Justified Representation in Approval-Based Committee Voting

Haris Aziz
NICTA and UNSW,
Sydney 2033, Australia

Markus Brill
Duke University
Durham, NC 27708, USA

Vincent Conitzer
Duke University
Durham, NC 27708, USA

Edith Elkind
University of Oxford
Oxford OX1 3QD, UK

Rupert Freeman
Duke University
Durham, NC 27708, USA

Toby Walsh
NICTA and UNSW,
Sydney 2033, Australia

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. We then check if this axiom is fulfilled by well-known approval-based voting rules. We show that the answer is negative for most of the rules we consider, with notable exceptions of PAV (Proportional Approval Voting), an extreme version of RAV (Reweighted Approval Voting), and, for a restricted preference domain, MAV (Minimax Approval Voting). We then introduce a stronger version of the JR axiom, which we call *extended justified representation (EJR)*, and show that PAV satisfies EJR, while other rules do not. We also consider several other questions related to JR and EJR, including the relationship between JR/EJR and unanimity, and the complexity of the associated algorithmic problems.

1 Introduction

Aggregation of preferences is a central problem in the field of social choice, and has received a considerable amount of attention from the artificial intelligence research community (see e.g., Conitzer 2010). 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 (committee)*: this includes domains such as parliamentary elections, the hiring of faculty members, or (automated) agents deciding on a set of plans (Elkind, Lang, and Saffidine 2014; LeGrand, Markakis, and Mehta 2007; Davis, Orrison, and Su 2014). The study of algorithmic complexity of voting rules that output committees is an active research direction (see, e.g., Procaccia, Rosenschein, and Zohar 2008; Meir, Procaccia, and Rosenschein 2008; Caragiannis, Kalaitzis, and Markakis 2010; Lu and Boutilier 2011; Betzler, Slinko, and Uhlmann 2013; Skowron, Faliszewski, and Slinko 2013; Cornaz, Galand, and Spanjaard 2012; Skowron et al. 2013).

Much of the prior work in AI on multi-winner rules focuses on the setting where voters' preferences are total orders of the candidates; notable exceptions are (LeGrand, Markakis, and Mehta 2007) and (Caragiannis, Kalaitzis, and Markakis 2010). In contrast, 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. One of the advantages of approval ballots is their simplicity: such 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, Kilgour, and Sanver 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 winners, 51% of the agents approve the same k candidates, and the remaining agents approve a disjoint set of k candidates, then the agents 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., Kilgour 2010). Under *Proportional Approval Voting (PAV)*, each agent's contribution to the committee's total score depends on how many candidates from the agent's approval set have been elected. A sequential variant of this rule is known as *Reweighted Approval Voting (RAV)*. Another way to modulate the approvals is through computing a satisfaction score for each agent 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. 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 more equitable representation in multi-winner elections (Brams and Kilgour 2014).

Based on their relative merits, approval-based multi-winner rules have been examined in great detail in both economics and computer science in recent years (Brams and Fishburn 2007; LeGrand, Markakis, and Mehta 2007; Meir, Procaccia, and Rosenschein 2008). The Handbook of Approval Voting discusses various approval-based multi-winner rules including PAV, RAV, SAV and MAV (Kilgour 2010). However, there has been limited axiomatic analysis of these rules from the perspective of representation.

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. 2014), 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 very few of the existing multi-winner approval-based rules satisfy it. Specifically, we demonstrate that AV, SAV, MAV and the standard variant of RAV do not satisfy JR. On the positive side, JR is satisfied by PAV and some of its variants, as well as an extreme variant of RAV. 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 very 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 PAV satisfies EJR. Moreover, it is computationally hard to check whether a given committee provides EJR. We conclude the paper by showing how JR can be used to formulate other attractive approval-based multi-winner rules, and by identifying several directions for future work. Some proofs are omitted due to space constraints, and can be found in the full version of this paper (Aziz et al. 2014).

2 Preliminaries

We consider a social choice setting with a set of agents (voters) $N = \{1, \dots, n\}$ and a set of candidates $C = \{c_1, \dots, c_m\}$. Each agent $i \in N$ submits an approval ballot $A_i \subseteq C$, which represents the subset of candidates that she approves of. We refer to the list $\mathcal{A} = (A_1, \dots, A_n)$ of approval ballots as the *ballot profile*. We will consider *approval-based multi-winner voting rules* that take as input $(N, C, \mathcal{A}, k)$, 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 such 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 a fixed priority order over size-k subsets; however, most of our results do not depend on the tie-breaking rule.

