among practitioners in the real world. Citizens’ Assembly of Scotland diversifies deliberation groups based on age, gender, and political affiliation [Gözl, 2022].

Random group: Each group is created by randomly sampling agents from the population (without replacement) with equal probability.

Large group: This is a special case where the deliberation process runs over the entire population of agents. Considering time constraints, limited attention spans, and other physical limitations, such a strategy is not typically used in practice. However, we include this strategy as a Utopian baseline because it ensures maximum exposure to the preferences of every other agent in the system.

Iterative Group-Formation Strategies

Iterative random: In each round, agents are randomly assigned to groups.

Iterative golfer: This strategy is a variant of the social golfer problem [Harvey, 2002; Liu *et al.*, 2019] from combinatorial optimization. The number of rounds, R , is fixed *a priori*, and the number of times any pair of agents meet more than once is minimized. We refer the reader to Appendix B for details. A similar approach is used in Sortition Foundation’s GroupSelect algorithm [Verpoort, 2020], which is used by several nonprofits for group-formation in PB sessions [Gözl, 2022].

4 Experimental Setup

Our election setup consists of 50 candidates ($|C| = 50$)⁵ and 100 voters, with 80 agents in the majority group (N_{maj}) and 20 in the minority group (N_{min}). Agents’ initial preferences are sampled using a Mallows model, with $\phi = 0.2$. Reference rankings, Π_{maj} and Π_{min} , are sampled uniformly from all linear orders over C . Due to this sampling process, agents in either the majority or minority group have fairly similar

⁵Typically, project proposals are invited from the participants in PB [Cabannes, 2004; Aziz and Shah, 2021]. So, there are a large number of candidate projects to choose from (e.g., PB instances in Warsaw, Poland had between 20-100 projects (36 on average). [Stolicki *et al.*, 2020; Fairstein *et al.*, 2022]).

preferences (as ϕ is relatively small) but the two groups themselves are distinct. To instantiate agents’ utility functions, we generate m samples independently from the uniform distribution $U(0, 1)$, sort it, and then map the utilities to the candidates according to the agent’s preference ranking. For the BC model, all three parameters ($\Delta_i, \alpha_i, \beta_i$) are sampled from uniform distributions over the full range for each parameter.⁶

When deliberating, agents are divided into 10 groups (except for the *large group* strategy). This is similar to the Citizens’ Assembly of Scotland, which ran over 16 sessions; in each session, the 104 participants were divided across 12 tables [Gözl, 2022]. For iterative deliberation, the deliberation continues for $R = 5$ rounds.

We consider different approval-based multi-winner voting rules to elect $k = 5$ winners. We use a flexible ballot size, such that each agent’s ballot is of size b_i , where b_i is sampled from $\mathcal{N}(2k, 1.0)$. Agent i ’s approval vote is then the set consisting of its b_i top-ranked candidates from its preference ranking P_i .

As a baseline, we apply every voting rule to the agent preferences *before* deliberation. We then run the different deliberation strategies, freezing agents’ utilities once deliberation has concluded. We then apply every voting rule to the updated preferences. We use the Python library *abcvoting* [Lackner *et al.*, 2021], and use random tie-breaking when a voting rule returns multiple winning committees.⁷

To avoid trivial profiles, *i.e.*, profiles where an almost perfect compromise between welfare and representation is easily achievable, we impose some eligibility conditions. An initial approval profile A^0 is eligible only if $RR(AV, A^0) < 0.9 \wedge UR(CC, A^0) < 0.9$. This is a common technique used in simulations comparing voting rules based on synthetic datasets [Lackner and Skowron, 2020].

This entire simulation is repeated 10,000 times and the average values are reported. To determine statistical significance while comparing any two sets of results, we used both the t -test and Wilcoxon signed-rank test, and we found the p -values to be roughly similar. All pairs of comparisons between deliberation group strategies for a given voting rule are statistically significant ($p < 0.05$) unless otherwise noted.

