Selecting Interlacing Committees

Chris Dong
Technical University of Munich
Munich, Germany
chris.dong@tum.de

Martin Bullinger
University of Oxford
Oxford, United Kingdom
martin.bullinger@cs.ox.ac.uk

Tomasz Wąs
University of Oxford
Oxford, United Kingdom
tomasz.was@cs.ox.ac.uk

Larry Birnbaum
Northwestern University
Evanston, United States
l-birnbaum@northwestern.edu

Edith Elkind
Northwestern University
Evanston, United States
edith.elkind@northwestern.edu

ABSTRACT

Polarization is a major concern for a well-functioning society. Often, mass polarization of a society is driven by polarizing political representation, even when the latter is easily preventable. The existing computational social choice methods for the task of committee selection are not designed to address this issue. We enrich the standard approach to committee selection by defining two quantitative measures that evaluate how well a given committee interconnects the voters. Maximizing these measures aims at avoiding polarizing committees. While the corresponding maximization problems are NP-complete in general, we obtain efficient algorithms for profiles in the voter-candidate interval domain. Moreover, we analyze the compatibility of our goals with other representation objectives, such as excellence, diversity, and proportionality. We identify trade-offs between approximation guarantees, and describe algorithms that achieve simultaneous constant-factor approximations.

KEYWORDS

Computational Social Choice; Approval-Based Committee Voting; Polarization

ACM Reference Format:
Chris Dong, Martin Bullinger, Tomasz Wąs, Larry Birnbaum, and Edith Elkind. 2025. Selecting Interlacing Committees. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Detroit, Michigan, USA, May 19 – 23, 2025*, IFAAMAS, 9 pages.

1 INTRODUCTION

In recent years, the increasing prevalence of polarization has been a major concern, discussed not just by social scientists, but by society at large, and accompanied by extensive media coverage [20, 27]. Polarization is commonly defined as the division of a group into clusters of completely different opinions or ideologies. It is a major concern for the modern society, which has to work towards a consensus when resolving global challenges, such as fighting poverty, climate change, or pandemics (see [28] and the references therein).

Importantly, polarization can occur as a phenomenon concerning an entire society or only at the level of political representation, e.g., when considering the distribution of opinions among the delegates in a parliament. The former is often referred to as *mass polarization*, while the latter is known as *elite polarization*, see, e.g., [1].

Academic literature broadly agrees that the phenomenon of elite polarization is on the rise. For example, when depicting the members of the US Congress in terms of their ideology on a scale ranging from the most liberal to the most conservative, one can observe a significant shift when comparing the 87th Congress in the 1960s and the 111th Congress around 2010, see Figure 2.1 in the book by Fiorina [19]. However, whether the society as a whole is polarized as well is less clear. Fiorina et al. [20] argue that there is no conclusive evidence for mass polarization, even when considering highly sensitive topics such as abortion. For instance, they provide evidence that the elite polarization among delegates is already much higher than the polarization among party identifiers [20, Table 2.1]. They argue that the media play an important role in creating an inaccurate picture of mass polarization [20]. Indeed, the media can have a significant effect on the perception of and conclusions drawn from elite polarization [27].

This view is opposed by Abramowitz and Saunders [1], who analyze data from the American National Election Studies. They provide extensive evidence that mass polarization has increased signficiantly since the 1970s. Moreover, their results suggest mass polarization based on geography (i.e., different ideologies across US states) or religious beliefs.

Against this background, we aim to offer a novel perspective on the intertwined phenomena of mass polarization at the broad level of a society as a whole and elite polarization at the level of the society's political, parliamentary representation. We highlight how an election can lead to a parliament that is far more polarized than the society it represents, and we propose quantitative measures that evaluate a set of representatives according to how well it interlaces the electorate. We believe that our ideas can be developed to prevent societies with broadly moderate opinions being represented by unnecessarily polarized parliaments.

We approach polarized democratic representation through the lens of social choice theory. In this line of research, parliamentary elections have been conceptualized as so-called multiwinner voting rules. Their formal study, especially in the approval-based setting, in which each voter's ballot specifies a set of approved candidates, has received extensive attention in recent years [18, 26].

Example 1.1. As a motivating example, consider the voting scenario illustrated in Figure 1. There are four voters, indicated by the

This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

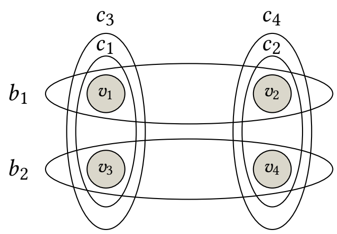

Figure 1: A preference profile with four voters $v_1, \ldots, v_4$ is depicted as hypergraph, where the voters are nodes and the candidates b_i, c_j are hyperedges connecting the voters approving them. In this profile, typical multiwinner voting rules do not distinguish between selecting $\{c_1, c_2, c_3, c_4\}$ and $\{c_1, c_2, b_1, b_2\}$.

gray circles, as well as six candidates. Each candidate is represented by an ellipse that covers the voters approving this candidate. For instance, candidate b_1 is approved by voters v_1 and v_2, whereas candidates c_1 and c_3 are both approved by the same set of voters, namely v_1 and v_3. In practice, this is likely to happen when c_1 and c_3 represent very similar ideologies.

Assume that we want to select a committee consisting of 4 candidates. Two reasonable choices would be to select $W = \{c_1, c_2, c_3, c_4\}$ or $W' = \{c_1, c_2, b_1, b_2\}$. Both selections lead to committees in which each voter approves exactly two selected candidates. Moreover, multiwinner voting rules typically considered in the literature, such as Thiele rules and their sequential variants [33], Phragmén's rule [32], or the more recently introduced method of equal shares [31], do not distinguish between these two choices. There is, however, a difference. While W divides the electorate into two perfectly separated subsets of voters, W' connects all voters. From the perspective of polarization, W looks polarizing while W' bridges all voters. Thus, we need novel voting rules that can tease out this distinction. In our paper, we aim to provide a principled approach that favors committees in the spirit W'.[1] ◁

We define two simple objectives that aim to measure how well a committee interlaces the voters. First, we consider maximizing the number of *pairs* of voters approving a common candidate (the Pairs objective). While optimizing this objective leads to the selection of W' in Example 1.1, it can still result in voters being split into large disconnected clusters (cf. Example 3.1). The reason is that Pairs only counts direct, but not indirect links. Hence, as a second objective we count the number of pairs of voters that are *connected* by a sequence of candidates (the Cons objective).

While both objectives immediately give rise to voting rules—select a committee that maximizes Pairs or Cons—we primarily view them as measures of polarization. Whenever they are high, polarization in the selected committee is low. Thus, we investigate the feasibility of maximizing our objectives, both on their own and in combination with the goals of diversity and proportionality.

[1] Of course, while we try to highlight the phenomenon at hand with a simple example, our construction extends to elections with many voters or candidates: e.g., each voter in the example might represent a quarter of a large electorate.

