

The Proportional Veto Principle for Approval Ballots

Daniel Halpern¹, Ariel D. Procaccia¹, Warut Suksompong²

¹Harvard University

²National University of Singapore

dhalpern@g.harvard.edu, arielpro@seas.harvard.edu, warut@comp.nus.edu.sg

Abstract

The proportional veto principle, which captures the idea that a candidate vetoed by a large group of voters should not be chosen, has been studied for ranked ballots in single-winner voting. We introduce a version of this principle for approval ballots, which we call flexible-voter representation (FVR). We show that while the approval voting rule and other natural scoring rules provide the optimal FVR guarantee only for some flexibility threshold, there exists a scoring rule that is FVR-optimal for all thresholds simultaneously. We also extend our results to multi-winner voting.

1 Introduction

Voting is one of the most prominent approaches to collective decision-making and has been studied extensively in several disciplines, not least in social choice theory [Brams and Fishburn, 2002; Zwicker, 2016]. Some applications of voting require selecting a single winner, e.g., the president of a country or the venue for an event. Other applications involve choosing multiple winners, such as members of a council, places to visit on a trip, or recipients of a research grant.

In several voting scenarios, a natural consideration is that groups of voters should have the right to veto candidates that they dislike, be it a politician whose views they disagree with or a venue they find unsuitable. This intuition is formalized by the concept of the *proportional veto core*, which was proposed by Moulin [1981] and further studied by a number of authors [Moulin, 1982; Ianovski and Kondratev, 2021; Kizilkaya and Kempe, 2023; Peters, 2023; Kondratev and Ianovski, 2024]. The proportional veto core is defined for single-winner voting when each voter expresses her preference over candidates in the form of a strict ranking. According to this principle, a candidate a can be vetoed by a group of voters T if there exists a sufficiently large set of candidates B such that every voter in T prefers all candidates in B to a —the required size of B depends on the size of T . A significant result by Moulin [1981] states that the proportional veto core is always non-empty, i.e., there always exists a candidate that is not vetoed by any group of voters.

While ranked ballots have been examined since the dawn of social choice theory, an alternative method for eliciting

voter preferences, which has also received substantial attention in the literature, is via *approval ballots* [Brams and Fishburn, 2007; Laslier and Sanver, 2010]. In an approval ballot, a voter may either approve or disapprove each candidate. In addition to their simplicity, approval ballots are quite expressive, as voters are allowed to approve as many or as few candidates as they wish. As Brandl and Peters [2022] discussed, approval ballots arise naturally when preferences are dichotomous—for example, when they correspond to whether workers are available during various time slots, or whether hiring committee members consider candidates to be capable of performing a clearly defined task. Approval ballots have been used in political contexts—including in the St. Louis, Missouri mayoral and Fargo, North Dakota municipal elections—as well as by professional groups such as the American Mathematical Society.

A tempting approach for applying the proportional veto principle to approval ballots is to convert the approval ballots into ranked ballots by breaking ties arbitrarily, and apply the result of Moulin [1981] to obtain a candidate in the proportional veto core of the ranked ballots. However, this approach may lead to patently undesirable outcomes. For instance, consider four candidates a, b, c, d and three voters who approve $\{b, c\}$, $\{b, d\}$, $\{c, d\}$, respectively. One can convert these approval ballots into the ranked ballots where voter 1 ranks $b \succ c \succ a \succ d$, voter 2 ranks $b \succ d \succ a \succ c$, and voter 3 ranks $c \succ d \succ a \succ b$. For these approval ballots, voter 1 vetoes d for being last, and similarly voters 2 and 3 veto c and b , respectively, leaving a as the only candidate in the proportional veto core. However, the choice of a is clearly suboptimal with respect to the original approval ballots, as it is approved by none of the voters.

To formalize the proportional veto principle for approval ballots, we introduce a concept of voter *flexibility*, which corresponds to the fraction of candidates that a voter approves. Intuitively, in single-winner voting, if a group of voters is sufficiently large and each voter in the group is sufficiently flexible, we would like to ensure that at least one voter in the group approves the chosen candidate. Put differently, a sufficiently large and sufficiently flexible group has the power to veto a candidate that the entire group disapproves. While considering flexibility is natural in any setting where approval voting is employed, it may be especially useful in applications such as participatory budgeting, where voters sometimes partici-

pate to support a single project that aligns with a particular cause. Indeed, rules that reward flexibility could encourage voters to approve more projects, lead to chosen projects with broader consensus, and provide greater legitimacy to the process.¹ As another example, consider a group of countries deciding the location(s) for hosting an international event. In this scenario, flexibility could help motivate the representatives of each country to vote beyond merely their own country or close allies. We refer to our principle as *flexible-voter representation* (FVR).

We highlight that FVR is more robust than proportional veto for ranked ballots. In particular, for a candidate a to be vetoed by a group T , FVR does not require every voter in T to approve all candidates in some set B and disapprove a —instead, it only requires all voters in T to approve sufficiently many (but not necessarily the same) candidates and disapprove a . We discuss this distinction further in Section 1.2, but note here that FVR is particularly suited to approval ballots, where voter flexibility can be defined naturally.