5 Results

Impact of Deliberation on Preferences: To understand how deliberation processes shape and change agents’ preferences, we compare the average variance in the agents’ utilities before (*initial*) and after deliberation (Figure 1). As expected, deliberation reduces disagreement amongst agents, moving all towards a consensus. Processes where agents are exposed to more, diverse, agents (e.g. the iterative variants and the *large group*) see the largest reduction in variance across the population. We also measure consensus as the average number of common approvals between majority and minority voters and observe the same trend (see Appendix E). While

⁶We ran experiments where all parameters were drawn from a normal distribution. There were no significant differences from the results reported here.

⁷The code for the experiments is available at: <https://github.com/kanav-mehra/deliberation-voting>.

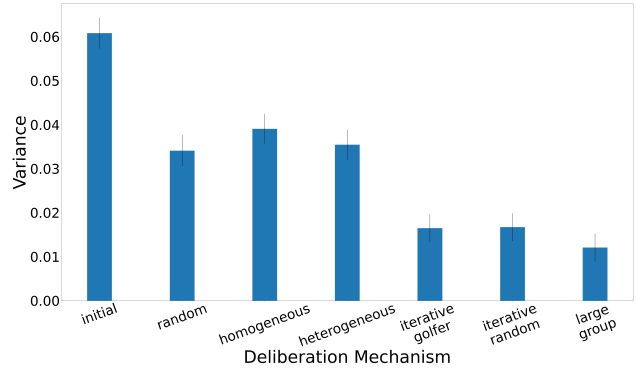


Figure 1: Average variance of agents’ utilities for candidates. Lower variance implies a higher degree of consensus in the population.




Figure 2: Utilitarian ratio across deliberation mechanisms.

achieving greater agreement is desirable, it should not be achieved by disregarding initial minority opinions. We delve into this topic in Section 6.

Utilitarian Ratio: Figure 2 reports the impact of deliberation on utilitarian social welfare. First, we compare the voting rules where there is no deliberation (see *initial* case denoted by the blue bars). AV achieves the highest utilitarian ratio, *i.e.* the utilitarian social welfare provided by AV is closest to the optimal social welfare. Both proportional rules (MES and PAV) are similar and obtain utilitarian ratios that are only slightly lower than AV. Finally, CC performs the worst in terms of welfare. In general, our results match the trends reported in previous work [Lackner and Skowron, 2020].

We now address our main point of interest – the effect of deliberation. As seen in Figure 2, deliberation improves social welfare over the *initial* baseline (blue). In single-round deliberation, both *random* (green) and *heterogeneous* (red) methods show similar results and outperform *homogeneous* (orange) for AV, MES, and PAV. Iterative deliberation exhibits further improvement for these rules. It is worth noting that *iterative golfer* (purple) and *iterative random* (pink) perform similarly and match the *large group* benchmark (brown). For CC, social welfare is always improved with deliberation, however, iterative deliberation is not as powerful. This is due to the nature of the rule (explained in detail in Appendix D).




Figure 3: Representation ratio across deliberation mechanisms.

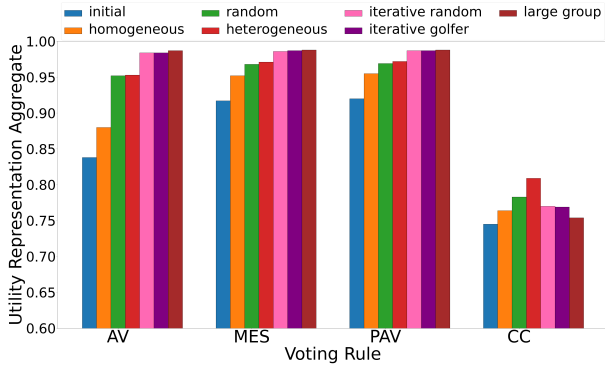


Figure 4: Utility-Representation aggregate score across deliberation mechanisms.