We first consider the computational problem of maximizing Pairs or Cons in isolation (Section 4). Unfortunately, for unrestricted preferences this problem is NP-hard. However, we obtain a polynomial-time algorithm for the structured domain of voter-candidate interval (VCI) preferences [22], where voters and candidates are represented by intervals on the real line and a voter approves a candidate if and only if their intervals intersect. Such preferences are reasonable in parliamentary elections where candidates can often be ordered on a left-right spectrum and voters approve candidates that are close to them on this spectrum.

In Section 5, we investigate whether one can select interlacing committees while achieving other desiderata. We first consider *excellence*, as measured by the *approval voting* (AV) score, i.e., the total number of approvals received by committee members. There is a straightforward way to obtain what is essentially an α-approximation of the Pairs objective together with an $(1-\alpha)$-approximation of the AV score: one can simply use an α-fraction of the committee for the former and an $(1-\alpha)$-fraction for the latter. Unfortunately, it turns out that this simple algorithm is essentially optimal: We prove that if a voting rule provides an α-approximation of the Pairs objective and a β-approximation of the AV score, then necessarily $\alpha + \beta \leq 1$. Next, we look at *diversity*, as captured by the *Chamberlin–Courant* (CC) score, which is the number of voters who approve at least one candidate in the committee. The CC score is closely related to the Pairs objective: the former measures the coverage of voters, while the latter measures the coverage of pairs of voters. Hence, it is quite surprising that we obtain the same trade-off as for Pairs and AV. Further, we study the compatibility with *proportionality*, as captured by the extended justified representation axiom (EJR). Again, we show the same tight trade-off: If a voting rule provides an α-approximation of the Pairs objective and a β-approximation of EJR, then $\alpha + \beta \leq 1$.

It is more challenging to combine the Cons objective with AV, CC, EJR, or even Pairs. This is due to an interesting qualitative difference between Pairs and Cons. While a constant fraction of the best candidates achieves a constant approximation of Pairs, for Cons this is not the case. Hence, we obtain worse trade-offs: If a voting rule provides an α^2-approximation of Cons and a β-approximation of AV, CC, EJR, or Pairs, then $\alpha + \beta \leq 1$. Note that since $\alpha < 1$, it holds that $\alpha^2 < \alpha$. Hence, for instance, $\alpha^2 = \frac{1}{3}$ and $\beta = \frac{1}{2}$ is already impossible. Moreover, for Cons and AV specifically, the trade-off that we obtain is even more subtle, which suggests that finding a matching lower bound might be challenging. Nevertheless, we make first steps towards this goal, by showing that under suitable domain restrictions there always exists a committee that achieves a $\frac{1}{4}$-approximation of Cons and a $\frac{1}{2}$-approximation of AV, CC, EJR, or Pairs, which matches our upper bound.

2 RELATED WORK

In the existing literature, multiwinner voting rules usually aim to guarantee the selection of the best candidates based on their individual quality [3, 14], representation of diverse opinions [8, 15], or proportional treatment of different groups of interests [29, 31–33]. An overview of the most common approval-based multiwinner voting rules is given in the book by Lackner and Skowron [26]. To the best of our knowledge, no rules were proposed so far with the explicit goal of reducing polarization or connecting voters.

A line of research in multiwinner voting looks at the possibility of combining various objectives as well as their inherent trade-offs, similar to our study in Section 5. Lackner and Skowron [25] provide worst-case bounds on AV and CC scores of committees output by popular voting rules. For ordinal preferences, Kocot et al. [23] analyze the complexity of finding committees that offer an optimal combination of approximations of two objectives. Moreover, a series of works look at AV and CC scores that can be guaranteed by committees that satisfy proportionality axioms [7, 13, 17].

A number of authors study the relationship between an electoral system (or, more narrowly, a voting rule) and the way the candidates choose to strategically place themselves on the political spectrum [5, 10, 24, 30]. Such an analysis can indicate whether a rule prevents, or reinforces, polarization. Our approach differs in that we analyze the direct effect of a voting rule on the polarization caused by a chosen committee, while the aforementioned works analyze how preferences evolve based on a given rule.

Delemazure et al. [11] pursue a goal that can be seen as opposite to ours: selecting a most polarizing committee of size 2; they focus on ordinal preferences. In a similar vein, Colley et al. [9] proposed measures of how *divisive*, or polarizing, a single candidate is.

3 MODEL

We start by introducing key notation and proposing two ways of measuring how well a committee interconnects the voters. For a positive integer $k \in \mathbb{N}$, define $[k] := \{1, \dots, k\}$.

3.1 Approval-Based Multiwinner Voting

We consider the standard setting of approval-based multiwinner voting [26]. Given a set of m candidates C, an *election instance* $\mathcal{E} = (V, A, k)$ consists of a set of n voters V, an approval profile $A = (A_v)_{v \in V}$ with $A_v \subseteq C$ for all $v \in V$, and a target committee size $k \in [m]$. Intuitively, a voter $v \in V$ approves precisely the candidates in A_v. Throughout the paper, we view a profile A as a hypergraph with vertex set V, and, for each $c \in C$, a hyperedge $V_c = \{v \in V : c \in A_v\}$. In the remainder of this section, we consider an election instance $\mathcal{E} = (V, A, k)$ over a candidate set C.

Besides the general setting, we also consider structured domains of spatial one-dimensional preferences. Specifically, we consider elections where all voters and all candidates can be mapped to intervals on the real line so that a voter approves a candidate if and only if their respective intervals intersect. Formally, following Godziszewski et al. [22], we say that an election (V, A, k) belongs to the *voter-candidate interval (VCI) domain* if there exist a collection of positions $\{x_c\}_{c \in C} \cup \{x_v\}_{v \in V} \subseteq \mathbb{R}$ and a collection of nonnegative radii $\{r_c\}_{c \in C} \cup \{r_v\}_{v \in V} \subseteq \mathbb{R}^+ \cup \{0\}$ such that for all $v \in V$, $c \in C$ it holds that $c \in A_v$ if and only if $|x_c - x_v| \le r_c + r_v$.

The VCI domain is the most general domain of one-dimensional approval preferences considered in the literature. In particular, it generalizes the voter interval (VI) and candidate interval (CI) domains, defined as follows [16]. An election belongs to the *voter interval (VI) domain* if there is an ordering of the voters $v_1, \dots, v_n$ such that each candidate is approved by some interval of this ordering, i.e., for each $c \in C$ there exist $i, j \in [n]$ such that $V_c = \{v_i, \dots, v_j\}$. Similarly, an election belongs to the *candidate interval (CI) domain* if there is an ordering of the candidates $c_1, \dots, c_m$ such that for each $v \in V$ there exist $i, j \in [m]$ such that $A_v = \{c_i, \dots, c_j\}$. It is easy to see that the VI and CI domains are contained in the VCI domain.[2]

A *feasible committee* for an instance (V, A, k) is a subset $W \subseteq C$ with $|W| = k$. A *(multiwinner) voting rule* f takes as input an instance (V, A, k) and outputs a feasible committee $f(V, A, k)$.

3.2 Classic Committee Selection

A popular classification of multiwinner voting rules is in terms of the main objective in electing the committee, with three most commonly studied objectives being *excellence*, *diversity*, and *proportionality* [18].

