Soc Choice Welf (2017) 48:461–485
DOI 10.1007/s00355-016-1019-3

ORIGINAL PAPER

Justified representation in approval-based committee voting

Haris Aziz[1] · **Markus Brill**[2] · **Vincent Conitzer**[3] · **Edith Elkind**[2] · **Rupert Freeman**[3] · **Toby Walsh**[1]

Published online: 9 January 2017

Abstract We consider approval-based committee voting, i.e. the setting where each voter approves a subset of candidates, and these votes are then used to select a fixed-size set of winners (committee). We propose a natural axiom for this setting, which we call justified representation (JR). This axiom requires that if a large enough group of voters exhibits agreement by supporting the same candidate, then at least one voter in this group has an approved candidate in the winning committee. We show that for every list of ballots it is possible to select a committee that provides JR. However, it turns out that several prominent approval-based voting rules may fail to output such a committee. In particular, while Proportional Approval Voting (PAV) always outputs a committee that provides JR, Sequential Proportional Approval Voting (SeqPAV), which is a tractable approximation to PAV, does not have this property. We then introduce a stronger ver-

✉ Haris Aziz
haris.aziz@data61.csiro.au

Markus Brill
mbrill@cs.ox.ac.uk

Vincent Conitzer
conitzer@cs.duke.edu

Edith Elkind
elkind@cs.ox.ac.uk

Rupert Freeman
rupert@cs.duke.edu

Toby Walsh
toby.walsh@data61.csiro.au

1 Data61, CSIRO and UNSW Australia, Sydney 2052, Australia

2 University of Oxford, Oxford OX1 3QD, UK

3 Duke University, Durham, NC 27708, USA

sion of the JR axiom, which we call extended justified representation (EJR), and show that PAV satisfies EJR, while other rules we consider do not; indeed, EJR can be used to characterize PAV within the class of weighted PAV rules. We also consider several other questions related to JR and EJR, including the relationship between JR/EJR and core stability, and the complexity of the associated computational problems.

1 Introduction

Aggregation of preferences is a central problem in the field of social choice. While the most-studied scenario is that of selecting a single candidate out of many, it is often the case that one needs to select a fixed-size *set of winners* (a *committee*): this includes domains such as parliamentary elections, the hiring of faculty members, or (automated) agents deciding on a set of plans (LeGrand et al. 2007; Davis et al. 2014; Elkind et al. 2015; Skowron et al. 2016; Elkind et al. 2017; Aziz et al. 2016). The study of the computational complexity of voting rules that output committees is an active research direction (Procaccia et al. 2008; Meir et al. 2008; Caragiannis et al. 2010; Lu and Boutilier 2011; Cornaz et al. 2012; Betzler et al. 2013; Skowron et al. 2015a, b).

In this paper we consider approval-based rules, where each voter lists the subset of candidates that she approves of. There is a growing literature on voting rules that are based on approval ballots: the Handbook on Approval Voting (Laslier and Sanver 2010) provides a very useful survey of pre-2010 research on this topic, and after this seminal book was published, various aspects of approval voting continued to attract a considerable amount of attention (see, e.g., the papers of Caragiannis et al. 2010; Endriss 2013; Duddy 2014). One of the advantages of approval ballots is their simplicity: compared to ranked ballots, approval ballots reduce the cognitive burden on voters (rather than providing a full ranking of the candidates, a voter only needs to decide which candidates to approve) and are also easier to communicate to the election authority. The most straightforward way to aggregate approvals is to have every approval for a candidate contribute one point to that candidate's score and select the candidates with the highest score. This rule is called *Approval Voting* (AV). AV has many desirable properties in the single-winner case (Brams and Fishburn 2007; Brams et al. 2006; Endriss 2013), including its "simplicity, propensity to elect Condorcet winners (when they exist), its robustness to manipulation and its monotonicity" (Brams 2010, p. viii). However, for the case of multiple winners, the merits of AV are "less clear" (Brams 2010, p. viii). For example, AV may fail proportional representation: if the goal is to select $k > 1$ winners, 51% of the voters approve the same k candidates, and the remaining voters approve a disjoint set of k candidates, then the voters in minority do not get any of their approved candidates selected.

As a consequence, over the years, several multi-winner rules based on approval ballots have been proposed (see, e.g., the survey by Kilgour 2010); we will now briefly describe the rules that will be considered in this paper (see Sect. 2 for formal definitions). Under Proportional Approval Voting (PAV), each voter's contribution to the committee's total score depends on how many candidates from the voter's approval set have been elected. In the canonical variant of this rule the marginal utility of the ℓ-th approved candidate is $\frac{1}{\ell}$, i.e. this rule is associated with the weight vector $(1, \frac{1}{2}, \frac{1}{3}, \dots)$; other weight vectors can be used as well, resulting in the family of *weighted* PAV*rules*. A sequential variant of PAV is known as *Sequential Proportional Approval Voting* (SeqPAV); again, by varying the weight vector, we obtain the family of *weighted* SeqPAV*rules*. Another way to modulate the approvals is through computing a satisfaction score for each voter based on the ratio of the number of their approved candidates appearing in the committee and their total number of approved candidates; this idea leads to Satisfaction Approval Voting (SAV). One could also use a distance-based approach: Minimax Approval Voting (MAV) selects a set of k candidates that minimizes the maximum Hamming distance from the submitted ballots. Finally, one could adapt classic rules that provide fully proportional representation, such as the Chamberlin–Courant rule (Chamberlin and Courant 1983) or the Monroe rule (Monroe 1995), to work with approval ballots, by using each voter's ballot as a scoring vector. All the rules informally described above have a more egalitarian objective than AV. For example, Steven Brams, a proponent of AV in single-winner elections, has argued that SAV is more suitable for equitable representation in multi-winner elections (Brams and Kilgour 2014).

The relative merits of approval-based multi-winner rules and the complexity of winner determination under these rules have been examined in great detail in both economics and computer science in recent years (Brams and Fishburn 2007; LeGrand et al. 2007; Meir et al. 2008; Caragiannis et al. 2010; Aziz et al. 2015; Byrka and Sornat 2014; Misra et al. 2015). On the other hand, there has been limited axiomatic analysis of these rules from the perspective of representation (see, however, Sect. 7).

In this paper, we introduce the notion of justified representation (JR) in approval-based voting. Briefly, a committee is said to provide justified representation for a given set of ballots if every large enough group of voters with shared preferences is allocated at least one representative. A rule is said to satisfy justified representation if it always outputs a committee that provides justified representation. This concept is related to the *Droop proportionality criterion* (Droop 1881) and Dummett's *solid coalition property* (Dummett 1984; Tideman and Richardson 2000; Elkind et al. 2017), but is specific to approval-based elections.

We show that every set of ballots admits a committee that provides justified representation; moreover, such a committee can be computed efficiently, and checking whether a given committee provides JR can be done in polynomial time as well. This shows that justified representation is a reasonable requirement. However, it turns out that many popular multi-winner approval-based rules fail JR; in particular, this is the case for AV, SAV, MAV and the canonical variant of SeqPAV. On the positive side, JR is satisfied by some of the weighted PAV rules, including the canonical PAV rule, as well as by the weighted SeqPAV rule associated with the weight vector $(1, 0, \dots)$ and by the Monroe rule. Also, MAV satisfies JR for a restricted domain of voters' preferences. We then consider a strengthening of the JR axiom, which we call extended justified representation (EJR). This axiom captures the intuition that a large group of voters with similar preferences may deserve not just one, but several representatives. EJR turns out to be a more demanding property than JR: of all voting rules considered in this paper, only the canonical PAV rule satisfies EJR. Thus, in particular, EJR characterizes the canonical PAV rule within the class of weighted PAV rules. However, we show that it is computationally hard to check whether a given committee provides EJR.

We also consider other strengthenings of JR, which we call *semi-strong justified representation* and *strong justified representation*; however, it turns out that for some inputs the requirements imposed by these axioms are impossible to satisfy. Finally, we explore the relationship between JR/EJR and core stability in a non-transferable utility game that can be associated with a multiwinner approval voting scenario. We show that, even though EJR may appear to be similar to core stability, it is, in fact, a strictly weaker condition. Indeed, the core stability condition appears to be too demanding, as none of the voting rules considered in our work is guaranteed to produce a core stable outcome, even when the core is known to be non-empty. We conclude the paper by discussing related work and identifying several directions for future work.

2 Preliminaries

We consider a social choice setting with a set $N = \{1, \ldots, n\}$ of voters and a set C of candidates. Each voter $i \in N$ submits an approval ballot $A_i \subseteq C$, which represents the subset of candidates that she approves of. We refer to the list $\mathbf{A} = (A_1, \ldots, A_n)$ of approval ballots as the *ballot profile*. We will consider *approval-based multi-winner voting rules* that take as input a tuple $(N, C, \mathbf{A}, k)$, where k is a positive integer that satisfies $k \leq |C|$, and return a subset $W \subseteq C$ of size k, which we call the *winning set*, or *committee* (Kilgour and Marshall 2012). We omit N and C from the notation when they are clear from the context. Several approval-based multi-winner rules are defined below. Whenever the description of the rule does not uniquely specify a winning set, we assume that ties are broken according to some deterministic procedure; however, most of our results do not depend on the tie-breaking rule.

2.1 Approval-based multi-winner rules

Approval voting (AV) Under AV, the winners are the k candidates that receive the largest number of approvals. Formally, the *approval score* of a candidate $c \in C$ is defined as $|\{i \in N \mid c \in A_i\}|$, and AV outputs a set W of size k that maximizes $\sum_{c \in W} |\{i \in N \mid c \in A_i\}|$. AV has been adopted by several academic and professional societies such as the Institute of Electrical and Electronics Engineers (IEEE), the International Joint Conference on Artificial Intelligence (IJCAI), and the Society for Social Choice and Welfare (SSCW).

Satisfaction approval voting (SAV) A voter's *satisfaction score* is the fraction of her approved candidates that are elected. SAV maximizes the sum of voters' satisfaction scores. Formally, SAV outputs a set $W \subseteq C$ of size k that maximizes $\sum_{i \in N} \frac{|W \cap A_i|}{|A_i|}$. This rule was proposed with the aim of "representing more diverse interests" than AV (Brams and Kilgour 2014).[1]

[1] SAV is equivalent to *equal and even cumulative voting* (e.g., see Alcalde-Unzu and Vorsatz 2009), where for each voter, a score of 1 is divided evenly among all candidates in the voter's approval set, and the k candidates with the highest total score are selected (see also Brams and Kilgour 2014, p. 328).

Proportional approval voting (PAV) Under PAV, a voter is assumed to derive a utility of $1 + \frac{1}{2} + \frac{1}{3} + \cdots + \frac{1}{j}$ from a committee that contains exactly j of her approved candidates, and the goal is to maximize the sum of the voters' utilities. Formally, the PAV-score of a set $W \subseteq C$ is defined as $\sum_{i \in N} r(|W \cap A_i|)$, where $r(p) = \sum_{j=1}^{p} \frac{1}{j}$, and PAV outputs a set $W \subseteq C$ of size k with the highest PAV-score. Though sometimes attributed to Forest Simmons, PAV was already proposed by the Danish polymath Thorvald N. Thiele in the nineteenth century (Thiele 1895).[2] PAV captures the idea of diminishing returns: an individual voter's preferences should count less the more she is satisfied.