Representation Ratio: Figure 3 shows the average representation ratio of the voting rules across different deliberation mechanisms. Since CC optimizes for diversity by design, $RR(CC) = 1.0$. Under no deliberation, AV has the lowest representation ratio. Since AV simply picks the candidates with the highest approval scores, it does not care about minority preferences. The proportional rules (MES and PAV), however, achieve much higher representation as they are designed to maintain a balance between welfare and diversity.

The effect of deliberation is more pronounced here compared to the utilitarian ratio results, particularly for AV. Within the single-round mechanisms, *homogeneous* achieves a slight improvement over the *initial* setup for all rules. However, specifically for AV, both *heterogeneous* and *random* achieve much higher representation over both *initial* and *homogeneous* setups. Again, both iterative mechanisms achieve further improvements compared to the single-round setups and almost match the *large group* benchmark.

Utility-Representation Aggregate Score: Figure 4 shows the average results for this objective. Under no deliberation (*initial* baseline), we see that the proportional rules (MES and PAV) perform the best, followed by AV, and then CC. This is consistent with earlier findings since the proportional rules are designed with this goal in mind. Both AV and CC perform poorly on this metric, pre-deliberation, since they do well on

Deliberation Strategy	EJR%		PJR%	
	AV	CC	AV	CC
Initial (no deliberation)	99.5	62.5	99.5	73.4
Homogeneous	96.4	69.9	96.4	75.1
Random	100	81.9	100	85.6
Heterogeneous	100	92.7	100	94.0
Iterative Random	100	31.4	100	53.6
Iterative Golfer	100	29.9	100	51.2
Large Group	100	6.10	100	23.4

Table 1: EJR- and PJR-satisfaction (AV and CC).

either welfare (AV) or representation (CC) but not both.

We observe a positive effect from deliberation and achieve a significant performance improvement over the *initial* baseline. Within the single-round mechanisms, *heterogeneous* and *random* perform similarly (except for CC where *heterogeneous* is better) and outperform the *homogeneous* setup. Iterative deliberation leads to further improvement as both iterative methods match the performance of the *large group*.

EJR and PJR Satisfaction: Table 1 shows the percentage of EJR- and PJR-satisfying committees returned by AV and CC. We focus only on AV and CC since the proportional rules MES and PAV guarantee EJR. Even under no deliberation (*initial*), AV satisfies EJR in almost all profiles, which further improves to perfect satisfaction with deliberation (except *homogeneous*). This is interesting since AV is not guaranteed to satisfy EJR.⁸ EJR and PJR satisfaction for CC also improves if single-round deliberation is supported, with *heterogeneous* achieving the best result. Iterative deliberation, however, does not perform well. We believe that this arises due to CC’s strong focus on representation (see Appendix D).

6 Discussion

Deliberation changes the quality of the outcomes produced by different multi-winner voting rules. We explore these observations in more detail.

Single-Round Deliberation: Even a single round of deliberation improved outcomes across all voting rules and all objectives. However, the choice of the deliberation structure was also important. For all objectives, *random* and *heterogeneous* consistently outperformed *homogeneous*. We hypothesize that this improvement was due to these deliberation strategies maximizing exposure to diverse opinions. Under *homogeneous* deliberation, the population sub-groups become more inwardly focused, leading to the formation of distinct *T*-cohesive groups. This was particularly problematic when used with AV, which picks candidates with the highest approval support and fails to ‘fairly’ represent the cohesive minority agents in some cases, thereby failing EJR (Table 1). By allowing majority and minority agents to interact, there

⁸Since the minority and majority agents have highly correlated approval sets, *T*-cohesive groups may exist only for a small set of minority- and majority-supported candidates, thereby making the EJR requirement easy to satisfy. Furthermore, previous research [Fairstein *et al.*, 2022; Bredereck *et al.*, 2019] shows that under many natural preference distributions (generated elections), there are many EJR-satisfying committees.