Both excellence and diversity are defined quantitatively: each of these objectives is associated with a function that assigns a numerical score to each feasible committee, with higher scores associated with better performance. Formally, given an instance $\mathcal{E} = (V, A, k)$ and a feasible committee W, we define

$$\mathrm{AV}(W, \mathcal{E}) := \sum_{v \in V} |A_v \cap W|,$$
$$\mathrm{CC}(W, \mathcal{E}) := |\{v \in V : A_v \cap W \neq \varnothing\}|.$$

For both objectives (as well as for the two novel objectives defined in Section 3.3) we omit $\mathcal{E}$ from the notation when it is clear from the context. The quantities AV and CC are referred to as, respectively, the *approval score* and the *Chamberlin–Courant score* of committee W in election $\mathcal{E}$. Intuitively, AV counts the number of approvals received by the members of W and is viewed as a measure of excellence, while CC counts the number of voters represented by W, i.e., voters who approve at least one member of W, and is viewed as a measure of diversity. The voting rule that outputs a committee maximizing AV (respectively, CC) is known as the *approval voting rule* (respectively, the *Chamberlin–Courant rule*[3]).

Consider a function S that assigns scores to feasible committees in a given election (e.g., S = AV or S = CC). Given $\alpha \in [0, 1]$, we say that a committee W^* *satisfies α-S* for an election $\mathcal{E} = (V, A, k)$ if

$$S(W^*, \mathcal{E}) \ge \alpha \cdot \max_{\substack{W \subseteq C, \\ |W| = k}} S(W, \mathcal{E}).$$

Moreover, we say that a voting rule f *satisfies α-S* if for every election $\mathcal{E}$ it holds that $f(\mathcal{E})$ satisfies α-S for $\mathcal{E}$. For instance, the Chamberlin–Courant rule satisfies 1-CC.

In contrast, proportionality is typically captured by representation axioms. A prominent axiom of this type is *extended justified representation* (EJR) [2]; intuitively, it states that sufficiently large groups of voters with similar preferences should be appropriately represented in the selected committee. We will now define what it means for a committee to satisfy approximate EJR.

Given an election (V, A, k) over C and $\alpha \in (0, 1]$, a committee $W \subseteq C$ is said to *satisfy α-EJR* if for every $\ell \in [k]$ and every subset

[2]For instance, given an election $\mathcal{E} = (V, A, k)$ in VI, as witnessed by voter ordering $v_1, \dots, v_n$, we can set $x_{v_i} = i$ and $r_{v_i} = 0$ for each $i \in [n]$. To position the candidates, for each $c \in C$ we compute $c^- = \min\{i : c \in A_{v_i}\}$ and $c^+ = \max\{i : c \in A_{v_i}\}$ and set $x_c = (c^- + c^+)/2$, $r_c = (c^+ - c^-)/2$. Clearly, these positions and radii certify that $\mathcal{E}$ belongs to the VCI domain. For CI, the construction is similar.

[3]Originally, Chamberlin and Courant [8] proposed their rule for linear preferences. However, the approval variant of this rule is commonly studied in the computational social choice literature.

$S \subseteq V$ such that $\alpha \cdot |S| \geq \frac{\ell}{k} \cdot |V|$ and $|\bigcap_{i \in S} A_i| \geq \ell$ there exists at least one voter $i \in S$ such that $|W \cap A_i| \geq \ell$. We say that a rule f *satisfies* α*-EJR* if for every election $\mathcal{E}$ it holds that $f(\mathcal{E})$ satisfies α-EJR. By setting α to 1, we obtain the standard EJR axiom.

3.3 Interlacing Committee Selection

We now define two new objectives, which assess committees based on how well they interlace voters.

Our first objective is the number of *pairs* of voters that jointly approve a selected candidate. Given an election $\mathcal{E} = (V, A, k)$, let $V^{(2)} := \{\{u, v\} \subseteq V : u \neq v\}$ be the set of all voter pairs. We set

$$\textsc{Pairs}(W, \mathcal{E}) := |\{\{u, v\} \in V^{(2)} : A_u \cap A_v \cap W \neq \varnothing\}|.$$

Note that for every instance $\mathcal{E} = (V, A, k)$ one can define the *associated pair instance* $\mathcal{E}^{(2)} = (V^{(2)}, A^{(2)}, k)$, where $A^{(2)}_{\{u,v\}} = A_u \cap A_v$ for every $\{u, v\} \in V^{(2)}$. For each instance $\mathcal{E}$ and committee $W \subseteq C$ we have $\textsc{Pairs}(W, \mathcal{E}) = \mathrm{CC}(W, \mathcal{E}^{(2)})$.

While the Pairs objective only considers direct links between voters, our second objective takes into account indirect connections as well. Given an instance $\mathcal{E} = (A, V, k)$ and a subset of candidates $W \subseteq C$, we say that two voters $u, v \in V$ are *connected by* W (and write $u \sim_W v$) if there is a sequence of voters $u = v_0, v_1, \ldots, v_s = v$ with $A_{v_{i-1}} \cap A_{v_i} \cap W \neq \varnothing$ for every $i \in [s]$. To evaluate a committee W, we count pairs of voters connected by W. Formally,

$$\textsc{Cons}(W, \mathcal{E}) := \left|\{\{u, v\} \in V^{(2)} : u \sim_W v\}\right|.$$

Since both Pairs and Cons assign scores to committees, we also consider their approximate versions, i.e., α-Pairs and α-Cons.

Our interest in Cons is motivated by the following example.

Example 3.1. Consider a profile with six voters $v_1, \ldots, v_6$, six *cycle* candidates $c_1, \ldots, c_6$, and two *diagonal* candidates d_1 and d_2, whose hypergraph is depicted in Figure 2. Each cycle candidate is approved by two consecutive voters: for $i = 1 \ldots, 5$ candidate c_i is approved by v_i and v_{i+1}, while c_6 is approved by v_1 and v_6. Also, d_1 is approved by v_2 and v_6 and d_2 by v_3 and v_5. Let $k = 6$.

Consider two committees: $W = \{c_1, c_2, c_3, c_4, c_5, c_6\}$ contains all cycle candidates, whereas in $W' = \{c_1, c_3, c_4, c_6, d_1, d_6\}$ two cycle candidates are exchanged for the diagonal candidates (W' is shown in red in Figure 2). Common voting rules, including the approval voting rule and the Chamberlin–Courant rule, do not distinguish between W and W', as each voter approves exactly two candidates in either committee. Moreover, the rule that maximizes Pairs is also unable to distinguish them, as both W and W' cover exactly 6 pairs of voters. However, intuitively, W' seems more polarizing: under W', there are two disconnected groups of voters, each supporting (though not fully) their own set of candidates.

In contrast, a rule that maximizes Cons is sensitive to the differences between the two committees. Under W, all 15 pairs of voters are connected, while W' only achieves 6 connections. ◁

4 COMPUTATION OF THE NEW OBJECTIVES

In this section, we show that maximizing Pairs and Cons is NP-hard in general, but tractable on well-structured domains. All proofs missing from this and subsequent sections can be found in the full version of our paper.