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 $|\{i \mid c \in A_i\}|$, and AV outputs a set W of size k that maximizes $\sum_{c \in W} |\{i \mid c \in A_i\}|$. AV has been adopted by several academic and professional societies such as the Institute of Electrical and Electronics Engineers (IEEE) and the International Joint Conference on Artificial Intelligence.

Satisfaction Approval Voting (SAV) An agent's *satisfaction score* is the fraction of her approved candidates that are elected. SAV maximizes the sum of agents' satisfaction scores. Formally, SAV finds 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).

Proportional Approval Voting (PAV) Under PAV, an agent is assumed to derive a utility of $1 + \frac{1}{2} + \frac{1}{3} + \dots + \frac{1}{j}$ from a committee that contains exactly j of her approved candidates, and the goal is to maximize the sum of the agents' 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. PAV was proposed by mathematician Forest Simmons in 2001, and captures the idea of diminishing returns—an individual agent's preferences should count less the more she is satisfied. It has recently been shown that computing PAV is NP-hard (Aziz et al. 2014; Skowron, Faliszewski, and Lang 2015). We can generalize the definition of PAV by using an arbitrary non-increasing score vector in place of $(1, \frac{1}{2}, \frac{1}{3}, \dots)$: for every vector $\mathbf{w} = (w_1, w_2, \dots)$, where $w_1, w_2, \dots$ are non-negative reals[1], $w_1 = 1$ and $w_1 \geq w_2 \geq \dots$, we define a voting rule $\mathbf{w}$-PAV that, given a ballot profile $(A_1, \dots, A_n)$ and a target number of winners k, 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 $r_{\mathbf{w}}(p) = \sum_{j=1}^{p} w_j$.

Reweighted Approval Voting (RAV) RAV converts AV into a multi-round rule, by selecting a candidate in each round and then reweighing the approvals for the subsequent rounds. Specifically, it starts by setting $W = \emptyset$. Then in round j, $j = 1, \dots, k$, it computes the *approval weight* of each candidate c 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. RAV

[1] It is convenient to think of $\mathbf{w}$ as an infinite vector; however, 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 algorithmic properties of such rules in this paper.

was invented by the Danish polymath Thorvald Thiele in the early 1900's. RAV has also been referred to as "*sequential proportional AV*" (Brams and Kilgour 2014), and was used briefly in Sweden during the early 1900's. Just as for PAV, we can extend the definition of RAV to score vectors other than $(1, \frac{1}{2}, \frac{1}{3}, \dots)$: every vector $\mathbf{w} = (w_1, w_2, \dots)$ with $w_1 = 1$ and $w_1 \geq w_2 \geq \dots$ defines a sequential voting rule $\mathbf{w}$-RAV, which proceeds as RAV, 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.

Minimax Approval Voting (MAV) MAV returns a committee W that minimizes the maximum *Hamming distance* between W and the agents' ballots. Formally, let $d(Q, T) = |Q \setminus T| + |T \setminus Q|$ and 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. Minimax approval voting was proposed by Brams, Kilgour, and Sanver (2007). Computing the outcome of MAV is known to be NP-hard (LeGrand, Markakis, and Mehta 2007).

3 Justified Representation

We will now define the main concept of this paper.

Definition 1 (Justified representation (JR)) *Given a ballot profile $\mathcal{A} = (A_1, \dots, A_n)$ over a candidate set C and a target committee size k, $k \leq |C|$, we say that a set of candidates W of size $|W| = k$* provides justified representation *for $(\mathcal{A}, k)$ 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 $\mathcal{A} = (A_1, \dots, A_n)$ and every target committee size k it outputs a winning set that provides justified representation for $(\mathcal{A}, k)$.*

The intuition behind this definition is that if k candidates are to be selected, then a set of $\frac{n}{k}$ voters that are completely unrepresented can demand that at least one of their unanimously approved candidates should be selected.

3.1 Existence and Computational Properties

We start our analysis of justified representation by observing that, for every ballot profile $\mathcal{A}$ and every value of k, there is a committee that provides justified representation for $(\mathcal{A}, k)$, and, moreover, such a committee can be computed efficiently given the voters' ballots.