In fact, Thiele (1895) not only introduced PAV, but a whole family of *weighted* PAV *rules*: for every vector[3] $\mathbf{w} = (w_1, w_2, \dots)$, where $w_1, w_2, \dots$ are non-negative reals, the voting rule $\mathbf{w}$-PAV operates as follows. Given a ballot profile $(A_1, \dots, A_n)$ and a target committee size k, $\mathbf{w}$-PAV returns a set W of size k with the highest $\mathbf{w}$-PAV score, defined by $\sum_{i \in N} r_{\mathbf{w}}(|W \cap A_i|)$, where $r_{\mathbf{w}}(p) = \sum_{j=1}^{p} w_j$. Usually, it is required that $w_1 = 1$ and $w_1 \geq w_2 \geq \dots$. The latter constraint is appropriate in the context of representative democracy: it is motivated by the intuition that once an agent already has one or more representatives in the committee, that agent should have less priority for further representation. In what follows, we will always impose the constraint $w_1 = 1$ (as we can always rescale the weight vector, this is equivalent to requiring that $w_1 > 0$; while the case $w_1 = 0$ may be of interest in some applications, we omit it in order to keep the length of the paper manageable)[4] and explicitly indicate which of our results require that $w_1 \geq w_2 \geq \dots$; in particular, for our characterization of PAV in Theorem 11 this is not the case.

Sequential proportional approval voting (*SeqPAV*) SeqPAV converts PAV into a multi-round rule, by selecting a candidate in each round and then reweighing the approvals for the subsequent rounds. Specifically, SeqPAV starts by setting $W = \emptyset$. Then in round j, $j = 1, \dots, k$, it computes the *approval weight* of each candidate c in $C \setminus W$ as $\sum_{i: c \in A_i} \frac{1}{1+|W \cap A_i|}$, selects a candidate with the highest approval weight, and adds him to W. After k rounds, it outputs the set W. Several papers (including an earlier version of our work) refer to SeqPAV as "Reweighted Approval Voting (RAV)"; this rule was used briefly in Sweden during the early 1900s.

Thiele (1895) proposed SeqPAV as a tractable approximation to PAV (see Sect. 2.2 for a discussion of the computational complexity of these rules and the relationship between them). We note that there are several other examples of voting rules that were conceived as approximate versions of other rules, yet became viewed as legitimate voting rules in and of themselves; two representative examples are the Simplified Dodgson rule of Tideman (2006), which was designed as an approximate version of the Dodgson rule (see the discussion by Caragiannis et al. 2014), and the Greedy

[2] We are grateful to Svante Janson and Xavier Mora for pointing this out to us.

[3] It is convenient to think of $\mathbf{w}$ as an infinite vector; note that for an election with m candidates only the first m entries of $\mathbf{w}$ matter. To analyze the complexity of $\mathbf{w}$-PAV rules, one would have to place additional requirements on $\mathbf{w}$; however, we do not consider computational properties of such rules in this paper.

[4] $\mathbf{w}$-PAV rules with $w_1 = 0$ have been considered by Fishburn and Pekec (2004) and Skowron et al. (2016). We note that such rules do not satisfy justified representation (as defined in Sect. 3).

Monroe rule of Skowron et al. (2015a), which approximates the Monroe rule (Monroe 1995).

Just as for PAV, the definition of SeqPAV can be extended to score vectors other than $(1, \frac{1}{2}, \frac{1}{3}, \dots)$: every vector $\mathbf{w} = (w_1, w_2, \dots)$ defines a sequential voting rule $\mathbf{w}$-SeqPAV, which proceeds as SeqPAV, except that it computes the approval weight of a candidate c in round j as $\sum_{i:c\in A_i} w_{|W\cap A_i|+1}$, where W is the winning set after the first $j-1$ rounds. Again, we impose the constraint $w_1 = 1$ (note that if $w_1 = 0$, then $\mathbf{w}$-SeqPAV can pick an arbitrary candidate at the first step, which is obviously undesirable).

A particularly interesting rule in this class is $(1, 0, \dots)$-SeqPAV: this rule, which we will refer to as Greedy Approval Voting (GreedyAV), is the sequential version of what Thiele (1895) called the "weak method." GreedyAV can be seen as a variant of the SweetSpotGreedy (SSG) algorithm of Lu and Boutilier (2011), and admits a very simple description: we pick candidates one by one, trying to 'cover' as many currently 'uncovered' voters as possible. In more detail, a winning committee under this rule can be computed by the following algorithm. We start by setting $C' = C$, $\mathbf{A}' = \mathbf{A}$, and $W = \emptyset$. As long as $|W| < k$ and $\mathbf{A}'$ is non-empty, we pick a candidate $c \in C'$ that has the highest approval score with respect to $\mathbf{A}'$, and set $W := W \cup \{c\}$, $C' := C'\setminus\{c\}$. Also, we remove from $\mathbf{A}'$ all ballots A_i such that $c \in A_i$. If at some point we have $|W| < k$ and $\mathbf{A}'$ is empty, we add an arbitrary set of $k - |W|$ candidates from C' to W and return W; if this does not happen, we terminate after having picked k candidates.

We will also consider a variant of GreedyAV, where, at each step, after selecting a candidate c, instead of removing all ballots in $N_c = \{i \in N \mid c \in A_i\}$ from $\mathbf{A}'$, we remove a subset of N_c of size $\min\left(\left\lceil \frac{n}{k} \right\rceil, |N_c|\right)$. This rule can be seen as an adaptation of the classic STV rule to approval ballots, and we will refer to it as *HareAV*.[5]

Minimax approval voting (MAV) MAV returns a committee W that minimizes the maximum *Hamming distance* between W and the voters' ballots; this rule was proposed by Brams et al. (2007). Formally, for any $Q, T \subseteq C$, let $d(Q, T) = |Q\setminus T| + |T\setminus Q|$. Define the MAV-score of a set $W \subseteq C$ as $\max(d(W, A_1), \dots, d(W, A_n))$. MAV outputs a size-$k$ set with the lowest MAV-score.

Chamberlin–Courant and Monroe approval voting (CCAV and MonroeAV) The Chamberlin–Courant rule (Chamberlin and Courant 1983) is usually defined for the setting where each voter provides a full ranking of the candidates. Each voter $i \in N$ is associated with a scoring vector $\mathbf{u}^i = (u^i_1, \dots, u^i_{|C|})$ whose entries are non-negative reals; we think of u^i_j as voter i's satisfaction from being represented by candidate c_j. A voter's satisfaction from a committee W is defined as $\max_{c_j\in W} u^i_j$, and the rule returns a committee of size k that maximizes the sum of voters' satisfactions. For the case of approval ballots, it is natural to define the scoring vectors by setting $u^i_j = 1$ if $c_j \in A_i$ and $u^i_j = 0$ otherwise; that is, a voter is satisfied by a committee if this committee contains one of her approved candidates. Thus, the resulting rule is equivalent to $(1, 0, \dots)$-PAV (and therefore we will not discuss it separately).

[5] For readability, we use the Hare quota $\lceil \frac{n}{k} \rceil$; this choice of quota motivates our name for this rule. However, all our proofs go through if we use the Droop quota $\left\lfloor \frac{n}{k+1} \right\rfloor + 1$ instead. For a discussion of differences between these two quotas, see the article of Tideman (1995).

The Monroe rule (Monroe 1995) is a modification of the Chamberlin–Courant rule where each committee member represents roughly the same number of voters. Just as under the Chamberlin–Courant rule, we have a scoring vector $\mathbf{u}^i = (u^i_1, \dots, u^i_{|C|})$ for each voter $i \in N$. Given a committee W of size k, we say that a mapping $\pi : N \to W$ is *valid* if it satisfies $|\pi^{-1}(c)| \in \left\{\left\lfloor \frac{n}{k} \right\rfloor, \left\lceil \frac{n}{k} \right\rceil\right\}$ for each $c \in W$. The Monroe score of a valid mapping π is given by $\sum_{i \in N} u^i_{\pi(i)}$, and the Monroe score of a committee W is the maximum Monroe score of a valid mapping from N to W. The Monroe rule returns a size-k committee with the maximum Monroe score. For approval ballots, we define the scoring vectors in the same manner as for the Chamberlin–Courant Approval Voting rule; we call the resulting rule the Monroe Approval Voting rule (MonroeAV).

We note that for $k = 1$, AV, PAV, SeqPAV, GreedyAV, HareAV and MonroeAV produce the same output if there is a unique candidate with the highest approval score. However, such a candidate need not be a winner under SAV or MAV.

2.2 Computational complexity

The rules listed above differ from a computational perspective. For some of these rules, namely, AV, SAV, SeqPAV, GreedyAV and HareAV, a winning committee can be computed in polynomial time; this is also true for $\mathbf{w}$-SeqPAV as long as the entries of the weight vector are rational numbers that can be efficiently computed given the number of candidates. In contrast, PAV, MAV, and MonroeAV are computationally hard (Aziz et al. 2015; Skowron et al. 2016; LeGrand et al. 2007; Procaccia et al. 2008); for $\mathbf{w}$-PAV, the hardness result holds for most weight vectors, including $(1, 0, \dots)$ (i.e. it holds for CCAV). However, both PAV and MAV admit efficient approximation algorithms (i.e. algorithms that output committees which are approximately optimal with respect to the optimization criteria of these rules) and have been analyzed from the perspective of parameterized complexity. Specifically, $\mathbf{w}$-PAV admits an efficient $\left(1 - \frac{1}{e}\right)$-approximation algorithm as long as the weight vector $\mathbf{w}$ is efficiently computable and non-increasing; in fact, such an algorithm is provided by $\mathbf{w}$-SeqPAV (Skowron et al. 2016). For MAV, LeGrand et al. (2007) proposed a simple 3-approximation algorithm; Caragiannis et al. (2010) improve the approximation ratio to 2 and Byrka and Sornat (2014) develop a polynomial-time approximation scheme. Misra et al. (2015) show that MAV is fixed-parameter tractable for a number of natural parameters; Elkind and Lackner (2015) obtain fixed-parameter tractability results for PAV when voters' preferences are, in some sense, single-dimensional. There is also a number of tractability results for CCAV, and, to a lesser extent, for MonroeAV; we refer the reader to the work of Skowron et al. (2016) and references therein.

3 Justified representation

We will now define one of the main concepts of this paper.

Definition 1 (*Justified representation (JR)*) Given a ballot profile $\mathbf{A} = (A_1, \dots, A_n)$ over a candidate set C and a target committee size k, we say that a set of candidates W of size $|W| = k$ *provides justified representation for* $(\mathbf{A}, k)$ if there does not exist a set of voters $N^* \subseteq N$ with $|N^*| \geq \frac{n}{k}$ such that $\bigcap_{i \in N^*} A_i \neq \emptyset$ and $A_i \cap W = \emptyset$ for all $i \in N^*$. We say that an approval-based voting rule *satisfies justified representation* (*JR*) if for every profile $\mathbf{A} = (A_1, \ldots, A_n)$ and every target committee size k it outputs a winning set that provides justified representation for $(\mathbf{A}, k)$.

The logic behind this definition is that if k candidates are to be selected, then, intuitively, each group of $\frac{n}{k}$ voters "deserves" a representative. Therefore, a set of $\frac{n}{k}$ voters that have at least one candidate in common should not be completely unrepresented. We refer the reader to Sect. 6 for a discussion of alternative definitions.