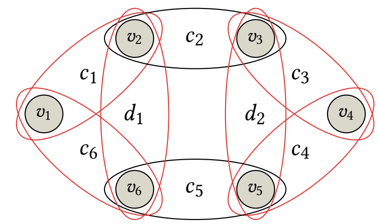

Figure 2: Illustration of Example 3.1.

4.1 General Preferences

Our hardness proofs are based on the NP-complete problem Exact Cover by 3-Sets (X3C) [21]. The proof idea for the Pairs objective is to represent every element in the ground set of an X3C instance by a pair of voters.

Theorem 4.1. *It is* NP*-complete to decide, given an election* $\mathcal{E} = (V, A, k)$ *and a threshold* $q \in \mathbb{N}$*, whether there exists a committee* W *of size at most* k *such that* $\textsc{Pairs}(W, \mathcal{E}) \geq q$.

A similar hardness result holds for Cons. The proof idea is to introduce an auxiliary voter that is the focal point in connecting all voters.

Theorem 4.2. *It is* NP*-complete to decide, given an election* $\mathcal{E} = (V, A, k)$ *and a threshold* $q \in \mathbb{N}$*, whether there exists a committee* W *of size* k *such that* $\textsc{Cons}(W, \mathcal{E}) \geq q$*. The hardness result holds even if* $q = \binom{n}{2}$*, i.e., if the goal is to connect all* n *voters.*

4.2 One-dimensional Preferences

We will now complement our hardness results by arguing that these problems can be solved in polynomial time on the VCI domain. We start by observing that, for the objectives we consider, a VCI instance can be transformed into a CI instance without changing the value of these objectives. To this end, we define a notion of dominance among candidates and prove that, in the absence of dominated candidates, every VCI instance is a CI instance.

4.2.1 From VCI to CI. Given an election $\mathcal{E} = (V, A, k)$ over a candidate set C, we say that candidate $c' \in C$ is *dominated* by a candidate $c \in C$ if every voter approving c' also approves c, and some voter approves c but not c', i.e., $V_{c'}$ is a proper subset of V_c.

It turns out that if an election in the VCI domain contains no dominated candidates, it belongs to the much simpler to analyze CI domain; this observation, which is implicit in the work of Elkind et al. [13, Lemma 4.7], may be of independent interest. Indeed, removal of dominated candidates from a winning committee does not affect the Pairs and Cons objectives, so we can simply remove all dominated candidates from the input instance.

Proposition 4.3. *Let* $\mathcal{E}$ *be an instance in the VCI domain. If* $\mathcal{E}$ *contains no dominated candidates, then it belongs to the CI domain.*

In what follows, we state our results for the VCI domain, but assume that the input election belongs to the CI domain, and we are explicitly given the respective candidate order. It will also be convenient to assume that this order is $c_1, \ldots, c_m$. This requires two preprocessing steps: first, we eliminate all dominated candidates (which, by Proposition 4.3, results in a CI election), and second, we compute an ordering of the candidates witnessing that our instance belongs to the CI domain. Both steps can be implemented in polynomial time (for the second step, see, e.g., [16]).

4.2.2 Efficient Algorithms. We are ready to present polynomial-time algorithms for Pairs and Cons. Since Pairs is identical to CC on the associated pair instance, we can compute Pairs by leveraging an existing algorithm for CC in the CI domain [4, 16].

Proposition 4.4. *In the VCI domain, a committee that maximizes Pairs can be computed in polynomial time.*

In the VCI domain, we can also compute a committee that maximizes Cons in polynomial time; however, the argument is significantly more complicated. Again, we assume that the input profile belongs to the CI domain, as witnessed by the candidate ordering $c_1, \dots, c_m$. A natural idea, then, is to use dynamic programming to compute, for each $b \in [k]$ and $i \in [m]$, an optimal subcommittee of size b with rightmost candidate c_i. For $b = 1$, the computation is straightforward, and for $b = k$, one of the resulting m committees globally maximizes Cons. However, computing the value of adding c_i to a committee of size $b - 1$ that has c_j as its rightmost candidate is a challenging task: this is because the number of connections that c_i adds depends on the size of the connected component associated with c_j. To handle this, we add a third dimension to the dynamic program: the number of voters $x \in [n]$ in the connected component of the last selected candidate. The resulting dynamic program has $\mathcal{O}(mnk)$ cells, and each cell can be filled in polynomial time given the values of the already-filled-in cells.

Theorem 4.5. *In the VCI domain, a committee that maximizes Cons can be computed in polynomial time.*

5 COMBINING OBJECTIVES

While interlacing objectives can be viewed in isolation, in many cases, standard objectives of excellence, diversity, or proportionality continue to be important for the selection of a committee. In this section, we investigate to what extent we can select committees that simultaneously perform well with respect to both interlacing and standard objectives.

5.1 Pairs Objective

First, we consider combining the Pairs objective with individual excellence of the committee members, as measured by AV. For every $\alpha \in [0, 1]$ and every election $\mathcal{E} = (V, A, k)$, there is a simple way to obtain a simultaneous $\lceil \alpha k \rceil / k$-approximation of Pairs and $\lfloor (1 - \alpha)k \rfloor / k$-approximation of AV. Indeed, we can split the k positions on the committee into two parts of size $k_1 = \lceil \alpha k \rceil$ and $k_2 = \lfloor (1 - \alpha)k \rfloor$, respectively, and then select k_1 candidates so as to maximize Pairs and k_2 candidates so as to maximize AV (if some candidate is selected both times, we replace their second copy by an arbitrary unselected candidate). Since the marginal gain for Pairs and AV objectives from each additional candidate is non-increasing, this procedure provides the desired guarantees. Note that Lackner and Skowron [25] propose a similar method for combining AV and CC.

Proposition 5.1. *For every $\alpha \in [0, 1]$ and election $\mathcal{E}$, there exists a committee that satisfies $\lceil \alpha k \rceil / k$-Pairs and $\lfloor (1 - \alpha)k \rfloor / k$-AV.*

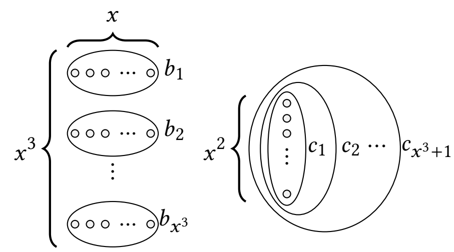

Figure 3: Illustration of the profile constructed in the proof of Proposition 5.3. The block voters are on the left and the central voters are on the right. Each block candidate is approved by x block voters, whereas each central candidate is approved by all central voters.

We can use the same technique to combine Pairs with the goal of diverse representation, as measured by CC.

Proposition 5.2. *For every $\alpha \in [0, 1]$ and election $\mathcal{E}$, there exists a committee that satisfies $\lceil \alpha k \rceil / k$-Pairs and $\lfloor (1 - \alpha)k \rfloor / k$-CC.*

Perhaps surprisingly, it turns out that, for both combinations, this is the best we can hope for.