To see this, consider the following greedy algorithm, which we will refer to as *Greedy Approval Voting* (GAV). We start by setting $C' = C$, $\mathcal{A}' = \mathcal{A}$, and $W = \emptyset$. As long as $|W| < k$ and $\mathcal{A}'$ is non-empty, we pick a candidate $c \in C'$ that has the highest approval score with respect to $\mathcal{A}'$, and set $W := W \cup \{c\}$, $C' := C' \setminus \{c\}$. Also, we remove from $\mathcal{A}'$ all ballots A_i such that $c \in A_i$. If at some point we have $|W| < k$ and $\mathcal{A}'$ 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. Observe that this algorithm runs in polynomial time. We will now show that it satisfies JR.

Theorem 1 *GAV satisfies JR.*

Proof: Suppose for the sake of contradiction that for some ballot profile $\mathcal{A} = (A_1, \dots, A_n)$ and some $k > 0$, GAV outputs a committee that does not provide justified representation for $(\mathcal{A}, k)$. 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 GAV terminates, every ballot A_i such that $i \in N^*$ is still in $\mathcal{A}'$. Consider some candidate $c \in \bigcap_{i \in N^*} A_i$. At every point in the execution of GAV, 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. Since at the end of each stage the algorithm removed from $\mathcal{A}'$ all ballots containing the candidate added to W at that stage, altogether the algorithm has removed at least $k \cdot \frac{n}{k}$ ballots from $\mathcal{A}'$. This contradicts the assumption that $\mathcal{A}'$ 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 $(\mathcal{A}, k)$ if and only if there exists a candidate c with $s(c) \geq \frac{n}{k}$.

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

3.2 JR and Unanimity

Another desirable property of approval-based voting rules is *unanimity*: we say that an approval-based rule is *unanimous* if, given a ballot profile $(A_1, \dots, A_n)$ with $\bigcap_{i \in N} A_i \neq \emptyset$ and a target committee of size k, it outputs a winning set W, $|W| = k$, such that $W \cap \bigcap_{i \in N} A_i \neq \emptyset$. While unanimity may appear to be similar to JR, the two properties are essentially unrelated. Specifically, for $k = 1$ unanimity implies JR, but for $k > 1$ this is not the case; JR does not imply unanimity either, even for $k = 1$. Examples showing this can be found in the full version of the paper.

4 JR 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, and RAV. We will also identify conditions on $\mathbf{w}$ that are sufficient/necessary for $\mathbf{w}$-PAV and $\mathbf{w}$-RAV 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 considered rules 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: We omit the proof of the first statement due to space restrictions. For $k \geq 3$, we let $C = \{c_0, c_1, \dots, c_k\}$, $n = k$, and consider the profile where the first voter approves c_0, whereas each of the remaining voters approves all of $c_1, \dots, c_k$. JR requires c_0 to be selected, but AV selects $\{c_1, \dots, c_k\}$. □

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 $X = \{x_1, \dots, x_k, x_{k+1}\}$, $Y = \{y_1, \dots, y_k\}$, $C = X \cup Y$, and consider the profile $(A_1, \dots, A_k)$, where $A_1 = X$, $A_2 = \{y_1, y_2\}$, $A_i = \{y_i\}$ for $i = 3, \dots, k$. 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 $X = \{x_1, \dots, x_k\}$, $Y = \{y_1, \dots, y_k\}$, $C = X \cup Y \cup \{z\}$, and consider the profile $(A_1, \dots, A_{2k})$, where $A_i = \{x_i, y_i\}$ for $i = 1, \dots, k$, $A_i = \{z\}$ for $i = k + 1, \dots, 2k$. 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. □

Interestingly, we can show that MAV satisfies JR if we assume that each agent 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.*

While Theorem 5 provides an example of a setting where a well-known voting rule 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 $\mathcal{A} = (A_1, \dots, A_n)$ and a $k > 0$ and let $s = \lceil \frac{n}{k} \rceil$. Let W be the output of PAV on $(\mathcal{A}, k)$. Suppose for the sake of contradiction that there exists a set $N^* \subset N$, $|N^*| \geq s$, 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 $w_j \leq \frac{1}{j}$ for all $j \geq 1$. In the full version of this paper, we show that this condition on $\mathbf{w}$ is also necessary for $\mathbf{w}$-PAV to satisfy JR.

Next, we consider RAV. 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 RAV satisfies JR as well. However, this turns out not to be the case, at least if k is sufficiently large.

Theorem 7 *RAV 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.