3.1 Existence and computational properties

We start our analysis of justified representation by observing that, for every ballot profile $\mathbf{A}$ and every value of k, there is a committee that provides justified representation for $(\mathbf{A}, k)$, and, moreover, such a committee can be computed efficiently. In fact, both GreedyAV and HareAV output a committee that provides JR.

Theorem 1 *GreedyAV and HareAV satisfy JR.*

Proof We present a proof that applies to both GreedyAV and HareAV. Suppose for the sake of contradiction that for some ballot profile $\mathbf{A} = (A_1, \ldots, A_n)$ and some $k > 0$, GreedyAV (respectively, HareAV) outputs a committee that does not provide justified representation for $(\mathbf{A}, k)$. Then there exists a set $N^* \subseteq N$ with $|N^*| \geq \frac{n}{k}$ such that $\bigcap_{i \in N^*} A_i \neq \emptyset$ and, when GreedyAV (respectively, HareAV) terminates, every ballot A_i with $i \in N^*$ is still in $\mathbf{A}'$. Consider some candidate $c \in \bigcap_{i \in N^*} A_i$. At every point in the execution of our algorithm, c's approval score is at least $|N^*| \geq \frac{n}{k}$. As c was not elected, at every stage the algorithm selected a candidate whose approval score was at least as high as that of c. Thus, at the end of each stage the algorithm removed from $\mathbf{A}'$ at least $\lceil \frac{n}{k} \rceil$ ballots containing the candidate added to W at that stage, so altogether the algorithm has removed at least $k \cdot \frac{n}{k}$ ballots from $\mathbf{A}'$. This contradicts the assumption that $\mathbf{A}'$ contains at least $\frac{n}{k}$ ballots when the algorithm terminates. $\square$

Theorem 1 shows that it is easy to find a committee that provides justified representation for a given ballot profile. It is also not too hard to check whether a given committee W provides JR. Indeed, while it may seem that we need to consider every subset of voters of size $\frac{n}{k}$, in fact it is sufficient to consider the candidates one by one, and, for each candidate c, compute $s(c) = |\{i \in N \mid c \in A_i, A_i \cap W = \emptyset\}|$; the set W fails to provide justified representation for $(\mathbf{A}, k)$ if and only if there exists a candidate c with $s(c) \geq \frac{n}{k}$. We obtain the following theorem.

Theorem 2 *There exists a polynomial-time algorithm that, given a ballot profile* $\mathbf{A}$ *over a candidate set C, and a committee W,* $|W| = k$*, decides whether W provides justified representation for* $(\mathbf{A}, k)$.

3.2 Justified representation and unanimity

A desirable property of single-winner approval-based voting rules is *unanimity*: a voting rule is unanimous if, given a ballot profile $(A_1, \ldots, A_n)$ with $\cap_{i \in N} A_i \neq \emptyset$, it outputs a candidate in $\bigcap_{i \in N} A_i$. This property is somewhat similar in spirit to JR, so the reader may expect that for $k = 1$ it is equivalent to JR. However, it turns out that the JR axiom is strictly weaker than unanimity for $k = 1$: while unanimity implies JR, the converse is not true, as illustrated by the following example.

Example 1 Let $N = \{1, \ldots, n\}$, $C = \{a, b_1, \ldots, b_n\}$, $A_i = \{a, b_i\}$ for $i \in N$. Consider a voting rule that for $k = 1$ outputs b_1 on this profile and coincides with GreedyAV in all other cases. Clearly, this rule is not unanimous; however, it satisfies JR, as it is impossible to find a group of $\frac{n}{k} = n$ unrepresented voters for $(A_1, \ldots, A_n)$.

It is not immediately clear how to define unanimity for multi-winner voting rules; however, any reasonable definition would be equivalent to the standard definition of unanimity when $k = 1$, and therefore would be different from justified representation.

We remark that a rule can be unanimous for $k = 1$ and provide JR for all values of k: this is the case, for instance, for GreedyAV.

4 Justified representation under approval-based rules

We have argued that justified representation is a reasonable condition: there always exists a committee that provides it, and, moreover, such a committee can be computed efficiently. It is therefore natural to ask whether prominent voting rules satisfy JR. In this section, we will answer this question for AV, SAV, MAV, PAV, SeqPAV, and MonroeAV. We will also identify conditions on $\mathbf{w}$ that are sufficient/necessary for $\mathbf{w}$-PAV and $\mathbf{w}$-SeqPAV to satisfy JR.

In what follows, for each rule we will try to identify the range of values of k for which this rule satisfies JR. Trivially, all rules that we consider satisfy JR for $k = 1$. It turns out that AV fails JR for $k > 2$, and for $k = 2$ the answer depends on the tie-breaking rule.

Theorem 3 *For $k = 2$, AV satisfies JR if ties are broken in favor of sets that provide JR. For $k \geq 3$, AV fails JR.*

Proof Suppose first that $k = 2$. Fix a ballot profile $\mathbf{A}$. If every candidate is approved by fewer than $\frac{n}{2}$ voters in $\mathbf{A}$, JR is trivially satisfied. If some candidate is approved by more than $\frac{n}{2}$ voters in $\mathbf{A}$, then AV selects some such candidate, in which case no group of $\lceil \frac{n}{2} \rceil$ voters is unrepresented, so JR is satisfied in this case as well. It remains to consider the case where $n = 2n'$, some candidates are approved by n' voters, and no candidate is approved by more than n' voters. Then AV necessarily picks at least one candidate approved by n' voters; denote this candidate by c. In this situation JR can only be violated if the n' voters who do not approve c all approve the same candidate (say, c'), and this candidate is not elected. But the approval score of c' is n', and, by our assumption, the approval score of every candidate is at most n', so this is a contradiction with our tie-breaking rule. This argument also illustrates why the assumption on the tie-breaking rule is necessary: it can be the case that n' voters approve c and c'', and the remaining n' voters approve c', in which case the approval score of $\{c, c''\}$ is the same as that of $\{c, c'\}$.

For $k \geq 3$, we let $C = \{c_0, c_1, \ldots, c_k\}$, $n = k$, and consider the profile where the first voter approves c_0, whereas each of the remaining voters approves all of $c_1, \ldots, c_k$. JR requires c_0 to be selected, but AV selects $\{c_1, \ldots, c_k\}$. □

On the other hand, SAV and MAV fail JR even for $k = 2$.

Theorem 4 *SAV and MAV do not satisfy JR for $k \geq 2$.*

Proof We first consider SAV. Fix $k \geq 2$, let $X = \{x_1, \ldots, x_k, x_{k+1}\}$, $Y = \{y_1, \ldots, y_k\}$, $C = X \cup Y$, and consider the profile $(A_1, \ldots, A_k)$, where $A_1 = X$, $A_2 = \{y_1, y_2\}$, $A_i = \{y_i\}$ for $i = 3, \ldots, k$. JR requires each voter to be represented, but SAV will choose Y: the SAV-score of Y is $k - 1$, whereas the SAV-score of every committee W with $W \cap X \neq \emptyset$ is at most $k - 2 + \frac{1}{2} + \frac{1}{k+1} < k - 1$. Therefore, the first voter will remain unrepresented.

For MAV, we use the following construction. Fix $k \geq 2$, let $X = \{x_1, \ldots, x_k\}$, $Y = \{y_1, \ldots, y_k\}$, $C = X \cup Y \cup \{z\}$, and consider the profile $(A_1, \ldots, A_{2k})$, where $A_i = \{x_i, y_i\}$ for $i = 1, \ldots, k$, $A_i = \{z\}$ for $i = k + 1, \ldots, 2k$. Every committee of size k that provides JR for this profile contains z. However, MAV fails to select z. Indeed, the MAV-score of X is $k + 1$: we have $d(X, A_i) = k$ for $i \leq k$ and $d(X, A_i) = k + 1$ for $i > k$. Now, consider some committee W with $|W| = k$, $z \in W$. We have $A_i \cap W = \emptyset$ for some $i \leq k$, so $d(W, A_i) = k + 2$. Thus, MAV prefers X to any committee that includes z. □

The constructions used in the proof of Theorem 4 show that MAV and SAV may behave very differently: SAV appears to favor voters who approve very few candidates, whereas MAV appears to favor voters who approve many candidates.

Interestingly, we can show that MAV satisfies JR if we assume that each voter approves exactly k candidates and ties are broken in favor of sets that provide JR.

Theorem 5 *If the target committee size is k, $|A_i| = k$ for all $i \in N$, and ties are broken in favor of sets that provide JR, then MAV satisfies JR.*

Proof Consider a profile $\mathbf{A} = (A_1, \ldots, A_n)$ with $|A_i| = k$ for all $i \in N$.

Observe that if there exists a set of candidates W with $|W| = k$ such that $W \cap A_i \neq \emptyset$ for all $i \in N$, then MAV will necessarily select some such set. Indeed, for any such set W we have $d(W, A_i) \leq 2k - 1$ for each $i \in N$, whereas if $W' \cap A_i = \emptyset$ for some set W' with $|W'| = k$ and some $i \in N$, then $d(W', A_i) = 2k$. Further, by definition, every set W such that $|W| = k$ and $W \cap A_i \neq \emptyset$ for all $i \in N$ provides justified representation for $(\mathbf{A}, k)$.

On the other hand, if there is no k-element set of candidates that intersects each A_i, $i \in N$, then the MAV-score of every set of size k is $2k$, and therefore MAV can pick an arbitrary size-k subset. Since we assumed that the tie-breaking rule favors sets that provide JR, our claim follows. □

While Theorem 5 provides an example of a setting where MAV satisfies JR, this result is not entirely satisfactory: first, we had to place a strong restriction on voters' preferences, and, second, we used a tie-breaking rule that was tailored to JR.

We will now show that PAV satisfies JR, for all ballot profiles and irrespective of the tie-breaking rule.

Theorem 6 *PAV satisfies JR.*

Proof Fix a ballot profile $\mathbf{A} = (A_1, \ldots, A_n)$ and a $k > 0$ and let $s = \lceil \frac{n}{k} \rceil$. Let W be the output of PAV on $(\mathbf{A}, k)$. Suppose for the sake of contradiction that there exists a set $N^* \subseteq N$, $|N^*| \geq s$, such that $\bigcap_{i \in N^*} A_i \neq \emptyset$, but $W \cap (\bigcup_{i \in N^*} A_i) = \emptyset$. Let c be some candidate approved by all voters in N^*.

For each candidate $w \in W$, define its *marginal contribution* as the difference between the PAV-score of W and that of $W \setminus \{w\}$. Let $m(W)$ denote the sum of marginal contributions of all candidates in W. Observe that if c were to be added to the winning set, this would increase the PAV-score by at least s. Therefore, it suffices to argue that the marginal contribution of some candidate in W is less than s: this would mean that swapping this candidate with c increases the PAV-score, a contradiction. To this end, we will prove that $m(W) \leq s(k-1)$; as $|W| = k$, our claim would then follow by the pigeonhole principle.