Proposition 5.3. *For every $\alpha, \beta \in [0, 1]$, if a voting rule satisfies α-Pairs and β-AV, then $\alpha + \beta \leq 1$.*

Proof. We will construct a family of instances that allows us to bound the sum of approximation ratios. For a given constant $x \in \mathbb{N}$, consider the election $\mathcal{E} = (V, A, k)$ defined as follows (see Figure 3 for an illustration). The set C consists of x^3 *block candidates* $b_1, \dots, b_{x^3}$, and $x^3 + 1$ *central candidates* $c_1, \dots, c_{x^3+1}$, so that $|C| = 2x^3 + 1$. The set V consist of x^3 groups of *block voters* $(V_i)_{i \in [x^3]}$ of size x each, and a single group of x^2 central voters. For $i \in [x^3]$ each voter in block V_i approves candidate b_i only, whereas each central voter approves all central candidates $c_1, \dots, c_{x^3+1}$. The target committee size is set to $k = x^3 + 1$.

Since there are fewer than k block candidates, every committee W contains at least one central candidate, who is approved by all x^2 central voters, and covers all $(x^2 - 1)x^2/2$ pairs of central voters. Then, every additional central candidate contributes x^2 to the AV objective and 0 to the Pairs objective, whereas every additional block candidate contributes x to the AV objective and $x(x - 1)/2$ to the Pairs objective.

Assume that for some $\gamma \in \{0, 1/x^3, 2/x^3, \dots, 1\}$ our rule selects a committee $W_\gamma \subseteq C$ with $\gamma x^3 + 1$ central candidates and $(1 - \gamma)x^3$ block candidates. Then, we obtain the following AV and Pairs scores:

$$\mathrm{AV}(W_\gamma, \mathcal{E}) = (\gamma x^3 + 1)x^2 + (1 - \gamma)x^4 = \gamma x^5 + \mathcal{O}(x^4), \text{ and}$$

$$\textsc{Pairs}(W_\gamma, \mathcal{E}) = \frac{x^4 - x^2}{2} + (1 - \gamma)\frac{x^5 - x^4}{2} = \frac{1 - \gamma}{2} x^5 + \mathcal{O}(x^4),$$

where the $\mathcal{O}(\cdot)$ upper bound holds for all values of γ. Observe that the maximum AV score is obtained when we take $\gamma = 1$, and the maximum Pairs score is obtained when $\gamma = 0$. Also,

$$\frac{\mathrm{AV}(W_\gamma, \mathcal{E})}{\mathrm{AV}(W_1, \mathcal{E})} + \frac{\textsc{Pairs}(W_\gamma, \mathcal{E})}{\textsc{Pairs}(W_0, \mathcal{E})} \leq 1 + \mathcal{O}(1/x).$$

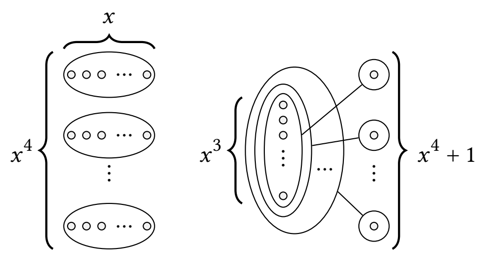

Figure 4: An illustration of the profile constructed in the proof of Proposition 5.4. Block voters are on the left, central voters in the middle, and arm voters on the right. Each block candidate is approved by x block voters, whereas each central candidate is approved by all central voters and one arm voter.

Hence, as x tends to infinity, the sum of approximation ratios for AV and Pairs becomes arbitrarily close to 1. □

One may expect the Pairs and CC objectives to be more aligned than Pairs and AV. Indeed, CC and Pairs count the number of voters (resp., pairs of voters) that (jointly) approve at least one candidate in the selected committee. However, surprisingly, the worst-case trade-off for these objectives is the same as for Pairs and AV.

Proposition 5.4. *For every $\alpha, \beta \in [0, 1]$, if a voting rule satisfies α-Pairs and β-CC, then $\alpha + \beta \le 1$.*

Proof sketch. The proof is similar to that of Proposition 5.3. This time, block candidates are used to achieve a large CC-score, while central candidates are used to achieve a large Pairs-score, see Figure 4.

We increase the number of block candidates from x^3 to x^4; as before, each block candidate is approved by a distinct block of x voters. Also, we increase the number of central candidates to $x^4 + 1$ and the number of central voters from x^2 to x^3; as before, all central voters approve all central candidates. Further, we add $x^4 + 1$ *arm voters* $a_1, \ldots, a_{x^4+1}$; each a_i approves the central candidate c_i. Finally, we set the target committee size to $k = x^4 + 1$.

Just as in the proof of Proposition 5.3, a winning committee contains $\gamma x^4 + 1$ central candidates and $(1 - \gamma)x^4$ block candidates for some $\gamma \in \{0, 1/x^4, \ldots, 1\}$. Then, by a similar analysis, we obtain

$$\frac{\text{CC}(W_\gamma, \mathcal{E})}{\text{CC}(W_1, \mathcal{E})} + \frac{\text{Pairs}(W_\gamma, \mathcal{E})}{\text{Pairs}(W_0, \mathcal{E})} \le 1 + \mathcal{O}(1/x),$$

which concludes the proof. □

Finally, we investigate how to combine the Pairs objective with proportional representation, as captured by the EJR axiom. Again, we can use the committee-splitting technique to show that for every election $\mathcal{E} = (V, A, k)$ there is a committee that satisfies $\lceil \alpha k \rceil / k$-Pairs and $(1 - \alpha)$-EJR. For this, we first need to show that we can guarantee $(1 - \alpha)$-EJR with a $(1 - \alpha)$-fraction of the committee seats. To this end, we employ a variant of the *method of equal shares* (MES) [31]. Briefly, this rule gives each voter k/n units of money; it then sequentially selects candidates that are best for voters that still have money, and subtracts money from the supporters of the selected candidates (see the full version of our paper for the formal definition). By adapting the proof that MES satisfies EJR (Peters and Skowron [31]), we show that, by executing MES while scaling the voters' budgets by α, we obtain α-EJR for the original instance; we believe that this result is of independent interest.[4]

Lemma 5.5. *Let $\alpha \le 1$ be given. For every election $\mathcal{E} = (V, A, k)$, executing MES on $(V, A, \alpha k)$ returns a committee of size $\lfloor \alpha k \rfloor$ satisfying α-EJR in polynomial time.*

We remark that the committee obtained as described in Lemma 5.5 satisfies an even stronger notion of proportionality, namely, α-EJR+ [6]. Using Lemma 5.5, we now easily obtain the desired guarantees.

Proposition 5.6. *For every $\alpha \in [0, 1]$ and election $\mathcal{E}$, there exists a committee that satisfies α-Pairs and $(1 - \alpha)$-EJR.*

Proof. Consider an election $\mathcal{E}$. By Lemma 5.5, we can satisfy $(1-\alpha)$-EJR using $\lfloor (1-\alpha)k \rfloor$ candidates. With the remaining $k - \lfloor (1-\alpha)k \rfloor = \lceil \alpha k \rceil$ candidates, we can guarantee α-CC on the associated pair instance $\mathcal{E}^{(2)}$. This is equivalent to satisfying α-Pairs on $\mathcal{E}$, concluding the proof. □

As before, we provide a matching upper bound. We note that our proof works even if, instead of EJR, we consider the much weaker axiom of *justified representation* (JR) [2].