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

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

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 RAV selects c_1 and c_4 first. This reduces the RAV 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 RAV score is 147, is selected next. Now, the RAV 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 RAV 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

$120 > \frac{n}{k}$. RAV will proceed to select $c_1, \dots, c_{10}$, followeed by $k - 10$ additional candidates, and c_{11} or one of the new candidates will remain unselected. $\square$

While RAV itself fails JR, one could hope that this can be fixed by modifying the weights, i.e., that $\mathbf{w}$-RAV satisfies JR for a suitable weight vector $\mathbf{w}$. However, Theorem 7 extends to $\mathbf{w}$-RAV for *every* weight vector $\mathbf{w}$ with $w_2 > 0$.

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

Theorem 8 partially subsumes Theorem 7: it implies that RAV fails JR, but the proof only shows that this is the case for $k \geq 18 \cdot 19 = 342$, while Theorem 7 states that RAV 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. We remark that it remains an open problem whether RAV satisfies JR for $k = 3, \dots, 9$.

As we require $w_1 \geq w_2 \geq \dots$, the only weight vector not captured by Theorem 8 is $(1, 0, \dots, 0)$. In fact, $(1, 0, \dots, 0)$-RAV satisfies JR: indeed, this rule is exactly the greedy rule GAV! We can extend this result somewhat, by allowing the entries of the weight vector to depend on the number of voters n: the argument used to show that GAV satisfies JR extends to $\mathbf{w}$-RAV where the weight vector $\mathbf{w}$ satisfies $w_2 \leq \frac{1}{n}$. In particular, the rule $(1, \frac{1}{n}, \frac{1}{n^2}, \dots,)$-$RAV$ is somewhat more appealing than GAV: for instance, if $\bigcap_{i \in N} A_i = \{c\}$ and $k > 1$, GAV will pick c, and then behave arbitrarily, whereas $(1, \frac{1}{n}, \frac{1}{n^2}, \dots,)$-$RAV$ will also pick c, but then it will continue to look for candidates approved by as many voters as possible.

5 Extended Justified Representation

We have identified two families of voting rules that satisfy JR for arbitrary ballot profiles: $\mathbf{w}$-PAV with $w_j \leq \frac{1}{j}$ (this includes the PAV rule) and $\mathbf{w}$-RAV with $w_2 \leq \frac{1}{n}$ (this includes the GAV rule). The obvious advantage of the greedy rule is that its output can be computed efficiently, whereas computing the output of PAV in NP-hard. However, arguably, GAV puts too much 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. For instance, if $k = 3$, there are 98 voters who approve a and b, while c and d are each approved by a single voter, the greedy rule 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).

This issue is not addressed by the JR axiom, as it 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 $(A_1, \dots, A_n)$ over a candidate set C, a target committee size k, $k \leq |C|$, and a positive integer ℓ, $\ell \leq k$, we say that a set of candidates W, $|W| = k$,* provides ℓ-justified representation for $(\mathcal{A}, k)$ *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 $(\mathcal{A}, k)$ *if it provides ℓ-JR for $(\mathcal{A}, k)$ for all ℓ, $1 \leq \ell \leq k$. We say that an approval-based voting rule* satisfies ℓ-justified representation (ℓ-JR) *if for every profile $\mathcal{A} = (A_1, \dots, A_n)$ and every target committee size k it outputs a committee that provides ℓ-JR for $(\mathcal{A}, k)$. Finally, we say that a rule* satisfies extended justified representation (EJR) *if it satisfies ℓ-JR for all ℓ, $1 \leq \ell \leq k$.*

Observe that 1-JR is simply JR. However, EJR is not implied by JR: this is illustrated by the 4-candidate 100-voter example earlier in this section. Further, although EJR is stronger than JR, it still does not imply unanimity.

5.1 EJR under Approval-based Rules

It is natural to ask which of the voting rules that satisfy JR also satisfy EJR. Our 4-candidate 100-voter example immediately shows that for GAV the answer is negative. Consequently, no $\mathbf{w}$-RAV rule such that the entries of $\mathbf{w}$ do not depend on n satisfies EJR: if $w_2 = 0$, this rule is GAV, and if $w_2 > 0$, this follows from Theorem 8. For MAV, it can be shown that EJR is violated even if each voter approves exactly k candidates (recall that MAV satisfies JR under this assumption). It remains to consider PAV.

Theorem 9 *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. □

Interestingly, Theorem 9 does not extend to weight vectors other than $(1, \frac{1}{2}, \frac{1}{3}, \dots)$: our next theorem shows that PAV is the unique $\mathbf{w}$-PAV rule that satisfies EJR.