Deliberation Strategy	Minority Opinion Preservation
Initial (no deliberation)	0
Homogeneous	0.20
Random	0.30
Heterogeneous	0.48
Iterative Random	0.65
Iterative Golfer	0.66
Large Group	0.92

Table 2: Minority Opinion Preservation: average number of initial (pre-deliberation) *minority-supported* candidates selected by AV in the final committee after deliberation.

was an opportunity for minority agents to influence the majority population. This translated to higher welfare, representation, and proportionality guarantees (Figure 4 and Table 1). **Iterative Deliberation:** In comparison to single-round methods, iterative deliberation further supports consensus (Figure 1) and improves all objectives for most voting rules. Furthermore, there was no statistical difference between the *iterative golfer* and *iterative random* methods. We view this as a positive result with practical design implications. While care does need to be taken in determining group sizes, a simple, computationally inexpensive mechanism is as effective as one that is more complex.

The exception to the observation is the CC rule. CC’s strong focus on representation and coverage makes it unsuitable for deliberation methods that drive higher degrees of consensus (such as iterative methods and *large group*) since it fails to represent population groups proportionally. We refer the interested reader to Appendix D.

Minority Opinion Preservation: While we have been extolling deliberation, there are caveats. In particular, it is important to ensure that deliberation processes are inclusive and encourage minority participation [Gherghina *et al.*, 2021]. Care must be taken to ensure that when moving toward consensus, initial minority preferences are not ignored. While consensus would imply better voting outcomes, it could come at the cost of ignoring minority opinions. We measure whether this is a concern in our experiments by studying whether minority-supported candidates were selected by AV under different deliberation mechanisms.⁹

A candidate is either *minority-supported* or *majority-supported* based on the *initial* approval profile. We say that a candidate *c* is *minority-supported* if (pre-deliberation) the fraction of minority voters who include *c* in their approval ballot is greater than the fraction of majority voters who include *c* in their approval ballot. Table 2 reports the average number of pre-deliberation (initial) *minority-supported* candidates selected by AV (post-deliberation) across deliberation strategies. This serves as an indicator of whether minority preferences are *preserved*.

In the *initial* setup (no deliberation), AV does not elect any *minority-supported* candidates. However, this improves as agents interact and deliberate with the broader population. Note that since the minority agents have similar preferences, and they constitute 20% of the population in our setup, a pro-

⁹This is not a concern for other rules since they are designed to achieve proportionality (MES and PAV) or diversity (CC).

Approval Voting	MES (initial) (0.917)	PAV (initial) (0.92)
Initial (0.838)	0.913	0.910
Homogeneous (0.88)	0.959	0.956
Random (0.952)	1.038	1.034
Heterogeneous (0.953)	1.039	1.035
Iterative Random (0.984)	1.073	1.069
Iterative Golfer (0.984)	1.073	1.069

Table 3: Average utility-representation aggregate score obtained by AV under different deliberation setups in comparison to the proportional rules under no deliberation.

portional committee would represent them with 1 (out of 5) candidate. As seen in Table 2, the *large group* setup comes close to the ideal outcome on average. Thus, with deliberation, AV can *preserve* and represent minority preferences.

“Simple” vs. “Complex” Voting Rules: We argue that the “complexity” of a voting rule can be measured along three axes. First, one can ask about the computational complexity of computing a winning outcome or committee (e.g. PAV is known to be NP-hard [Aziz *et al.*, 2014], whereas AV is polynomial). Second, there is growing work in better understanding the ramifications of ballot design and voting rules on the cognitive load of voters [Benadè *et al.*, 2021]. Finally, there is value in using simple explainable voting rules. Explainability engenders trust in the system (which in turn may impact engagement in participatory democratic processes).