Proposition 5.7. *For every $\alpha, \beta \in [0, 1]$, if a voting rule satisfies α-Pairs and β-EJR, then $\alpha + \beta \le 1$.*

To conclude this section, we note that the guarantees offered by Propositions 5.1, 5.2, and 5.6 are established by combining algorithms for the two objectives in question. Some of these objectives (in particular, CC and Pairs) do not admit polynomial-time algorithms unless P=NP. To achieve polynomial runtime, we can plug in greedy approximation algorithms for CC and Pairs (for Pairs, we use the sequential Chamberlin–Courant rule on the associated pair instance); however, this comes at the expense of a factor of $(1 - 1/e)$ in the approximation guarantees [25].

5.2 Cons Objective

An important reason why we obtained good approximations of Pairs, AV, and CC was that these objectives are subadditive, i.e., for every two committees W and W', the value for committee $W \cup W'$ is never larger than the sum of the values for W and W'. As a consequence, these objectives are sublinear with respect to the committee size, in the sense that if we only use an α-fraction of the k committee seats, we can obtain at least an α-fraction of the original value for a committee of size k (up to rounding).

In contrast, the Cons objective is not subadditive, so we cannot use the same technique. In fact, the following result shows that the trade-off between Cons and any of AV, CC, or Pairs is strictly worse (on the side of the Cons) than the trade-offs we have established in Section 5.1. Notably, our upper bound applies even to instances that belong to the VI domain.

Proposition 5.8. *For every $\alpha, \beta \in [0, 1]$, if a voting rule satisfies α^2-Cons and β-AV, β-CC, or β-Pairs, then $\alpha + \beta \le 1$. This already holds in the VI domain.*

[4] A similar observation was made by Dong and Peters [12], but requires $\lceil (1 - \alpha)k \rceil$ seats, which in our case would allow only for a rounded-down Pairs guarantee.

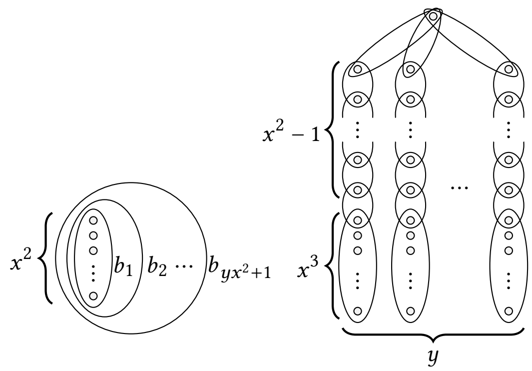

Figure 5: An illustration of the profile constructed in the proof of Proposition 5.10.

Proof sketch. For the proof of all three statements, consider an instance with x^3 blocks, with each block consisting of x voters approving the corresponding block candidate. Further, we have $x^3 + 1$ central voters ordered on a line, with each pair of adjacent central voters approving a designated central candidate.

This instance belongs to the VI domain, as we can place the voters within each block on the line, followed by the central voters. The remainder of the proof consists of two parts. In the first part, we show that, to satisfy β-Pairs, β-CC, or β-AV, we require at least $\beta x^3 - \mathcal{O}(x^2)$ block candidates. In the second part, we show that, with the remaining candidates, we can obtain at most a $(1-\beta)^2$-approximation of Cons. □

We derive a similar result for EJR by reducing the number of blocks from x^3 to slightly fewer than βx^3.

Proposition 5.9. *For every $\alpha, \beta \in [0,1]$, if a voting rule satisfies α^2-Cons and β-EJR, then $\alpha + \beta \le 1$. This already holds in the VI domain.*

However, in some cases the trade-off is even worse than the one in Propositions 5.8 and 5.9. Consider a stepwise function $s\colon [0,1] \to [0,1]$ given by $s(\alpha) = 1/(\lceil 2/\alpha \rceil - 1)$; see Figure 6. Intuitively, it finds the smallest $p \in \mathbb{N}$ such that $\alpha \ge 2/p$ and returns $1/(p-1)$. We then have the following trade-off between Cons and AV.

Proposition 5.10. *For every $\beta \in [0,1]$, if a voting rule f satisfies β-AVthen it satisfies at most $s(1-\beta)$-Cons.*

Proof. Let $y = \frac{1}{s(1-\beta)} = \lceil \frac{2}{1-\beta} \rceil - 1$; note that y is an integer. For an arbitrary constant $x \in \mathbb{N}$, consider the election $\mathcal{E} = (V, A, k)$ defined as follows (see Figure 5 for an illustration). The set C contains $yx^2 + 1$ *block* candidates $(b_i)_{i \in [yx^2+1]}$, y *arm* candidates $(a_i)_{i \in [y]}$, and yx^2 *chain* candidates $(c_{i,j})_{i \in [x^2], j \in [y]}$. The set V consist of x^2 *block* voters, yx^3 *arm* voters split into y arms $A_1, \dots, A_y$ of size x^3 each, $y(x^2-1)$ *chain* voters $(h_{i,j})_{i \in [x^2-1], j \in [y]}$ split into y arms $H_1, \dots, H_y$ of size $x^2 - 1$ each, and one *central* voter v.

The voters have the following preferences. All block voters approve all block candidates. For each $j \in [y]$, each voter in arm A_j approves the arm candidate a_j, and additionally, exactly one voter in A_j approves the chain candidate $c_{x^2,j}$. For each $i \in [x^2-1]$ and $j \in [y]$, the chain voter $h_{i,j}$ approves the chain candidates $c_{i,j}$ and $c_{i+1,j}$. Finally, the central voter approves the chain candidates $c_{1,j}$ for each $j \in [y]$.

We set the target committee size to $k = y(x^2+1) + 1$. The high-level idea of the proof is that for Cons it is important to connect the arm voters through the selection of chain candidates. However, if we select a β-fraction of block candidates in order to guarantee β-AV, then we cannot connect arm voters from different arms.

Consider a size-k committee W. Note that W contains at least one block candidate, as there are only $y + yx^2 = k - 1$ arm and chain candidates. Moreover, suppose there is an arm candidate a not included in W. Then removing a block candidate from W and adding a instead increases both AV and Cons. Thus, we can assume that W contains all y arm candidates. For each $z \in [yx^2]$, let W_z be a committee that selects y arm candidates, $z+1$ block candidates, and $yx^2 - z$ chain candidates.

Each arm candidate in W_z contributes x^3 to the AV score, each block candidate contributes x^2, and each chain candidate contributes 2. Thus, we obtain

$$\mathrm{AV}(W_z, \mathcal{E}) = yx^3 + zx^2 + x^2 + 2(yx^2 - z).$$

Observe that AV is maximized when $z = yx^2$; thus, for W_z to provide β-AV, it has to be the case that $z \ge \beta yx^2 - \mathcal{O}(x)$.