Theorem 10 *For every weight vector* $\mathbf{w}$ *with* $\mathbf{w} \neq (1, \frac{1}{2}, \frac{1}{3}, \dots)$*, the rule* $\mathbf{w}$-PAV *does not satisfy* EJR*.*

5.2 Computational Issues

In Section 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 for a fixed value of ℓ we can efficiently compute a committee that provides ℓ-JR, see the full version of this paper. For the problem of checking whether a given committee provides EJR for a given input, we can establish a formal hardness result.

Theorem 11 *Given a ballot profile* $\mathcal{A}$*, a target committee size* k*, and a committee* W*,* $|W| = k$*, it is* coNP*-complete to check whether* W *provides* EJR *for* $(\mathcal{A}, k)$*.*

Proof Sketch: It is easy to see that this problem is in coNP. To prove coNP-completeness, we reduce the classic BALANCED BICLIQUE problem ([GT24] in Garey and Johnson 1979) 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 $s \geq 3$, $\ell \geq 3$. 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 $\ell \geq 3$). 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 $\mathcal{A}$. 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 $\frac{n}{k} = s$.

Then, it can be proven that we have a "yes"-instance of BALANCED BICLIQUE iff W does not satisfy EJR. □

6 Discussion

We have formulated a desirable property of approval-based committee selection rules, which we called justified representation (JR). JR seems to have some merit over previous approaches towards fair representation. In particular, it seems more attractive than the related notion of *threshold representation* mentioned by Kilgour (2010). This notion requires that the winning set should represent all voters, if at all possible. It can be argued that threshold representation is overly egalitarian, as it ignores the relative numbers of agents supporting different candidates. Also, it is immediate that finding a committee that provides threshold representaiton is NP-hard (Fishburn and Pekec 2004). Another similar notion is that of *representativeness* (Duddy 2014); however, it applies to probabilistic voting rules, whereas JR can be stated for deterministic rules.

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 (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). 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. Also, it would be interesting to see if EJR, in combination with other natural axioms, can be used to axiomatize PAV.

Our analysis can be extended to approval-based variants of rules that provide fully proportional representation, such as Chamberlin–Courant's rule (Chamberlin and Courant 1983) and Monroe's rule (Monroe 1995). Specifically, under a natural adaptation of these rules to approval ballots, where the scoring function associated with each voter is identical to her ballot, Chamberlin–Courant's rule is simply $(1, 0, \dots)$-PAV, and hence satisfes JR, but not EJR; also, Monroe's rule can be shown to satisfy JR, but not EJR. We omit the definitions of these rules and the formal statements and proofs of the respective results, as the focus of this paper is the analysis of classic approval-based rules.

Justified representation can also be used to formulate new approval-based rules. We mention two rules that seem particularly attractive: The *utilitarian JR rule* returns a committee that, among all committees that satisfy JR, has the highest AV score. The *egalitarian JR rule* returns a committee that, among all committees that satisfy JR, maximizes the number of representatives of the agent who has the least number of representatives in the winning committee. The computational complexity of winner determination for these rules is an interesting problem. Finally, analyzing the compatibility of JR with other important properties, such as, e.g., strategyproofness for dichotomous preferences, is another avenue of future research.

Acknowledgments NICTA is funded by the Australian Government through the Department of Communications and the Australian Research Council through the ICT Centre of Excellence Program. Brill, Conitzer, and Freeman were supported by NSF and ARO under grants CCF-1101659, IIS-0953756, CCF-1337215, W911NF-12-1-0550, and W911NF-11-1-0332, and by a Feodor Lynen research fellowship of the Alexander von Humboldt Foundation. 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

Aziz, H.; Brill, M.; Conitzer, V.; Elkind, E.; Freeman, R.; and Walsh, T. 2014. Justified representation in approval-based committee voting. Technical Report 1407.8269, arXiv.org.

Aziz, H.; Gaspers, S.; Gudmundsson, J.; Mackenzie, S.; Mattei, N.; and Walsh, T. 2014. Computational aspects of multi-winner approval voting. In *MPREF-14*.

Betzler, N.; Slinko, A.; and Uhlmann, J. 2013. On the computation of fully proportional representation. *Journal of Artificial Intelligence Research* 47:475–519.

T. Lu and C. Boutilier. 2011. Budgeted social choice: From consensus to personalized decision making. In *IJCAI-11*, pages 280–286.

Brams, S. J., and Fishburn, P. C. 2007. *Approval Voting*. Springer-Verlag, 2nd edition.