While two of these dimensions are, somewhat subjective, we argue that AV can be viewed as being *simple* across all three, whereas PAV and MES are *complex* along at least one dimension. Our hypothesis is that *simple* rules coupled with deliberation processes can do as well as more *complex* voting rules. To this end, we compare AV with deliberation to MES and PAV without deliberation, using the utility-representation aggregate score ($URagg(R)$) as our measure (Table 3). Values greater than 1.0 indicate that AV with the corresponding deliberation mechanism achieves a better $URagg$ score than MES/PAV without deliberation. These findings support our argument that “simple” rules coupled with effective deliberation strategies can be as effective as the “complex” rules.

7 Conclusion

We presented an empirical study of the relationship between deliberation and voting rules in approval-based multi-winner elections. Deliberation generally improves voting outcomes with respect to welfare, representation, and proportionality guarantees. Effectively designed mechanisms that increase exposure to diverse groups and opinions enhance the quality of deliberation, protect minority preferences, and in turn, achieve better outcomes. Importantly, we show that in the presence of effective deliberation, ‘simpler’ voting rules such as AV can be as powerful as more ‘complex’ rules without deliberation. Our hope is that our findings can further support the design of effective citizen-focused democratic processes.

References

- [Aziz and Shah, 2021] Haris Aziz and Nisarg Shah. Participatory budgeting: Models and approaches. In *Pathways Between Social Science and Computational Social Science*, pages 215–236. Springer, 2021.
- [Aziz et al., 2014] Haris Aziz, Serge Gaspers, Joachim Gudmundsson, Simon Mackenzie, Nicholas Mattei, and Toby Walsh. Computational aspects of multi-winner approval voting. In *Workshops at the Twenty-Eighth AAAI Conference on Artificial Intelligence*, 2014.
- [Aziz et al., 2017] Haris Aziz, Markus Brill, Vincent Conitzer, Edith Elkind, Rupert Freeman, and Toby Walsh. Justified representation in approval-based committee voting. *Social Choice and Welfare*, 48(2):461–485, 2017.
- [Aziz et al., 2018] Haris Aziz, Edith Elkind, Shenwei Huang, Martin Lackner, Luis Sánchez-Fernández, and Piotr Skowron. On the complexity of extended and proportional justified representation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- [Benadè et al., 2021] Gerdus Benadè, Swaprava Nath, Ariel D Procaccia, and Nisarg Shah. Preference elicitation for participatory budgeting. *Management Science*, 67(5):2813–2827, 2021.
- [Brams and Fishburn, 2007] Steven Brams and Peter C Fishburn. *Approval voting*. Springer Science & Business Media, 2007.
- [Bredereck et al., 2019] Robert Bredereck, Piotr Faliszewski, Andrzej Kaczmarczyk, and Rolf Niedermeier. An experimental view on committees providing justified representation. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence, IJCAI’19*, page 109–115. AAAI Press, 2019.
- [Brill et al., 2017] Markus Brill, Rupert Freeman, Svante Janson, and Martin Lackner. Phragmén’s voting methods and justified representation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31, 2017.
- [Cabannes, 2004] Yves Cabannes. Participatory budgeting: a significant contribution to participatory democracy. *Environment and Urbanization*, 16(1):27–46, 2004.