Consider the set $N \setminus N^*$; we have $n \leq sk$, so $|N \setminus N^*| \leq n - s \leq s(k-1)$. Pick a voter $i \in N \setminus N^*$, and let $j = |A_i \cap W|$. If $j > 0$, this voter contributes exactly $\frac{1}{j}$ to the marginal contribution of each candidate in $A_i \cap W$, and hence her contribution to $m(W)$ is exactly 1. If $j = 0$, this voter does not contribute to $m(W)$ at all. Therefore, we have $m(W) \leq |N \setminus N^*| \leq s(k-1)$, which is what we wanted to prove. □

The reader may observe that the proof of Theorem 6 applies to all voting rules of the form $\mathbf{w}$-PAV where the weight vector satisfies $w_1 = 1$ and $w_j \leq \frac{1}{j}$ for all $j > 1$. In Sect. 5 we will see that this condition on $\mathbf{w}$ is also necessary for $\mathbf{w}$-PAV to satisfy JR.

Next, we consider SeqPAV. As this voting rule can be viewed as a tractable approximation of PAV (recall that PAV is NP-hard to compute), one could expect that SeqPAV satisfies JR as well. However, this turns out not to be the case, at least if k is sufficiently large.

Theorem 7 *SeqPAV satisfies JR for $k = 2$, but fails it for $k \geq 10$.*

Proof For $k = 2$, we can use essentially the same argument as for AV; however, we do not need to assume anything about the tie-breaking rule. This is because if there are three candidates, c, c', and c'', such that c and c'' are approved by the same $\frac{n}{2}$ voters, whereas c' is approved by the remaining $\frac{n}{2}$ voters, and SeqPAV selects c in the first round, then in the second round SeqPAV favors c' over c''.

Now, suppose that $k = 10$. Consider a profile over a candidate set $C = \{c_1, \ldots, c_{11}\}$ with 1199 voters who submit the following ballots:

$81 \times \{c_1, c_2\}$,	$81 \times \{c_1, c_3\}$,	$80 \times \{c_2\}$,	$80 \times \{c_3\}$,
$81 \times \{c_4, c_5\}$,	$81 \times \{c_4, c_6\}$,	$80 \times \{c_5\}$,	$80 \times \{c_6\}$,
$49 \times \{c_7, c_8\}$,	$49 \times \{c_7, c_9\}$,	$49 \times \{c_7, c_{10}\}$,	
$96 \times \{c_8\}$,	$96 \times \{c_9\}$,	$96 \times \{c_{10}\}$,	$120 \times \{c_{11}\}$.

Candidates c_1 and c_4 are each approved by 162 voters, the most of any candidate, and these blocks of 162 voters do not overlap, so SeqPAV selects c_1 and c_4 first. This reduces the SeqPAV scores of c_2, c_3, c_5 and c_6 from $80 + 81 = 161$ to $80 + 40.5 = 120.5$, so c_7, whose SeqPAV score is 147, is selected next. Now, the SeqPAV scores of c_8, c_9 and c_{10} become $96+24.5 = 120.5$. The selection of any of c_2, c_3, c_5, c_6, c_8, c_9 or c_{10} does not affect the SeqPAV score of the others, so all seven of these candidates will be selected before c_{11}, who has 120 approvals. Thus, after the selection of 10 candidates, there are $120 > \frac{1199}{10} = \frac{n}{k}$ unrepresented voters who jointly approve c_{11}.

To extend this construction to $k > 10$, we create $k - 10$ additional candidates and $120(k - 10)$ additional voters such that for each new candidate, there are 120 new voters who approve that candidate only. Note that we still have $120 > \frac{n}{k}$. SeqPAV will proceed to select $c_1, \ldots, c_{10}$, followed by $k - 10$ additional candidates, and c_{11} or one of the new candidates will remain unselected. □

While SeqPAV itself does not satisfy JR, one could hope that this can be fixed by tweaking the weights, i.e. that $\mathbf{w}$-SeqPAV satisfies JR for a suitable weight vector $\mathbf{w}$. However, it turns out that $(1, 0, \ldots)$ is essentially the only weight vector for which this is the case: Theorem 7 extends to $\mathbf{w}$-SeqPAV for *every* weight vector $\mathbf{w}$ with $w_1 = 1$, $w_2 > 0$.

Theorem 8 *For every vector* $\mathbf{w} = (w_1, w_2, \ldots)$ *with* $w_1 = 1$, $w_2 > 0$, *there exists a value of* $k_0 > 0$ *such that* $\mathbf{w}$*-SeqPAV does not satisfy JR for* $k > k_0$.

Proof Pick a positive integer $s \geq 8$ such that $w_2 \geq \frac{1}{s}$. Let $C = C_1 \cup C_2 \cup \{x, y\}$, where

$$C_1 = \{c_{i,j} \mid i = 1, \ldots, 2s + 3, j = 1, \ldots, 2s + 1\} \quad \text{and}$$
$$C_2 = \{c_i \mid i = 1, \ldots, 2s + 3\}.$$

For each $i = 1, \ldots, 2s + 3$ and each $j = 1, \ldots, 2s + 1$ we construct $2s^3 - s$ voters who approve $c_{i,j}$ only and s^2 voters who approve $c_{i,j}$ and c_i only. Finally, we construct $2s^3 - 1$ voters who approve x only and $s^2 - 7s - 5$ voters who approve y only (note that the number of voters who approve y is positive by our choice of s).

Set $k_0 = (2s + 2)(2s + 3) = |C_1 \cup C_2|$. Note that the number of voters n is given by

$$\begin{aligned}&(2s + 3)(2s + 1)(2s^3 + s^2 - s) + (2s^3 - 1) + (s^2 - 7s - 5)\\&\quad = (2s + 2)(2s + 3)(2s^3 - 1) = (2s^3 - 1)k_0,\end{aligned}$$

and hence $\frac{n}{k_0} = 2s^3 - 1$.

Under $\mathbf{w}$-SeqPAV initially the score of each candidate in C_2 is $s^2(2s+1) = 2s^3+s^2$, the score of each candidate in C_1 is $2s^3+s^2-s$, the score of x is $2s^3-1$, and the score of y is s^2-7s-5, so in the first $2s+3$ rounds the candidates from C_2 get elected. After that, the score of every candidate in C_1 becomes $2s^3-s+w_2s^2 \geq 2s^3-s+s = 2s^3$, while the scores of x and y remains unchanged. Therefore, in the next $(2s+3)(2s+1)$ rounds the candidates from C_1 get elected. At this point, k candidates are elected, and x is not elected, even though the $2s^3-1 = \frac{n}{k_0}$ voters who approve him do not approve of any of the candidates in the winning set.

To extend this argument to larger values of k, we proceed as in the proof of Theorem 7: for $k > k_0$, we add $k - k_0$ new candidates, and for each new candidate we construct $2s^3 - 1$ new voters who approve that candidate only. Let the resulting number of voters be n'; we have $\frac{n'}{k} = 2s^3 - 1$, so $\mathbf{w}$-SeqPAV will first select the candidates in C_2, followed by the candidates in C_1, and then it will choose $k - k_0$ winners among the new candidates and x. As a result, either x or one of the new candidates will remain unselected. □

Remark 1 Theorem 8 partially subsumes Theorem 7: it implies that SeqPAV fails JR, but the proof only shows that this is the case for $k \geq 18 \cdot 19 = 342$, while Theorem 7 states that SeqPAV fails JR for $k \geq 10$ already. We chose to include the proof of Theorem 7 because we feel that it is useful to know what happens for relatively small values of k. Note, however, that Theorem 7 leaves open the question of whether SeqPAV satisfies JR for $k = 3, \ldots, 9$. Very recently, Sánchez-Fernández et al. (2016, 2017) have answered this question by showing that SeqPAV satisfies JR for $k \leq 5$ and fails it for $k \geq 6$.

If we allow the entries of the weight vector to depend on the number of voters n, we can obtain another class of rules that provide justified representation: the argument used to show that GreedyAV satisfies JR extends to $\mathbf{w}$-SeqPAV where the weight vector $\mathbf{w}$ satisfies $w_1 = 1$, $w_j \leq \frac{1}{n}$ for $j > 1$. In particular, the rule $(1, \frac{1}{n}, \frac{1}{n^2}, \ldots,)$-SeqPAV is somewhat more appealing than GreedyAV: for instance, if $\bigcap_{i \in N} A_i = \{c\}$ and $k > 1$, GreedyAV will pick c, and then behave arbitrarily, whereas $(1, \frac{1}{n}, \frac{1}{n^2}, \ldots,)$-SeqPAV will also pick c, but then it will continue to look for candidates approved by as many voters as possible.

We conclude this section by showing that MonroeAV satisfies JR.

Theorem 9 *MonroeAV satisfies JR.*

Proof Fix a ballot profile $\mathbf{A} = (A_1, \ldots, A_n)$ and a $k > 0$. Let W be an output of MonroeAV on $(\mathbf{A}, k)$. If $A_i \cap W \neq \emptyset$ for all $i \in N$, then W provides justified representation for $(\mathbf{A}, k)$. Thus, assume that this is not the case, i.e. there exists some voter i with $A_i \cap W = \emptyset$. Consider a valid mapping $\pi : N \to W$ whose Monroe score equals the Monroe score of W, let $c = \pi(i)$, and set $s = |\pi^{-1}(c)|$; note that $s \in \left\{\left\lfloor \frac{n}{k} \right\rfloor, \left\lceil \frac{n}{k} \right\rceil\right\}$.

Suppose for the sake of contradiction that W does not provide justified representation for $(\mathbf{A}, k)$. Then by our choice of s there exists a set $N^* \subseteq N$, $|N^*| = s$, such that $\bigcap_{i \in N^*} A_i \neq \emptyset$, but $W \cap (\bigcup_{i \in N^*} A_i) = \emptyset$. Let c' be some candidate approved by all voters in N^*, and set $W' = (W \setminus \{c\}) \cup \{c'\}$. To obtain a contradiction, we will argue that W' has a higher Monroe score than W.

To this end, we will modify π by first swapping the voters in N^* with voters in $\pi^{-1}(c)$ and then assigning the voters in N^* to c'. Formally, let $\sigma : \pi^{-1}(c) \setminus N^* \to N^* \setminus \pi^{-1}(c)$ be a bijection between $\pi^{-1}(c) \setminus N^*$ and $N^* \setminus \pi^{-1}(c)$. We construct a mapping $\hat{\pi} : N \to W'$ by setting

$$\hat{\pi}(i) = \begin{cases} \pi(\sigma(i)) & \text{for } i \in \pi^{-1}(c) \setminus N^*, \\ c' & \text{for } i \in N^*, \\ \pi(i) & \text{for } i \notin \pi^{-1}(c) \cup N^*. \end{cases}$$

Note that $\hat{\pi}$ is a valid mapping: we have $|\hat{\pi}^{-1}(c')| = s$ and $|\hat{\pi}^{-1}(c'')| = |\pi^{-1}(c'')|$ for each $c'' \in W' \setminus \{c'\}$. Now, let us consider the impact of this modification on the Monroe score. The s voters in N^* contributed nothing to the Monroe score of π, and they contribute s to the Monroe score of $\hat{\pi}$. By our choice of c, the voters in $\pi^{-1}(c)$ contributed at most $s-1$ to the Monroe score of π, and their contribution to the Monroe score of $\hat{\pi}$ is non-negative. For all other voters their contribution to the Monroe score of π is equal to their contribution to the Monroe score of π'. Thus, the total Monroe score of $\hat{\pi}$ is higher than that of π. Since the Monroe score of W is equal to the Monroe score of π, and, by definition, the Monroe score of W' is at least the Monroe score of $\hat{\pi}$, we obtain a contradiction. □

5 Extended justified representation

We have identified four (families of) voting rules that satisfy JR for arbitrary ballot profiles: **w**-PAV with $w_1 = 1$, $w_j \leq \frac{1}{j}$ for $j > 1$ (this class includes PAV), **w**-SeqPAV with $w_1 = 1$, $w_j \leq \frac{1}{n}$ for $j > 1$ (this class includes GreedyAV), HareAV and MonroeAV. The obvious advantage of GreedyAV and HareAV is that their output can be computed efficiently, whereas computing the outputs of PAV or MonroeAV is NP-hard. However, GreedyAV puts considerable emphasis on representing *every* voter, at the expense of ensuring that large sets of voters with shared preferences are allocated an adequate number of representatives. This approach may be problematic in a variety of applications, such as selecting a representative assembly, or choosing movies to be shown on an airplane, or foods to be provided at a banquet (see the discussion by Skowron et al. 2016). In particular, it may be desirable to have several assembly members that represent a widely held political position, both to reflect the popularity of this position, and to highlight specific aspects of it, as articulated by different candidates. Consider, for instance, the following example.