1.1 Overview of Results

In Section 2, we focus on single-winner voting. For $s \in (0, 1)$, we say that a voter is s -flexible if she approves at least an s -fraction of the candidates. For each voting rule R and each s , we denote by $\text{FVR}(R, s)$ the smallest value of r such that whenever a group of voters constitutes more than an r -fraction of the voters and each voter in the group is s -flexible, at least one voter in the group approves the candidate chosen by R . The lower $\text{FVR}(R, s)$ is, the stronger the guarantee that the rule R provides for the threshold s .

We first derive a general lower bound: for any rule R and any s , it holds that $\text{FVR}(R, s) \geq 1 - s$. Moreover, this bound is tight in the sense that for each s , the rule $R_{s\text{-threshold}}$, which chooses a candidate with the most approvals among s -flexible voters, achieves $\text{FVR}(R_{s\text{-threshold}}, s) = 1 - s$. However, as each rule $R_{s\text{-threshold}}$ is tailored to a specific value of s , a rule with strong guarantees for all values of s simultaneously would be more desirable. We show that the classic approval voting rule, which chooses a candidate with the most approvals overall, yields $\text{FVR}(R, s) = \frac{1}{1+s}$ for all s ; this essentially matches the lower bound for s close to 0, but becomes further away as s increases. To achieve better guarantees for larger s , we consider placing more weight on voters with higher flexibility. Specifically, we let the contribution of a voter to the score of each of her approved candidates be the voter’s flexibility raised to the p -th power, where $p > 0$ is a given parameter. For each p , we determine the exact FVR of the p -power scoring rule—we find that it matches the lower bound at $s = \frac{p}{1+p}$, but not at other values of s .

From the aforementioned results, one may be tempted to believe that there is an inevitable trade-off between optimizing for different thresholds s . Strikingly, we show that this is not the case: there is in fact a rule R_{OPT} that is *FVR-optimal*, meaning that $\text{FVR}(R, s) = 1 - s$ for all s simultaneously. One may interpret this result as demonstrating that

an approval-ballot analog of the proportional veto core is always non-empty. Like p -power scoring rules, R_{OPT} can be described as a scoring rule, with the scores chosen carefully. We also establish a general theorem that can be used to determine FVR guarantees of arbitrary scoring rules—this theorem allows us to characterize R_{OPT} as the unique scoring rule (up to scaling) that satisfies FVR-optimality.

In Section 3, we turn our attention to the more general setting of multi-winner voting, where the goal is to choose a set of k candidates (called a *committee*) from the m given candidates; single-winner voting corresponds to $k = 1$. For $1 \leq t \leq k$, we say that a voter t -approves a committee if she approves at least t candidates in the committee. Unlike in the single-winner setting, it will be more convenient to state our bounds including their dependency on m (rather than in the worst case over all m). Specifically, we let $\text{FVR}(R, s, t, m)$ be the smallest value of r such that whenever a group of voters constitutes more than an r -fraction of the voters, each voter is s -flexible, and there are m candidates, then at least one voter in the group t -approves the committee chosen by R .

Similarly to single-winner voting, we provide a lower bound on $\text{FVR}(R, s, t, m)$, which can be written in terms of a hypergeometric random variable; for any s, t, m, k , we present a rule for which this bound is tight. Moreover, if k and t are fixed, we show that there exists a rule that yields the optimal guarantee for all m and s . This rule operates by iteratively choosing a candidate using a scoring rule; however, the weight assigned to each voter in this procedure depends non-trivially on the voter’s preference and the candidates added thus far. On the other hand, we prove that if we fix s (and k, m), it may not be possible to achieve FVR-optimality for different values of t at the same time. We also demonstrate the incompatibility between FVR-optimality and other notions of representation in the multi-winner literature.

1.2 Additional Discussion

As mentioned earlier, in single-winner voting, FVR can be viewed as an approval-ballot analog of the proportional veto core for ranked ballots [Moulin, 1981]. However, in order for a group of voters to veto a candidate a , FVR does not require the voters in the group to commonly approve a set of candidates and disapprove a —instead, it is enough that these voters approve sufficiently many candidates and disapprove a . To highlight the difference that this distinction makes, note that strengthening the proportional veto core along the lines of our FVR definition may lead to an empty core. Indeed, if there are two voters with rankings $a \succ b \succ c$ and $c \succ b \succ a$ over three candidates, then this more demanding notion would rule out c (as the first voter ranks it last), a (as the second voter ranks it last), and b (as both voters rank it second-last), leaving no viable candidate.² In the approval setting, FVR takes advantage of the dichotomous nature of the preferences to achieve a stronger guarantee.

Another related concept is *justified representation* (JR) [Aziz et al., 2017], which has received significant attention

¹We also remark that proportionality is a central guarantee behind the *method of equal shares*, which has been used for participatory budgeting in Poland and Switzerland [Boehmer et al., 2024].