- [Castellano et al., 2009] Claudio Castellano, Santo Fortunato, and Vittorio Loreto. Statistical physics of social dynamics. *Rev. Mod. Phys.*, 81:591–646, May 2009.
- [Chamberlin and Courant, 1983] John R Chamberlin and Paul N Courant. Representative deliberations and representative decisions: Proportional representation and the borda rule. *American Political Science Review*, 77(3):718–733, 1983.
- [Chung and Duggan, 2020] Hun Chung and John Duggan. A formal theory of democratic deliberation. *American Political Science Review*, 114(1):14–35, 2020.
- [Dryzek and List, 2003] John S. Dryzek and Christian List. Social choice theory and deliberative democracy: A reconciliation. *British Journal of Political Science*, 33(1):1–28, 2003.
- [Elkind et al., 2021] Edith Elkind, Davide Grossi, Ehud Shapiro, and Nimrod Talmon. United for change: deliberative coalition formation to change the status quo. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 5339–5346, 2021.
- [Elstub et al., 2022] Stephen Elstub, Oliver Escobar, Ailsa Henderson, Tamara Thorne, Nick Bland, and Evelyn Bowes. *Citizens’ Assembly of Scotland: Research Report*. Scottish Government Social Research, January 2022.
- [Fain et al., 2017] Brandon Fain, Ashish Goel, Kamesh Munagala, and Sukolsak Sakshuwong. Sequential deliberation for social choice. In *International Conference on Web and Internet Economics*, pages 177–190. Springer, 2017.
- [Fain et al., 2018] Brandon Fain, Kamesh Munagala, and Nisarg Shah. Fair allocation of indivisible public goods. In *Proceedings of the 2018 ACM Conference on Economics and Computation, EC ’18*, page 575–592, New York, NY, USA, 2018. Association for Computing Machinery.
- [Fairstein et al., 2022] Roy Fairstein, Dan Vilenchik, Reshef Meir, and Kobi Gal. Welfare vs. representation in participatory budgeting. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems, AAMAS ’22*, page 409–417, Richland, SC, 2022. International Foundation for Autonomous Agents and Multiagent Systems.
- [Faliszewski et al., 2017] Piotr Faliszewski, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. Multiwinner voting: A new challenge for social choice theory. *Trends in Computational Social Choice*, 74(2017):27–47, 2017.
- [Fishkin, 2009] James Fishkin. *When the people speak: Deliberative democracy and public consultation*. Oxford University Press, 2009.
- [Gherghina et al., 2021] Sergiu Gherghina, Monika Mokre, and Sergiu Miscoiu. Introduction: Democratic deliberation and under-represented groups. *Political Studies Review*, 19(2):159–163, 2021.
- [Goel and Lee, 2016] Ashish Goel and David T Lee. Towards large-scale deliberative decision-making: Small groups and the importance of triads. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 287–303, 2016.
- [Gölz, 2022] Paul Gölz. *Social Choice for Social Good: Proposals for Democratic Innovation from Computer Science*. PhD thesis, Carnegie Mellon University, 2022.
- [Habermas, 1996] Jürgen Habermas. *Between Facts and Norms: Contributions to a Discourse Theory of Law and Democracy*. Polity, 1996.
- [Harvey, 2002] Warwick Harvey. CSPLib problem 010: Social golfers problem. <http://www.csplib.org/Problems/prob010>, 2002. Accessed: 2022-12-14.
- [Hegselmann and Krause, 2002] Rainer Hegselmann and Ulrich Krause. Opinion Dynamics and Bounded Confidence Models, Analysis and Simulation. *Journal of Artificial Societies and Social Simulation*, 5(3):1–2, 2002.