Brams, S. J., and Kilgour, D. M. 2014. Satisfaction approval voting. In *Voting Power and Procedures*, Studies in Choice and Welfare. Springer. 323–346.

Brams, S. J., and Pothoff, R. F. 1990. Constrained approval voting: A voting system to elect a governing board. *Interfaces* 20(2):67–80.

Brams, S. J.; Kilgour, D. M.; and Sanver, M. R. 2006. How to elect a representative committee using approval balloting. In Simeone, B., and Pukelsheim, F., eds., *Mathematics and democracy: Recent advances in voting systems and collective choice*. Springer. 83–96.

Brams, S. J.; Kilgour, D. M.; and Sanver, M. R. 2007. A minimax procedure for electing committees. *Public Choice* 132:401–420.

Brams, S. J. 2010. Preface. In Laslier, J.-F., and Sanver, M. R., eds., *Handbook on Approval Voting*, Studies in Choice and Welfare. Springer. vii–ix.

Caragiannis, I.; Kalaitzis, D.; and Markakis, E. 2010. Approximation algorithms and mechanism design for minimax approval voting. In *AAAI-10*, 737–742.

D. Cornaz, L. Galand, and O. Spanjaard. 2012. Bounded single-peaked width and proportional representation. In *ECAI-12*, pages 270–275.

Chamberlin, J. R., and Courant, P. N. 1983. Representative deliberations and representative decisions: Proportional representation and the Borda rule. *The American Political Science Review* 77(3):718–733.

Conitzer, V. 2010. Making decisions based on the preferences of multiple agents. *Communications of the ACM* 53(3):84–94.

Davis, M.; Orrison, M. E.; and Su, F. E. 2014. Voting for committees in agreeable societies. Technical Report arXiv:1402.0861, arXiv.org.

Duddy, C. 2014. Electing a representative committee by approval ballot: An impossibility result. *Economics Letters* 124(1):14–16.

Droop, H. R. 1881. On methods of electing representatives. *Journal of the Statistical Society of London* 44(2):141–196.

Dummett, M. 1984. *Voting Procedures*. Oxford University Press.

Elkind, E.; Faliszewski, P.; Skowron, P.; and Slinko, A. 2014. Properties of multiwinner voting rules. In *AAMAS-14*, 53–60.

Elkind, E.; Lang, J.; and Saffidine, A. 2014. Condorcet winning sets. In *Social Choice and Welfare*, *Online First*.

Endriss, U. 2013. Sincerity and manipulation under approval voting. Theory and Decision 74 (4), 335–355.

Fishburn, P. C., and Pekec, S. 2004. Approval voting for committees: Threshold approaches.

Garey, M., and Johnson, D. 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company.

Kilgour, D. M. 2010. Approval balloting for multi-winner elections. In *Handbook on Approval Voting*. Springer. Chapter 6.

LeGrand, R.; Markakis, E.; and Mehta, A. 2007. Some results on approximating the minimax solution in approval voting. In *AAMAS-07*, 1193–1195.

Kilgour, D. M. and Marshall, E. 2012. Approval balloting for fixed-size committees. In *Electoral Systems*, Studies in Choice and Welfare. Springer. Chapter 12, 305–326.

Meir, R.; Procaccia, A. D.; and Rosenschein, J. S. 2008. A broader picture of the complexity of strategic behavior in multi-winner elections. In *AAMAS-08*, 991–998.

Monroe, B. L. 1995. Fully proportional representation. *The American Political Science Review* 89(4):925–940.

Procaccia, A. D.; Rosenschein, J. S.; and Zohar, A. 2008. On the complexity of achieving proportional representation. *Social Choice and Welfare* 30:353–362.

Skowron, P. K.; Faliszewski, P.; and Lang, J. 2014. Finding a collective set of items: From proportional multirepresentation to group recommendation. In *COMSOC-14*.

Skowron, P. K.; Faliszewski, P.; and Slinko, A. 2013. Fully proportional representation as resource allocation: Approximability results. In *IJCAI-13*, pages 353–359.

P. Skowron, L. Yu, P. Faliszewski, and E. Elkind. 2013. The complexity of fully proportional representation for single-crossing electorates. In *SAGT-13*, pages 1–12.

Straszak, A.; Libura, M.; Sikorski, J.; and Wagner, D. 1993. Computer-assisted constrained approval voting. *Group Decision and Negotiation* 2(4):375–385.

Tideman, N., and 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.