Example 2 Let $k = 3$, $C = \{a, b, c, d\}$, and $n = 100$. One voter approves c, one voter approves d, and 98 voters approve a and b. GreedyAV would include both c and d in the winning set, whereas in many settings it would be more reasonable to choose both a and b (and one of c and d); indeed, this is exactly what HareAV would do.

This issue is not addressed by the JR axiom, as this axiom does not care if a given voter is represented by one or more candidates. Thus, if we want to capture the intuition that large cohesive groups of voters should be allocated several representatives, we need a stronger condition. Recall that JR says that each group of $\frac{n}{k}$ voters that all approve the same candidate "deserves" at least one representative. It seems reasonable to scale this idea and say that, for every $\ell > 0$, each group of $\ell \cdot \frac{n}{k}$ voters that all approve the same ℓ candidates "deserves" at least ℓ representatives. This approach can be formalized as follows.

Definition 2 (*Extended justified representation (EJR)*) Given a ballot profile $\mathbf{A} = (A_1, \ldots, A_n)$ over a candidate set C, a target committee size k, $k \leq |C|$, and an integer ℓ, $1 \leq \ell \leq k$, we say that a set of candidates W, $|W| = k$, *provides* ℓ-*justified representation for* $(\mathbf{A}, k)$ if there does not exist a set of voters $N^* \subseteq N$ with $|N^*| \geq \ell \cdot \frac{n}{k}$ such that $|\bigcap_{i \in N^*} A_i| \geq \ell$, but $|A_i \cap W| < \ell$ for each $i \in N^*$; we say that W *provides extended justified representation* (EJR) *for* $(\mathbf{A}, k)$ if it provides ℓ-*JR* for $(\mathbf{A}, k)$ for all ℓ, $1 \le \ell \le k$. We say that an approval-based voting rule *satisfies* ℓ-*justified representation* (ℓ-JR) if for every profile $\mathbf{A} = (A_1, \ldots, A_n)$ and every target committee size k it outputs a committee that provides ℓ-JR for $(\mathbf{A}, k)$. Finally, we say that a rule *satisfies extended justified representation* (EJR) if it satisfies ℓ-*JR* for all ℓ, $1 \le \ell \le k$.

Observe that EJR implies JR, because the latter coincides with 1-JR.

The definition of EJR interprets "a group N^* deserves at least ℓ representatives" as "at least one voter in N^* gets ℓ representatives." Of course, other interpretations are also possible: for instance, we can require that each voter in N^* is represented by ℓ candidates in the winning committee or, alternatively, that the winning committee contains at least ℓ candidates each of which is approved by some member of N^*. However, the former requirement is too strong: it differs from JR for $\ell = 1$ and in Sect. 6 we show that there are ballot profiles for which committees with this property do not exist even for $\ell = 1$. The latter approach, which was very recently proposed by Sánchez-Fernández et al. (2016, 2017) (see the discussion in Sect. 7), is not unreasonable; in particular, it coincides with JR for $\ell = 1$. However, it is strictly less demanding than the approach we take: clearly, every rule that satisfies EJR also satisfies this condition. As it turns out (Theorem 10) that every ballot profile admits a committee that provides EJR, the EJR axiom offers more guidance in choosing a good winning committee than its weaker cousin, while still leaving us with a non-empty set of candidate committees to choose from. Finally, the EJR axiom in its present form is very similar to a core stability condition for a natural NTU game associated with the input profile (see Sect. 5.2); it is not clear if the axiom of Sánchez-Fernández et al. (2016, 2017) admits a similar interpretation.

5.1 Extended justified representation under approval-based rules

It is natural to ask which of the voting rules that satisfy JR also satisfy EJR. Example 2 immediately shows that for GreedyAV the answer is negative. Consequently, no $\mathbf{w}$-SeqPAV rule such that the entries of $\mathbf{w}$ do not depend on n satisfies EJR: if $w_2 = 0$, this rule behaves like GreedyAV on the ballot profile from Example 2 and if $w_2 > 0$, our claim follows from Theorem 8. Moreover, Example 2 also implies that $\mathbf{w}$-SeqPAV rules with $w_j \le \frac{1}{n}$ for $j > 1$ fail EJR as well.

The next example shows that MAV fails EJR even if each voter approves exactly k candidates (recall that under this assumption MAV satisfies JR).

Example 3 Let $k = 4$, $C = C_1 \cup C_2 \cup C_3 \cup C_4$, where $|C_1| = |C_2| = |C_3| = |C_4| = 4$ and the sets C_1, C_2, C_3, C_4 are pairwise disjoint. Let $\mathbf{A} = (A_1, \ldots, A_8)$, where $A_i = C_i$ for $i = 1, 2, 3$, and $A_i = C_4$ for $i = 4, 5, 6, 7, 8$. MAV will select exactly one candidate from each of the sets C_1, C_2, C_3 and C_4, but EJR dictates that at least two candidates from C_4 are chosen.

Further, MonroeAV fails EJR as well.

Example 4 Let $k = 4$, $C = \{c_1, c_2, c_3, c_4, a, b\}$, $N = \{1, \ldots, 8\}$, $A_i = \{c_i\}$ for $i = 1, \ldots, 4$, $A_i = \{c_{i-4}, a, b\}$ for $i = 5, \ldots, 8$. MonroeAV outputs $\{c_1, c_2, c_3, c_4\}$

on this profile, as this is the unique set of candidates with the maximum Monroe score. Thus, every voter is represented by a single candidate, though the voters in $N^* = \{5, 6, 7, 8\}$ "deserve" two candidates.

Example 4 illustrates the conflict between the EJR axiom and the requirement to represent all voters whenever possible. We discuss this issue in more detail in Sect. 7.

For HareAV, it is not hard to construct an example where this rule fails EJR for some way of breaking intermediate ties.

Example 5 Let $N = \{1, \dots, 8\}$, $C = \{a, b, c, d, e, f\}$, $A_1 = A_2 = \{a\}$, $A_3 = A_4 = \{a, b, c\}$, $A_5 = A_6 = \{d, b, c\}$, $A_7 = \{d, e\}$, $A_8 = \{d, f\}$. Suppose that $k = 4$. Note that all voters in $N^* = \{3, 4, 5, 6\}$ approve b and c, and $|N^*| = 2 \cdot \frac{n}{k}$. Under HareAV, at the first step candidates a, b, c and d are tied, so we can select a and remove voters 3 and 4. Next, we have to select d; we can then remove voters 5 and 6. In the remaining two steps, we add e and f to the committee. The resulting committee violates EJR, as each voter in $N^* = \{3, 4, 5, 6\}$ is only represented by a single candidate.

We note that in Example 5 we can remove voters 1 and 2 after selecting a, which enables us to select b or c in the second step and thereby obtain a committee that provides EJR. In fact, we were unable to construct an example where HareAV fails EJR for *all* ways of breaking intermediate ties; we now conjecture that it is always possible to break intermediate ties in HareAV so as to satisfy EJR. However, it is not clear if a tie-breaking rule with this property can be formulated in a succinct manner. Thus, HareAV does not seem particularly useful if we want to find a committee that provides EJR: even if our conjecture is true, we may have to explore all ways of breaking intermediate ties.

In contrast, we will now show that PAV satisfies EJR irrespective of the tie-breaking rule.

Theorem 10 *PAV satisfies EJR.*

Proof Suppose that PAV violates EJR for some value of k, and consider a ballot profile $A_1, \dots, A_n$, a value of $\ell > 0$ and a set of voters N^*, $|N^*| = s \geq \ell \cdot \frac{n}{k}$, that witness this. Let W, $|W| = k$, be the winning set. We know that at least one of the ℓ candidates approved by all voters in N^* is not elected; let c be some such candidate. Each voter in N^* has at most $\ell - 1$ representatives in W, so the marginal contribution of c (if it were to be added to W) would be at least $s \cdot \frac{1}{\ell} \geq \frac{n}{k}$. On the other hand, the argument in the proof of Theorem 6 can be modified to show that the sum of marginal contributions of candidates in W is at most n.

Now, consider some candidate $w \in W$ with the smallest marginal contribution; clearly, his marginal contribution is at most $\frac{n}{k}$. If it is strictly less than $\frac{n}{k}$, we are done, as we can improve the total PAV-score by swapping w and c, a contradiction. Therefore suppose it is exactly $\frac{n}{k}$, and therefore the marginal contribution of each candidate in W is exactly $\frac{n}{k}$. Since PAV satisfies JR, we know that $A_i \cap W \neq \emptyset$ for some $i \in N^*$. Pick some candidate $w' \in W \cap A_i$, and set $W' = (W \setminus \{w'\}) \cup \{c\}$. Observe that after w' is removed, adding c increases the total PAV-score by at least $(s - 1) \cdot \frac{1}{\ell} + \frac{1}{\ell - 1} > \frac{n}{k}$. Indeed, i approves at most $\ell - 2$ candidates in $W \setminus \{w'\}$ and therefore adding c to $W \setminus \{w'\}$ contributes at least $\frac{1}{\ell - 1}$ to her satisfaction. Thus, the PAV-score of W' is higher than that of W, a contradiction again. $\square$

Interestingly, Theorem 10 does not extend to weight vectors other than $(1, \frac{1}{2}, \frac{1}{3}, \dots)$: our next theorem shows that PAV is essentially the unique **w**-PAV rule that satisfies EJR.

Theorem 11 *For every weight vector* $\mathbf{w}$ *with* $w_1 = 1$, $\mathbf{w} \neq (1, \frac{1}{2}, \frac{1}{3}, \dots)$, *the rule* **w**-*PAV does not satisfy EJR.*

Theorem 11 follows immediately from Lemmas 1 and 2, which are stated below.

Lemma 1 *Consider a weight vector* $\mathbf{w}$ *with* $w_1 = 1$. *If* $w_j > \frac{1}{j}$ *for some* $j > 1$, *then* **w**-*PAV fails JR.*

Proof Suppose that $w_j = \frac{1}{j} + \varepsilon$ for some $j > 1$ and $\varepsilon > 0$. Pick $k > \left\lceil \frac{1}{\varepsilon j} \right\rceil + 1$ so that j divides k; let $t = \frac{k}{j}$. Let $C = C_0 \cup C_1 \cup \dots \cup C_t$, where $C_0 = \{c\}$, $|C_1| = \dots = |C_t| = j$, and the sets $C_0, C_1, \dots, C_t$ are pairwise disjoint. Note that $|C| = tj + 1 = k + 1$. Also, construct $t + 1$ pairwise disjoint groups of voters $N_0, N_1, \dots, N_t$ so that $|N_0| = k$, $|N_1| = \dots = |N_t| = j(k-1)$, and for each $i = 0, 1, \dots, t$ the voters in N_i approve the candidates in C_i only. Observe that the total number of voters is given by $n = k + tj(k-1) = k^2$.