- [Lackner and Skowron, 2018] Martin Lackner and Piotr Skowron. Consistent approval-based multi-winner rules. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 47–48, 2018.
- [Lackner and Skowron, 2020] Martin Lackner and Piotr Skowron. Utilitarian welfare and representation guarantees of approval-based multiwinner rules. *Artificial Intelligence*, 288:103366, 2020.
- [Lackner et al., 2021] Martin Lackner, Peter Regner, Benjamin Krenn, Elvi Cela, Jonas Kompauer, Florian Lackner, Stanisław Szufa, and Stefan Schlomo Forster. *abcvoting*: A Python library of approval-based committee voting rules, 2021. Current version: <https://github.com/martinlackner/abcvoting>.
- [Liu et al., 2019] Ke Liu, Sven Löffler, and Petra Hofstedt. Social golfer problem revisited. In Jaap van den Herik, Ana Paula Rocha, and Luc Steels, editors, *Agents and Artificial Intelligence*, pages 72–99, Cham, 2019. Springer International Publishing.
- [Lorenz, 2007] Jan Lorenz. Continuous opinion dynamics under bounded confidence: A survey. *International Journal of Modern Physics C*, 18(12):1819–1838, 2007.
- [Mackie et al., 1990] D M Mackie, L T Worth, and A G Asuncion. Processing of persuasive in-group messages. *J Pers Soc Psychol*, 58(5):812–822, May 1990.
- [Mallows, 1957] C. L. Mallows. Non-null Ranking Models. I. *Biometrika*, 44(1-2):114–130, 06 1957.
- [Mehra et al., 2023] Kanav Mehra, Nanda Kishore Sreenivas, and Kate Larson. Deliberation and voting in approval-based multi-winner elections. *arXiv preprint arXiv: 2305.08970*, 2023.
- [Perote-Peña and Piggins, 2015] Juan Perote-Peña and Ashley Piggins. A model of deliberative and aggregative democracy. *Economics & Philosophy*, 31(1):93–121, 2015.
- [Peters and Skowron, 2020] Dominik Peters and Piotr Skowron. Proportionality and the limits of welfarism. In *Proceedings of the 21st ACM Conference on Economics and Computation*, EC ’20, page 793–794, New York, NY, USA, 2020. Association for Computing Machinery. Extended version at <https://arxiv.org/abs/1911.11747>.
- [Peters et al., 2021] Dominik Peters, Grzegorz Pierczyński, and Piotr Skowron. Proportional participatory budgeting with additive utilities. *Advances in Neural Information Processing Systems*, 34:12726–12737, 2021.
- [Procaccia and Rosenschein, 2006] Ariel D Procaccia and Jeffrey S Rosenschein. The distortion of cardinal preferences in voting. In *International Workshop on Cooperative Information Agents*, pages 317–331. Springer, 2006.
- [Rad and Roy, 2021] Soroush Rafiee Rad and Olivier Roy. Deliberation, single-peakedness, and coherent aggregation. *American Political Science Review*, 115(2):629–648, 2021.
- [Sánchez-Fernández et al., 2017] Luis Sánchez-Fernández, Edith Elkind, Martin Lackner, Norberto Fernández, Jesús Fisteus, Pablo Basanta Val, and Piotr Skowron. Proportional justified representation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31, 2017.
- [Skowron et al., 2016] Piotr Skowron, Piotr Faliszewski, and Jérôme Lang. Finding a collective set of items: From proportional multirepresentation to group recommendation. *Artificial Intelligence*, 241:191–216, 2016.
- [Skowron et al., 2017] Piotr Skowron, Martin Lackner, Markus Brill, Dominik Peters, and Edith Elkind. Proportional rankings. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, pages 409–415, 2017.
- [Skowron, 2021] Piotr Skowron. Proportionality degree of multiwinner rules. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 820–840, 2021.
- [Stolicki et al., 2020] Dariusz Stolicki, Stanisław Szufa, and Nimrod Talmon. Pabulib: A participatory budgeting library. *arXiv preprint arXiv:2012.06539*, 2020.
- [Thiele, 1895] Thorvald N Thiele. Om flerfoldevalg. *Over-sigt over det Kongelige Danske Videnskabernes Selskabs Forhandlinger*, 1895:415–441, 1895.
- [Verpoort, 2020] Philipp Verpoort. GroupSelect by Sortition Foundation. <https://github.com/sortitionfoundation/groupselect-app>, 2020. Accessed: 2022-12-12.
- [Weisbuch et al., 2002] Gérard Weisbuch, Guillaume Defuant, Frédéric Amblard, and Jean-Pierre Nadal. Meet, discuss, and segregate! *Complexity*, 7(3):55–63, 2002.
- [Wilder, 1990] D A. Wilder. Some determinants of the persuasive power of in-groups and out-groups: Organization of information and attribution of independence. *Journal of Personality and Social Psychology*, 59(6):1202–1213, 1990.
- [Zvi et al., 2021] Gil Ben Zvi, Eyal Leizerovich, and Nimrod Talmon. Iterative deliberation via metric aggregation. In *International Conference on Algorithmic Decision Theory*, pages 162–176. Springer, 2021.