Now, for Cons, we claim that for large enough x, with the remaining $yx^2 - z$ chain candidates, we cannot connect arm voters from two different arms. Indeed, recall that $y = \frac{1}{s(1-\beta)} = \lceil \frac{2}{1-\beta} \rceil - 1 < \frac{2}{1-\beta}$. Hence, $(1-\beta)y < 2$. Consequently, there is an $\varepsilon > 0$ such that $(1-\beta)y = 2 - \varepsilon$. Since $z \ge \beta yx^2 - \mathcal{O}(x)$, we choose at most $yx^2 - z \le (1-\beta)yx^2 + \mathcal{O}(x) < 2x^2 - \varepsilon x^2 + \mathcal{O}(x)$ chain candidates. Thus, for large enough x we have strictly fewer than $2x^2 - 2$ chain candidates. This proves the claim, as it takes $2(x^2-1)$ chain candidates to connect two arm voters from different arms.

Let us now calculate the value of the Cons objective. The connections among block voters contribute at most $\binom{x^2}{2}$, and the connections among chain voters contribute at most $\binom{yx^2}{2}$. Connections between chain voters and arm voters contribute at most yx^5, as each chain voter can only be connected to arm voters in a single arm. We connect all arm voters within the same arm, but, as argued above, we do not connect arm voter from different arms, so connections among arm voters contribute $y\binom{x^3}{2}$. The central voter can contribute $\mathcal{O}(yx^3)$ connections. In total,

$$\mathrm{Cons}(W_z, \mathcal{E}) = \frac{yx^6}{2} + \mathcal{O}(x^5).$$

On the other hand, the maximum value of Cons is obtained when we select all $yx^2 - y$ chain voters: in this case, all yx^3 arm voters are connected, and Cons is at least

$$\frac{y^2x^6}{2} + \mathcal{O}(x^5).$$

Thus, the fraction of Cons we can obtain while satisfying β-AV converges to $\frac{1}{y} = s(1-\beta)$ for $x \to \infty$. As $s(\alpha)$ is monotonically non-decreasing in α (and hence $\alpha > 1 - \beta$ implies $s(\alpha) \ge s(1-\beta)$), this concludes the proof. □

Consider the two upper bounds that we have obtained for α-approximation of Cons given that a voting rule satisfies β-AV, given by Propositions 5.8 and 5.10 (their plots are presented in Figure 6). Since the upper bounds intersect several times and they are of different nature (stepwise vs. continuous), it seems that establishing tight trade-offs might be a challenging and interesting problem. Similarly, because Cons is not subadditive, finding a general lower bound for these trade-offs seems highly non-trivial as well. Nevertheless, we conclude this section with a positive result on guarantees that we can obtain for a combination of Pairs and Cons objectives in the VI domain, which matches our upper bound.

Proposition 5.11. *For every instance (V, A, k) in the VI domain with even k, there exists a committee that satisfies $\frac{1}{4}$-Cons and any one of the criteria $\frac{1}{2}$-AV, $\frac{1}{2}$-CC, $\frac{1}{2}$-EJR, or $\frac{1}{2}$-Pairs.*

Proof sketch. It suffices to split the committee into two halves and focus on satisfying $\frac{1}{4}$-Cons with $\frac{k}{2}$ committee members. With the other $\frac{k}{2}$ slots, we can use the methods in Propositions 5.1 and 5.2 and Lemma 5.5 to obtain $\frac{1}{2}$-AV, $\frac{1}{2}$-CC, $\frac{1}{2}$-EJR, or $\frac{1}{2}$-Pairs.

Let $\mathcal{E}$ be an election with even k that belongs to the VI domain. Take an optimal committee W of size k with respect to Cons. We will choose half of W in the following fashion. In the corresponding hypergraph, W consists of one or more connected components. For components in which we have an even number of candidates, we can show that half of these candidates connect more than half of the voters covered by the whole component. This gives us at least $\frac{1}{4}$ of connections inside the component.

Now, since k is even, there is an even number of components with an odd number of candidates. We arbitrarily group them into pairs. Then, for every pair, we show that it is possible to select half of the candidates rounded up in one of them and half of the candidates rounded down in the other so that in total we cover $\frac{1}{4}$ of the connections from the two components. □

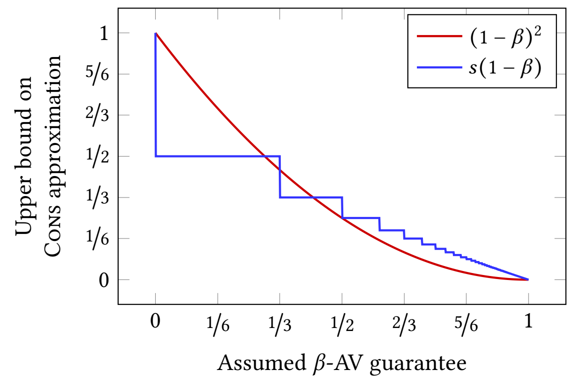

Figure 6: Two different upper bounds on the possible α-approximation of Cons for rules that satisfy β-AV. The $(1-\beta)^2$ upper bound is the result of Proposition 5.8 and $s(1-\beta)$ is implied by Proposition 5.10.

6 CONCLUSION

Our paper sheds new light on the interdependency of mass and elite polarization. We observe that the selection of a representative committee can significantly influence elite polarization independently of mass polarization. With the aim of avoiding polarization at the level of the representation, we have introduced Pairs and Cons, two numerical objectives that measure how well a committee interlaces the electorate.

We show that, while maximizing both objectives is NP-hard, a committee maximizing either of them can be computed in polynomial time on the voter-candidate interval domain. Also, we study the compatibility of our objectives with measures of excellence, diversity, and proportionality. We find approximation trade-offs suggesting that there is nothing better than dividing the committee seats among different objectives and trying to maximize each objective with its designated share of the committee: in the worst case, the synergies are negligible. For almost all objectives we study, a subcommittee yields a fraction of the optimal value proportional to its size. Only for Cons, the dependency is quadratic (or even worse), leading to inferior guarantees.

We believe that our work offers an important perspective that has been missing from the social choice literature on multiwinner voting. As such, it calls for further research; in what follows, we suggest some promising directions.

An immediate open question is to determine the exact trade-off between Cons and other objectives. While we have a bound for α^2-Cons and $(1-\alpha)$-approximations of other objectives, Proposition 5.10 shows that the picture is more nuanced.

Going beyond our base model, another direction is to consider our objectives in the broader context of participatory budgeting (PB), where each candidate has a cost, and the committee needs to stay within a given budget. In this setting, candidates are usually projects, such as a playground, a community garden, or a cycling path. Interlacing voters by projects in PB has an additional interpretation: the funded projects may lead to interaction among the agents who use them (e.g., working together in a community garden). This seems quite desirable in the context of PB, where one of the goals is community building.

Moreover, it would be interesting to explore the compatibility of our objectives and the canonical desiderata in the context of real-life instances: It is plausible that on realistic data one can achieve much better trade-offs than in the worst case.