We have $|N_0| = k = \frac{n}{k}$, so every committee that provides justified representation for this profile must elect c. However, we claim that **w**-PAV elects all candidates in $C \backslash \{c\}$ instead. Indeed, if we replace an arbitrary candidate in $C \backslash \{c\}$ with c, then under **w**-PAV the total score of our committee changes by

$$k - j(k-1) \cdot \left(\frac{1}{j} + \varepsilon\right) = 1 - j(k-1)\varepsilon < 1 - j\varepsilon \left\lceil \frac{1}{\varepsilon j} \right\rceil \leq 0,$$

i.e. $C \backslash \{c\}$ has a strictly higher score than any committee that includes c. □

Lemma 2 *Consider a weight vector* $\mathbf{w}$ *with* $w_1 = 1$. *If* $w_j < \frac{1}{j}$ *for some* $j > 1$, *then* **w**-*PAV fails* j-*JR.*

Proof Suppose that $w_j = \frac{1}{j} - \varepsilon$ for some $j > 1$ and $\varepsilon > 0$. Pick $k > j + \left\lceil \frac{1}{\varepsilon} \right\rceil$. Let $C = C_0 \cup C_1$, where $|C_0| = j$, $C_1 = \{c_1, \dots, c_{k-j+1}\}$ and $C_0 \cap C_1 = \emptyset$. Note that $|C| = k + 1$. Also, construct $k - j + 2$ pairwise disjoint groups of voters $N_0, N_1, \dots, N_{k-j+1}$ so that $|N_0| = j(k-j+1)$, $|N_1| = \dots = |N_{k-j+1}| = k - j$, the voters in N_0 approve the candidates in C_0 only, and for each $i = 1, \dots, k-j+1$ the voters in N_i approve c_i only. Note that the number of voters is given by $n = j(k-j+1) + (k-j+1)(k-j) = k(k-j+1)$.

We have $\frac{n}{k} = k - j + 1$ and $|N_0| = j \cdot \frac{n}{k}$, so every committee that provides EJR must select all candidates in C_0. However, we claim that **w**-PAV elects all candidates from C_1 and $j - 1$ candidates from C_0 instead. Indeed, let c be some candidate in C_0, let c' be some candidate in C_1, and let $W = C \backslash \{c\}$, $W' = C \backslash \{c'\}$. The difference between the total score of W and that of W' is

$$j(k-j+1)\left(\frac{1}{j} - \varepsilon\right) - (k-j) < 1 - j \cdot \frac{1}{\varepsilon} \cdot \varepsilon < 1 - j < 0,$$

i.e. **w**-PAV assigns a higher score to W. As this argument does not depend on the choice of c in C_0 and c' in C_1, the proof is complete. □

5.2 JR, EJR and core stability

One can view (extended) justified representation as a stability condition, by associating committees that provide JR/EJR with outcomes of a certain NTU game that are resistant to certain types of deviations.

Specifically, given a pair $(\mathbf{A}, k)$, where $\mathbf{A} = (A_1, \ldots, A_n)$, we define an NTU game $\mathcal{G}(\mathbf{A}, k)$ with the set of players N as follows. We assume that each coalition of size x, $\ell\frac{n}{k} \le x < (\ell+1)\frac{n}{k}$, where $\ell \in \{1, \ldots, k\}$, can "purchase" ℓ alternatives. Moreover, each player evaluates a committee of size ℓ, $\ell \in \{1, \ldots, k\}$, using the PAV utility function, i.e. i derives a utility of $1 + \frac{1}{2} + \cdots + \frac{1}{j}$ from a committee that contains exactly j of her approved alternatives (the argument goes through for $\mathbf{w}$-PAV utilities, as long as $w_1 \ge \cdots \ge w_k > 0$). Thus, for each coalition S with $\ell\frac{n}{k} \le |S| < (\ell+1)\frac{n}{k}$ a payoff vector $\mathbf{x} \in \mathbb{R}^n$ is considered to be feasible for S if and only if there exists a committee $W \subseteq C$ with $|W| \le \ell$ such that $x_i = u_i(W)$ for each $i \in S$, where $u_i(W) = 1 + \cdots + \frac{1}{|A_i \cap W|}$. We denote the set of all payoff vectors that are feasible for a coalition $S \subseteq N$ by $V(S)$.

We say that a coalition $S \subseteq N$ has a *profitable deviation* from a payoff vector $\mathbf{x} \in V(N)$ if there exists a payoff vector $\mathbf{y} \in V(S)$ such that $y_i > x_i$ for all $i \in S$. A payoff vector $\mathbf{x}$ is *stable* if it is feasible for N and no coalition $S \subseteq N$ has a profitable deviation from it; the set of all stable payoff vectors is the *core* of $\mathcal{G}(\mathbf{A}, k)$.

The following theorem describes the relationship between JR, EJR, and profitable deviations in $\mathcal{G}(\mathbf{A}, k)$.

Theorem 12 *A committee W, $|W| = k$, provides justified representation for $(\mathbf{A}, k)$ if and only if no coalition of size $\lceil\frac{n}{k}\rceil$ or less has a profitable deviation from the payoff vector $\mathbf{x}$ associated with W. Moreover, W provides extended justified representation for $(\mathbf{A}, k)$ if and only if for every integer $\ell \ge 0$ no coalition N^* with $\ell \cdot \frac{n}{k} \le |N^*| < (\ell+1) \cdot \frac{n}{k}$, $|\cap_{i\in N^*} A_i| \ge \ell$ has a profitable deviation from $\mathbf{x}$.*

Proof Suppose that W fails to provide justified representation for $(\mathbf{A}, k)$, i.e. there exists a set of voters N^*, $|N^*| = \lceil\frac{n}{k}\rceil$, who all approve some candidate $c \notin W$, but none of them approves any of the candidates in W. Then we have $x_i = 0$ for each $i \in N^*$, and players in N^* can successfully deviate: the payoff vector $\mathbf{y}$ that is associated with the committee $\{c\}$ is feasible for N^* and satisfies $y_i = 1$ for each $i \in N^*$.

Conversely, suppose that W provides justified representation for $(\mathbf{A}, k)$, and consider a coalition N^*. If $|N^*| < \lceil\frac{n}{k}\rceil$, then for every $\mathbf{y} \in V(N^*)$ we have $y_i = 0$ for all $i \in N^*$, so N^* cannot profitably deviate. On the other hand, if $|N^*| = \lceil\frac{n}{k}\rceil$, then every payoff vector $\mathbf{y} \in V(N^*)$ is associated with a committee of size 1. Hence, for every $\mathbf{y} \in V(N^*)$ we have $y_i \le 1$ for all $i \in N^*$, and if $y_i = 1$ for all $i \in N^*$, then $\cap_{i\in N^*} A_i \ne \emptyset$, and therefore, since W provides JR, we have $x_i \ge 1$ for some $i \in N^*$. Thus, N^* has no profitable deviation.

For EJR the argument is similar. If W fails to provide extended justified representation for $(\mathbf{A}, k)$, there exists an $\ell > 0$ and a set of voters N^*, $|N^*| \ge \ell \cdot \frac{n}{k}$, such that $|\bigcap_{i\in N^*} A_i| \ge \ell$, but $|A_i \cap W| < \ell$ for each $i \in N^*$. Then we have $x_i < 1 + \cdots + \frac{1}{\ell}$ for each $i \in N^*$, and players in N^* can successfully deviate: if S is a committee that consists of some ℓ candidates in $\bigcap_{i \in N^*} A_i$, then the payoff vector $\mathbf{y}$ that is associated with S is feasible for N^* and satisfies $y_i = 1 + \cdots + \frac{1}{\ell}$ for each $i \in N^*$.

Conversely, suppose that W provides extended justified representation for $(\mathbf{A}, k)$, and consider some $\ell \geq 0$ and some coalition N^* with $\ell \cdot \frac{n}{k} \leq |N^*| < (\ell+1)\frac{n}{k}$. We have argued above that if $\ell = 0$, then N^* cannot profitably deviate. Thus, assume $\ell > 0$. Every payoff vector $\mathbf{y} \in V(N^*)$ is associated with a committee of size ℓ. Hence, for every $\mathbf{y} \in V(N^*)$ we have $y_i \leq 1 + \cdots + \frac{1}{\ell}$ for all $i \in S$, and if $y_i = 1 + \cdots + \frac{1}{\ell}$ for all $i \in N^*$, then $|\cap_{i \in N^*} A_i| \geq \ell$. Since W provides EJR, we have $x_i \geq 1 + \cdots + \frac{1}{\ell}$ for some $i \in N^*$. Again, this implies that N^* has no profitable deviation. □

The second part of Theorem 12 considers deviations by cohesive coalitions. The reader may wonder if it can be strengthened to *arbitrary* coalitional deviations, i.e. whether a committee provides EJR *if and only if* the associated payoff vector is in the core of $\mathcal{G}(\mathbf{A}, k)$. The following example shows that this is not the case.

Example 6 Let $k = 10$, $C = \{x_1, x_2, \ldots, x_{10}, y, z\}$, $N = \{1, 2, \ldots, 20\}$, and

$$\begin{aligned} A_1 &= A_2 = A_3 = \{x_1, y\}, \\ A_4 &= A_5 = A_6 = \{x_1, z\}, \\ A_7 &= \ldots = A_{20} = \{x_2, \ldots, x_{10}\}. \end{aligned}$$

Then PAV outputs the committee $W = \{x_1, x_2, \ldots, x_{10}\}$ for $(\mathbf{A}, k)$; in particular, W provides EJR for $(\mathbf{A}, k)$. However, the associated payoff vector $\mathbf{x}$ is not in the core, as the players in $\{1, 2, 3, 4, 5, 6\}$, a coalition of size $3\frac{n}{k}$, can successfully deviate: the payoff vector associated with $\{x_1, y, z\}$ is feasible for $\{1, 2, 3, 4, 5, 6\}$ and provides a higher payoff than $\mathbf{x}$ to each of the first six players. We remark that the core of $\mathcal{G}(\mathbf{A}, k)$ is not empty: in particular, it contains the payoff vector associated with $\{x_1, \ldots, x_8, y, z\}$.

It remains an open question whether the core of $\mathcal{G}(\mathbf{A}, k)$ is non-empty for every pair $(\mathbf{A}, k)$. Further, while it would be desirable to have a voting rule that outputs a committee whose associated payoff vector is in the core whenever the core is not empty, we are not aware of any such rule: every voting rule that fails EJR also fails this more demanding criterion, and Example 6 illustrates that PAV fails this criterion as well.

5.3 Computational issues

In Sect. 3 we have argued that it is easy to find a committee that provides JR for a given ballot profile, and to check whether a specific committee provides JR. In contrast, for EJR these questions appear to be computationally difficult. Specifically, we were unable to design an efficient algorithm for computing a committee that provides EJR; while PAV is guaranteed to find such a committee, computing its output is NP-hard. We remark, however, that when ℓ is bounded by a constant, we can efficiently compute a committee that provides ℓ-JR.