Finally, while Pairs and Cons offer some insight into the polarization induced by a committee, there are settings where they fail to provide useful information: For example, if some candidate is approved by all voters, any committee containing this candidate maximizes both objectives. Therefore, further insights could be gained by studying refined versions of our objectives; e.g., one can consider the strength of the connections or, in case of the Cons objective, the length of the (shortest) path between a pair of voters.

ACKNOWLEDGMENTS

Most of this work was done when Edith Elkind was at the University of Oxford. Chris Dong was supported by the Deutsche Forschungsgemeinschaft under grants BR 2312/11-2 and BR 2312/12-1. Martin Bullinger was supported by the AI Programme of The Alan Turing Institute. Tomasz Wąs and Edith Elkind were supported by the UK Engineering and Physical Sciences Research Council (EPSRC) under grant EP/X038548/1.

REFERENCES

[1] Alan I. Abramowitz and Kyle L. Saunders. 2008. Is polarization a myth? *The Journal of Politics* 70, 2 (2008), 542–555.

[2] Haris Aziz, Markus Brill, Vincent Conitzer, Edith Elkind, Rupert Freeman, and Toby Walsh. 2017. Justified representation in approval-based committee voting. *Social Choice and Welfare* 48, 2 (2017), 461–485.

[3] Salvador Barberà and Danilo Coelho. 2008. How to choose a non-controversial list with k names. *Social Choice and Welfare* 31, 1 (2008), 79–96.

[4] Nadja Betzler, Arkadii Slinko, and Johannes Uhlmann. 2013. On the computation of fully proportional representation. *Journal of Artificial Intelligence Research* 47 (2013), 475–519.

[5] Damien Bol, Konstantinos Matakos, Orestis Troumpounis, and Dimitrios Xefteris. 2019. Electoral rules, strategic entry and polarization. *Journal of Public Economics* 178 (2019), 104065.

[6] Markus Brill and Jannik Peters. 2023. Robust and verifiable proportionality axioms for multiwinner voting. In *Proceedings of the 24th ACM Conference on Economics and Computation (ACM-EC)*.

[7] Markus Brill and Jannik Peters. 2024. Completing priceable committees: Utilitarian and representation guarantees for proportional multiwinner voting. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence*. 9528–9536.

[8] John R. Chamberlin and Paul N. Courant. 1983. Representative deliberations and representative decisions: Proportional representation and the Borda rule. *American Political Science Review* 77, 3 (1983), 718–733.

[9] Rachael Colley, Umberto Grandi, César Hidalgo, Mariana Macedo, and Carlos Navarrete. 2023. Measuring and controlling divisiveness in rank aggregation. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*. 2616–2623.

[10] Gary W. Cox. 1985. Electoral equilibrium under approval voting. *American Journal of Political Science* (1985), 112–118.

[11] Théo Delemazure, Lukasz Janeczko, Andrzej Kaczmarczyk, and Stanisław Szufa. 2024. Selecting the most conflicting pair of candidates. In *Proceedings of the 33rd International Joint Conference on Artificial Intelligence (IJCAI)*. 2766–2773.

[12] Chris Dong and Jannik Peters. 2024. *Proportional multiwinner voting with dynamic candidate sets*. Technical Report. https://pub.dss.in.tum.de/brandt-research/dynamic_multiwinner.pdf.

[13] Edith Elkind, Piotr Faliszewski, Ayumi Igarashi, Pasin Manurangsi, Ulrike Schmidt-Kraepelin, and Warut Suksompong. 2024. The price of justified representation. *ACM Transactions on Economics and Computation* 12, 3 (2024), 1–27.

[14] Edith Elkind, Piotr Faliszewski, Piotr Skowron, and Arkadii Slinko. 2017. Properties of multiwinner voting rules. *Social Choice and Welfare* 48 (2017), 599–632.

[15] Edith Elkind and Anisse Ismaili. 2015. OWA-based extensions of the Chamberlin–Courant rule. In *Proceedings of the 4th International Conference on Algorithmic Decision Theory (ADT)*. 486–502.

[16] Edith Elkind and Martin Lackner. 2015. Structure in dichotomous preferences. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*. 2019–2025.

[17] Roy Fairstein, Dan Vilenchik, Reshef Meir, and Kobi Gal. 2022. Welfare vs. representation in participatory budgeting. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 409–417.

[18] Piotr Faliszewski, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. 2017. Multiwinner voting: A new challenge for social choice theory. In *Trends in Computational Social Choice*, Ulle Endriss (Ed.). Chapter 2.

[19] Morris P. Fiorina. 2017. *Unstable majorities: Polarization, party sorting, and political stalemate.* Hoover press.

[20] Morris P. Fiorina, Samuel J. Abrams, and Jeremy C. Pope. 2011. *Culture war? The myth of a polarized America* (3rd ed.). Longman.

[21] Michael R. Garey and David S. Johnson. 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness.* W. H. Freeman.

[22] Michał T. Godziszewski, Paweł Batko, Piotr Skowron, and Piotr Faliszewski. 2021. An analysis of approval-based committee rules for 2D-Euclidean elections. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence*. 5448–5455.

[23] Maciej Kocot, Anna Kolonko, Edith Elkind, Piotr Faliszewski, and Nimrod Talmon. 2019. Multigoal committee selection. In *Proceedings of the 28th International Joint Conferences on Artificial Intelligence (IJCAI)*. 385–391.

[24] Anna-Sophie Kurella and Salvatore Barbaro. 2024. *On the polarization premium for radical parties in PR electoral systems.* Technical Report. Gutenberg School of Management and Economics, Johannes Gutenberg-Universität Mainz.

[25] Martin Lackner and Piotr Skowron. 2020. Utilitarian welfare and representation guarantees of approval-based multiwinner rules. *Artificial Intelligence* 288 (2020), 103366.

[26] Martin Lackner and Piotr Skowron. 2023. *Multi-Winner Voting with Approval Preferences*. Springer Nature.

[27] Matthew Levendusky and Neil Malhotra. 2016. Does media coverage of partisan polarization affect political attitudes? *Political Communication* 33, 2 (2016), 283–301.

[28] Simon A Levin, Helen V Milner, and Charles Perrings. 2021. The dynamics of political polarization. *Proceedings of the National Academy of Sciences* 118, 50 (2021), e2116950118.

[29] Burt L. Monroe. 1995. Fully proportional representation. *American Political Science Review* 89, 4 (1995), 925–940.

[30] Roger B. Myerson and Robert J. Weber. 1993. A theory of voting equilibria. *American Political Science Review* 87, 1 (1993), 102–114.

[31] Dominik Peters and Piotr Skowron. 2020. Proportionality and the limits of welfarism. In *Proceedings of the 21nd ACM Conference on Economics and Computation (ACM-EC)*. 793–794.

[32] Edvard Phragmén. 1899. Till frågan om en proportionell valmetod. *Statsvetenskaplig Tidskrift* 2, 2 (1899), 297–305.

[33] Thorvald N. Thiele. 1895. Om Flerfoldsvalg. *Oversigt over det Kongelige Danske Videnskabernes Selskabs Forhandlinger* (1895), 415–441.