Theorem 13 *A committee that provides ℓ-JR can be computed in time polynomial in n and $|C|^{\ell}$.*

Proof Consider the following greedy algorithm, which we will refer to as ℓ-GreedyAV. We start by setting $C' = C$, $\mathbf{A}' = \mathbf{A}$, and $W = \emptyset$. As long as $|W| \le k - \ell$, we check if there exists a set of candidates $\{c_1, \dots, c_\ell\} \subseteq C'$ that is unanimously approved by at least $\ell \frac{n}{k}$ voters in $\mathbf{A}'$ (this can be done in time $n \cdot |C|^{\ell+1}$). If such a set exists, we set $W := W \cup \{c_1, \dots, c_\ell\}$ and we remove from $\mathbf{A}'$ all ballots A_i such that $|A_i \cap W| \ge \ell$ (note that this includes all ballots A_i with $\{c_1, \dots c_\ell\} \subseteq A_i$). If at some point we have $|W| \le k - \ell$ and no set $\{c_1, \dots, c_\ell\}$ satisfies our criterion or $|W| > k - \ell$, we add an arbitrary subset of $k - |W|$ candidates from C' to W and return W; if this does not happen, we terminate after having picked k candidates.

Suppose for the sake of contradiction that for some profile $\mathbf{A} = (A_1, \dots, A_n)$ and some $k > 0$, ℓ-GreedyAV outputs a committee that does not provide ℓ-JR for $(\mathbf{A}, k)$. Then there exists a set $N^* \subseteq N$ with $|N^*| \ge \ell \frac{n}{k}$ such that $|\bigcap_{i \in N^*} A_i| \ge \ell$ and, when ℓ-GreedyAV terminates, every ballot A_i such that $i \in N^*$ is still in $\mathbf{A}'$. Consider some subset of candidates $\{c_1, \dots, c_\ell\} \subseteq \bigcap_{i \in N^*} A_i$. At every point in the execution of ℓ-GreedyAV this subset is unanimously approved by at least $|N^*| \ge \ell \frac{n}{k}$ ballots in $\mathbf{A}'$. As at least one of $\{c_1, \dots, c_\ell\}$ was not elected, at every stage the algorithm selected a set of ℓ candidates that was approved by at least $\ell \frac{n}{k}$ ballots (until more than $k - \ell$ candidates were selected). Since at the end of each stage the algorithm removed from $\mathbf{A}'$ all ballots containing the candidates that had been added to W at that stage, it follows that altogether the algorithm has removed at least $\lfloor \frac{k}{\ell} \rfloor \cdot \ell \frac{n}{k} > \left(\frac{k}{\ell} - 1\right) \cdot \ell \frac{n}{k} = n - \ell \frac{n}{k}$ ballots from $\mathbf{A}'$. This is a contradiction, since we assumed that, when the algorithm terminates, the $\ell \frac{n}{k}$ ballots $(A_i)_{i \in N^*}$ are still in $\mathbf{A}'$. □

For the problem of checking whether a given committee provides EJR for a given input, we can establish a formal hardness result.

Theorem 14 *Given a ballot profile* $\mathbf{A}$*, a target committee size* k*, and a committee* W*,* $|W| = k$*, it is* coNP*-complete to check whether* W *provides EJR for* $(\mathbf{A}, k)$.

Proof It is easy to see that this problem is in coNP: to show that W does not provide EJR for $(\mathbf{A}, k)$, it suffices to guess an integer $\ell > 0$ and a set of voters N^* of size at least $\ell \cdot \frac{n}{k}$ such that $|\bigcap_{i \in N^*} A_i| \ge \ell$, but $|A_i \cap W| < \ell$ for all $i \in N^*$.

To prove coNP-completeness, we reduce the classic Balanced Biclique problem (Garey and Johnson 1979, [GT24]) to the complement of our problem. An instance of Balanced Biclique is given by a bipartite graph (L, R, E) with parts L and R and edge set E, and an integer ℓ; it is a "yes"-instance if we can pick subsets of vertices $L' \subseteq L$ and $R' \subseteq R$ so that $|L'| = |R'| = \ell$ and $(u, v) \in E$ for each $u \in L'$, $v \in R'$; otherwise, it is a "no"-instance.

Given an instance $\langle (L, R, E), \ell \rangle$ of Balanced Biclique with $R = \{v_1, \dots, v_s\}$, we create an instance of our problem as follows. Assume without loss of generality that $s \ge 3$, $\ell \ge 3$. We construct 4 pairwise disjoint sets of candidates C_0, C_1, C_1', C_2, so that $C_0 = L$, $|C_1| = |C_1'| = \ell - 1$, $|C_2| = s\ell + \ell - 3s$, and set $C = C_0 \cup C_1 \cup C_1' \cup C_2$. We then construct 3 sets of voters N_0, N_1, N_2, so that $N_0 = \{1, \dots, s\}$, $|N_1| = \ell(s - 1)$, $|N_2| = s\ell + \ell - 3s$ (note that $|N_2| > 0$ as we assume that $\ell \ge 3$). For each $i \in N_0$ we set $A_i = \{u_j \mid (u_j, v_i) \in E\} \cup C_1$, and for each $i \in N_1$ we set $A_i = C_0 \cup C_1'$. The candidates in C_2 are matched to voters in N_2: each voter in N_2 approves exactly one candidate in C_2, and each candidate in C_2 is approved by exactly one voter in N_2.

Denote the resulting list of ballots by $\mathbf{A}$. Finally, we set $k = 2\ell - 2$, and let $W = C_1 \cup C'_1$. Note that the number of voters n is given by $s + \ell(s-1) + s\ell + \ell - 3s = 2s(\ell - 1)$, so $\frac{n}{k} = s$.

Suppose first that we started with a "yes"-instance of Balanced Biclique, and let (L', R') be the respective ℓ-by-ℓ biclique. Let $C^* = L'$, $N^* = N_1 \cup \{i \in N_0 \mid v_i \in R'\}$. Then $|N^*| = \ell s$, all voters in N^* approve all candidates in C^*, $|C^*| = \ell$, but each voter in N^* is only represented by $\ell - 1$ candidates in W. Hence, W fails to provide ℓ-justified representation for $(\mathbf{A}, k)$.

Conversely, suppose that W fails to provide EJR for $(\mathbf{A}, k)$. That is, there exists a value $j > 0$, a set N^* of js voters and a set C^* of j candidates so that all voters in N^* approve of all candidates in C^*, but for each voter in N^* at most j of her approved candidates are in W. Note that, since $s > 1$, we have $N^* \cap N_2 = \emptyset$. Further, each voter in $N \setminus N_2$ is represented by $\ell - 1$ candidates in W, so $j \geq \ell$. As $N^* = js \geq \ell s \geq s$, it follows that $|N^* \cap N_0| \geq \ell$, $|N^* \cap N_1| > 0$. Since N^* contains voters from both N_0 and N_1, it follows that $C^* \subseteq C_0$. Thus, there are at least ℓ voters in $N^* \cap N_0$ who approve the same $j \geq \ell$ candidates in C_0; any set of ℓ such voters and ℓ such candidates corresponds to an ℓ-by-ℓ biclique in the input graph. □

6 Variants of justified representation

The definition of JR requires that if there is a group of $\lceil \frac{n}{k} \rceil$ voters who jointly approve some candidate, then the elected committee has to contain at least one candidate approved by *some* member of this group. This condition may appear to be too weak; it may seem more natural to require that *every* group member approves some candidate in the committee, or—stronger yet—that the committee contains at least one candidate approved by *all* group members. This intuition is captured by the following definitions.

Definition 3 Given a ballot profile $\mathbf{A} = (A_1, \ldots, A_n)$ and a target committee size k, we say that a committee W of size k provides

- *semi-strong justified representation for* $(\mathbf{A}, k)$ if for each group $N^* \subseteq N$ with $|N^*| \geq \frac{n}{k}$ and $\bigcap_{i \in N^*} A_i \neq \emptyset$ it holds that $W \cap A_i \neq \emptyset$ for all $i \in N^*$.
- *strong justified representation for* $(\mathbf{A}, k)$ if for each group $N^* \subseteq N$ with $|N^*| \geq \frac{n}{k}$ and $\bigcap_{i \in N^*} A_i \neq \emptyset$ it holds that $W \cap (\cap_{i \in N^*} A_i) \neq \emptyset$.

By definition, a committee providing strong justified representation also provides semi-strong justified representation, and a committee providing semi-strong justified representation also provides (standard) justified representation.

However, it turns out that satisfying these stronger requirements is not always feasible: there are ballot profiles for which no committee provides semi-strong justified representation.

Example 7 Let $k = 3$ and consider the following profile with $C = \{a, b, c, d\}$ and $n = 9$.

$$A_1 = A_2 = \{a\} \qquad A_3 = \{a, b\} \qquad A_4 = \{b\} \qquad A_5 = \{b, c\}$$

$$A_6 = \{c\} \qquad A_7 = \{c, d\} \qquad A_8 = A_9 = \{d\}$$

For each candidate $x \in C$, there are $\frac{n}{k} = 3$ voters such that $\cap_i A_i = \{x\}$, and at least one of those voters has $A_i = \{x\}$. Thus, a committee that satisfies semi-strong justified representation would have to contain all four candidates, which is impossible.

While Example 7 shows that no approval-based voting rule can always find a committee that provides strong or semi-strong justified representation, it may be interesting to identify voting rules that output such committees whenever they exist.

Finally, we remark that strong justified representation does not imply EJR.

Example 8 Let $C = \{a, b, c, d, e\}$, $n = 4$, $k = 4$, and consider the following ballot profile.

$$A_1 = \{a, b\} \qquad A_2 = \{a, b\} \qquad A_3 = \{c\} \qquad A_4 = \{d, e\}$$

EJR requires that we choose both a and b, but $\{a, c, d, e\}$ provides strong justified representation.

7 Related work

It is instructive to compare JR and EJR to alternative approaches towards fair representation, such as *representativeness* (Duddy 2014) and *proportional justified representation* (Sánchez-Fernández et al. 2016, 2017).

Duddy (2014) proposes the notion of *representativeness*, which applies to probabilistic voting rules. The property Duddy proposes is incomparable with JR: in situations he considers ($k = 2$, n voters approve x, $n + 1$ voters approve y and z), JR requires that one of y and z should be selected, whereas Duddy requires x to be selected with positive probability. Both are reasonable requirements, but they address different concerns. Duddy's axiom say nothing about situations where voters are split equally (say, n voters approve $\{x, y\}$, n voters approve $\{z, t\}$), whereas JR requires that each voter is represented. Another obvious difference is that he allows for randomized rules.

Very recently (after the conference version of our paper was published), Sánchez-Fernández et al. (2016, 2017) came up with the notion of *proportional justified representation* (PJR), which can be seen as an alternative to EJR. A committee is said to provide PJR for a ballot profile $(A_1, \ldots, A_n)$ over a candidate set C and a target committee size k if, for every positive integer ℓ, $\ell \le k$, there does not exist a set of voters $N^* \subseteq N$ with $|N^*| \ge \ell \cdot \frac{n}{k}$ such that $|\bigcap_{i \in N^*} A_i| \ge \ell$, but $|(\bigcup_{i \in N^*} A_i) \cap W| < \ell$. In contrast to EJR, the PJR condition does not require one of the voters in N^* to have ℓ representatives. Rather, a committee provides PJR as long as it contains ℓ candidates that are approved by (possibly different) voters in N^*, for every group N^* satisfying the size and cohesiveness constraints. An attractive feature of PJR is that it is compatible with the idea of *perfect representation*: a committee W provides perfect representation for a group of n voters and a target committee size k if $n = ks$ for some positive integer s and the voters can be split into k pairwise disjoint groups $N_1, \ldots, N_k$ of size s each in such a way that there is a one-to-one mapping $\mu : W \to \{N_1, \ldots, N_k\}$

Table 1 Satisfaction of JR and EJR and computational complexity of approval-based voting rules; the superscript '$*$' indicates that the rule fails the respective axiom for some way of breaking intermediate ties

Rule	JR	EJR	Complexity
AV	–	–	in P
SAV	–	–	in P
MAV	–	–	NP-hard
SeqPAV	–	–	in P
GreedyAV	✓	–	in P
HareAV	✓	–*	in P
MonroeAV	✓	–	NP-hard
PAV	✓	✓	NP-hard

such that for each candidate $c \in W$ all voters in $\mu(c)$ approve c. Sánchez-Fernández et al. (2016, 2017) prove that every committee that provides perfect representation also provides PJR; in contrast, EJR may rule out all committees that provide perfect representation, as illustrated by Example 4. It is easily seen that PJR is a weaker requirement than EJR, and a stronger one than JR. Interestingly, Sánchez-Fernández et al. (2016, 2017) show that many results that we have established for EJR also hold for PJR: in particular, $\mathbf{w}$-SeqPAV violates PJR for every weight vector $\mathbf{w}$ with $w_1 = 1$, and $\mathbf{w}$-PAV with $w_1 = 1$ satisfies PJR if and only if $\mathbf{w} = (1, \frac{1}{2}, \frac{1}{3}, \ldots)$.

8 Conclusions

We have formulated a desirable property of approval-based committee selection rules, which we called justified representation (JR). While JR is fairly easy to satisfy, it turns out that many well-known approval-based rules fail it. A prominent exception is the PAV rule, which also satisfies a stronger version of this property, namely extended justified representation (EJR). Indeed, EJR characterizes PAV within the class of $\mathbf{w}$-PAV rules, and we are not aware of any other natural voting rule that satisfies EJR irrespective of the tie-breaking rule (of course, we can construct voting rules that differ from PAV, yet satisfy EJR, by modifying the output of PAV on profiles on which EJR places no constraints on the output). Table 1 summarizes the representation properties and computational complexity of approval-based voting rules.

Perhaps the most pressing open question suggested by our work is whether there is an efficient algorithm for finding a committee that provides EJR for a given profile. In particular, we would like to understand whether we can break ties in the execution of HareAV to always produce such a committee, and whether some tie-breaking rule with this property is polynomial-time computable. Also, it would be interesting to see if EJR, in combination with other natural axioms, can be used to axiomatize PAV. Concerning (semi-)strong justified representation, an interesting computational problem is whether there are efficient algorithms for checking the existence of committees satisfying these requirements.

JR and EJR can also be used to formulate new approval-based rules. We mention two types of rules that seem particularly attractive:

The *utilitarian* $(E)JR$ *rule* returns a committee that, among all committees that satisfy $(E)JR$, has the highest AV score.
The *egalitarian* $(E)JR$ *rule* returns a committee that, among all committees that satisfy $(E)JR$, maximizes the number of representatives of the voter who has the least number of representatives in the winning committee.

The computational complexity of winner determination for these rules is an interesting problem. Since PAV is NP-hard to compute, our study also provides additional motivation for the use of approximation and parameterized algorithms to compute PAV outcomes. Finally, analyzing the compatibility of $(E)JR$ with other important properties, such as, e.g., strategyproofness for dichotomous preferences, is another avenue of future research.

Acknowledgements The authors thank the anonymous reviewers of the *Multidisciplinary Workshop on Advances in Preference Handling (MPREF 2014)*, and the *Twenty-ninth AAAI Conference (AAAI 2015)* for their helpful feedback on earlier versions of the paper. We further thank Svante Janson, Martin Lackner, Xavier Mora, and Piotr Skowron for valuable discussions. Brill, Conitzer, and Freeman were supported by NSF and ARO under Grants CCF-1101659, IIS-0953756, IIS-1527434, CCF-1337215, W911NF-12-1-0550, and W911NF-11-1-0332, by a Feodor Lynen research fellowship of the Alexander von Humboldt Foundation, and by COST Action IC1205 on Computational Social Choice. Aziz is supported by a Julius Career Award. Brill and Elkind were partially supported by ERC-StG 639945. Conitzer is supported by a Guggenheim Fellowship. Walsh also receives support from the Asian Office of Aerospace Research and Development (AOARD 124056) and the German Federal Ministry for Education and Research through the Alexander von Humboldt Foundation.

References

Alcalde-Unzu J, Vorsatz M (2009) Size approval voting. J Econ Theory 144(3):1187–1210
Aziz H, Gaspers S, Gudmundsson J, Mackenzie S, Mattei N, Walsh T (2015) Computational aspects of multi-winner approval voting. In: Proceedings of the 14th international conference on autonomous agents and multiagent systems (AAMAS), pp 107–115
Aziz H, Lang J, Monnot J (2016) Computing Pareto optimal committees. In: Proceedings of the 25th international joint conference on artificial intelligence (IJCAI), pp 60–66
Betzler N, Slinko A, Uhlmann J (2013) On the computation of fully proportional representation. J Artif Intell Res 47:475–519
Brams SJ (2010) Preface. In: Laslier JF, Sanver MR (eds), Handbook on approval voting, studies in choice and welfare, Springer, New York, pp vii–ix
Brams SJ, Fishburn PC (2007) Approval voting, 2nd edn. Springer, New York
Brams SJ, Kilgour DM (2014) Satisfaction approval voting. In: Fara R, Leech D, Salles M (eds) Voting power and procedures, studies in choice and welfare. Springer, New York, pp 323–346
Brams SJ, Kilgour DM, Sanver MR (2006) How to elect a representative committee using approval balloting. In: Simeone B, Pukelsheim F (eds) Mathematics and democracy: recent advances in voting systems and collective choice. Springer, New York, pp 83–96
Brams SJ, Kilgour DM, Sanver RM (2007) A minimax procedure for electing committees. Public Choice 132(3–4):401–420
Byrka J, Sornat K (2014) PTAS for minimax approval voting. In: Proceedings of the 10th international conference on web and internet economics (WINE), pp 203–217
Caragiannis I, Kalaitzis D, Markakis E (2010) Approximation algorithms and mechanism design for minimax approval voting. In: Proceedings of the 24th AAAI conference on artificial intelligence (AAAI), pp 737–742
Caragiannis I, Kaklamanis C, Karanikolas N, Procaccia AD (2014) Socially desirable approximations for Dodgson's voting rule. ACM Trans Algorithm 10(2). doi:10.1145/2556950

Chamberlin B, Courant P (1983) Representative deliberations and representative decisions: proportional representation and the Borda rule. Am Polit Sci Rev 77(3):718–733

Cornaz D, Galand L, Spanjaard O (2012) Bounded single-peaked width and proportional representation. In: Proceedings of the 20th European conference on artificial intelligence (ECAI), pp 270–275

Davis M, Orrison ME, Su FE (2014) Voting for committees in agreeable societies. Contemp Math 624:147–157

Droop HR (1881) On methods of electing representatives. J Stat Soc Lond 44(2):141–196

Duddy C (2014) Electing a representative committee by approval ballot: an impossibility result. Econ Lett 124:14–16

Dummett M (1984) Voting procedures. Oxford University Press, Oxford

Elkind E, Lackner M (2015) Structure in dichotomous preferences. In: Proceedings of the 24th international joint conference on artificial intelligence (IJCAI), pp 2019–2025

Elkind E, Faliszewski P, Skowron P, Slinko A (2017) Properties of multiwinner voting rules. Soc Choice Welf (**forthcoming**)

Elkind E, Lang J, Saffidine A (2015) Condorcet winning sets. Soc Choice Welf 44(3):493–517

Endriss U (2013) Sincerity and manipulation under approval voting. Theor Decis 74(4):335–355

Fishburn PC, Pekec AS (2004) Approval voting for committees: threshold approaches. Techn Rept

Garey M, Johnson D (1979) Computers and intractability: a guide to the theory of NP-completeness. W. H Freeman and Company, New York

Kilgour DM (2010) Approval balloting for multi-winner elections. In: Laslier JF, Sanver MR (eds) Handbook on approval voting, chapter 6, Springer, New York, pp 105–124

Kilgour DM, Marshall E (2012) Approval balloting for fixed-size committees. In: Felsenthal DS, Machover M (eds), Electoral systems, studies in choice and welfare, Springer, New York, pp 305–326

Laslier J-F, Sanver MR (eds) (2010) Handbook on approval voting. Studies in choice and welfare. Springer, New York

LeGrand R, Markakis E, Mehta A (2007) Some results on approximating the minimax solution in approval voting. In: Proceedings of the 6th international conference on autonomous agents and multi-agent systems (AAMAS), pp 1185–1187

Lu T, Boutilier C (2011) Budgeted social choice: from consensus to personalized decision making. In: Proceedings of the 22nd international joint conference on artificial intelligence (IJCAI), AAAI Press, Cambridge, pp 280–286

Meir R, Procaccia AD, Rosenschein JS, Zohar A (2008) Complexity of strategic behavior in multi-winner elections. J Artif Intell Res 33:149–178

Misra N, Nabeel A, Singh H (2015) On the parameterized complexity of minimax approval voting. In: Proceedings of the 14th international conference on autonomous agents and multiagent systems (AAMAS), pp 97–105

Monroe B (1995) Fully proportional representation. Am Polit Sci Rev 89(4):925–940

Procaccia A, Rosenschein J, Zohar A (2008) On the complexity of achieving proportional representation. Soc Choice Welf 30(3):353–362

Sánchez-Fernández L, Fernández N, Fisteus JA, Basanta-Val P (2016) Some notes on justified representation. In: Proceedings of the 10th multidisciplinary workshop on advances in preference handling (MPREF)

Sánchez-Fernández L, Elkind E, Lackner M, Fernández N, Fisteus JA, Basanta-Val P, Skowron P (2017) Proportional justified representation. In: Proceedings of the 31st AAAI conference on artificial intelligence (AAAI)

Skowron P, Faliszewski P, Slinko A (2015a) Achieving fully proportional representation: approximability results. Artif Intell 222:67–103

Skowron P, Yu L, Faliszewski P, Elkind E (2015b) The complexity of fully proportional representation for single-crossing electorates. Theor Comput Sci 569:43–57

Skowron P, Faliszewski P, Lang J (2016) Finding a collective set of items: from proportional multirepresentation to group recommendation. Artif Intell 241:191–216

Thiele TN (1895) Om flerfoldsvalg. In: Oversigt over det Kongelige Danske Videnskabernes Selskabs Forhandlinger, pp 415–441

Tideman N (1995) The single transferable vote. J Econ Perspect 9(1):27–38

Tideman N (2006) Collective decisions and voting. Ashgate, Surrey

Tideman N, Richardson D (2000) Better voting methods through technology: the refinement-manageability trade-off in the single transferable vote. Public Choice 103(1–2):13–34