²We formalize this in the full version of our paper [Halpern et al., 2025]. Note that the proportional veto core does not rule out b , since the two voters do not agree on a candidate that is preferred to b .

recently in the context of multi-winner (approval) voting. JR captures the idea that if a group of voters is sufficiently large and all voters in the group approve a common candidate, then at least one candidate approved by some voter in the group should be selected. While FVR bears some similarity to JR, a major difference is that FVR does not require voters in the group to approve any candidate in common, but only that these voters be sufficiently flexible. Note that in single-winner voting, JR does not provide any meaningful guarantee, and neither do its various strengthenings [Aziz *et al.*, 2017; Sánchez-Fernández *et al.*, 2017; Peters *et al.*, 2021; Brill and Peters, 2023].

2 Single-Winner Voting

2.1 Preliminaries

For each positive integer t , let $[t] := \{1, 2, \dots, t\}$. There is a set N of n voters and a set M of m candidates. Each voter i has an approval set $A_i \subseteq M$. An instance consists of N , M , and $(A_i)_{i \in N}$. For a candidate $a \in M$, let $N_a := \{i \in N \mid a \in A_i\}$ be the set of voters who approve it. For a voter $i \in N$, let $f_i := |A_i|/m$ be the fraction of candidates that i approves. In single-winner voting, a rule chooses a single candidate from each instance.

For $s \in (0, 1)$, we say that a voter i is s -flexible if $f_i \geq s$. Given a rule R , we are interested in, for each $s \in (0, 1)$, the smallest $r \in [0, 1]$ (as a function of s) such that whenever a group of strictly more than rn voters (i.e., more than an r -fraction of the voters) are all s -flexible, at least one voter in the group approves the candidate chosen by the rule.³ We refer to this condition as the *FVR condition*, and denote this value of r by $\text{FVR}(R, s)$. Note that $\text{FVR}(R, s)$ is always well-defined. Indeed, $r = 1$ satisfies the FVR condition vacuously; if some r satisfies the condition then every $r' \geq r$ does as well; and if for some r it holds that every $r' > r$ satisfies the condition, then r also satisfies the condition (because the condition is phrased as “strictly more than rn voters”). The smaller the value $\text{FVR}(R, s)$, the stronger the FVR guarantee provided by the rule R for the threshold s .

Example 2.1. Consider a rule R that always returns the first candidate, regardless of the voters’ approvals. We claim that $\text{FVR}(R, s) = 1$ for all $s \in (0, 1)$, i.e., this rule offers the worst possible FVR guarantees. Indeed, for any $s \in (0, 1)$, consider an instance where each of the n voters approves all m candidates except the first candidate, where $m > \frac{1}{1-s}$. All voters are s -flexible but none of them approves the candidate chosen by R , so the FVR condition holds only for $r = 1$.

2.2 Lower Bound

We first derive a lower bound on the FVR guarantee achievable by any rule.

Theorem 2.2. *For any rule R and any $s \in (0, 1)$, we have $\text{FVR}(R, s) \geq 1 - s$. Moreover, for each $s \in (0, 1)$, there exists a rule R such that $\text{FVR}(R, s) = 1 - s$.*

³If $s = 1$, the answer is clearly $r = 0$ regardless of the rule, so we do not consider this trivial case.

Proof. For the first statement, consider any rule R and any $s \in (0, 1)$, and take any $r < 1 - s$. It suffices to show that there exists an instance where a group of at least rn voters are all s -flexible but do not approve the candidate chosen by R . Consider an instance with sufficiently large n and m (to be made more precise later). Assume that each of the n voters approves exactly $\lceil sm \rceil$ candidates (so all voters are s -flexible), and the voters distribute their approvals as equally as possible.⁴ Hence, each candidate receives at most $\left\lceil \frac{n \cdot \lceil sm \rceil}{m} \right\rceil$ approvals. Observe that

$$\left\lceil \frac{n \cdot \lceil sm \rceil}{m} \right\rceil < \frac{n(sm + 1)}{m} + 1 = n \cdot \frac{sm + 1}{m} + 1.$$

As m grows, $\frac{sm+1}{m}$ converges to s , so $n \cdot \frac{sm+1}{m} + 1$ converges to $sn + 1$. Since $s < 1 - r$, we have $sn + 1 < (1 - r)n$ for sufficiently large n . This means that for sufficiently large n and m , each candidate is approved by at most $(1 - r)n$ voters. Consequently, no matter which candidate R chooses, at least rn voters disapprove it.

For the second statement, fix $s \in (0, 1)$. Consider the rule $R_{s\text{-threshold}}$ that chooses a candidate a with the highest number of approvals among s -flexible voters, breaking ties arbitrarily; in particular, the rule ignores voters who are not s -flexible. It suffices to show that whenever a group of more than $(1 - s)n$ voters are all s -flexible, at least one voter in the group approves a . Since each s -flexible voter approves at least an s -fraction of the candidates, the candidates are approved by at least an s -fraction of the s -flexible voters on average. By definition of the rule $R_{s\text{-threshold}}$, the candidate a is approved by at least an s -fraction of the s -flexible voters, which means that at most a $(1 - s)$ -fraction of the s -flexible voters disapprove a . Since the number of s -flexible voters is at most n , this implies that at most $(1 - s)n$ voters who are s -flexible disapprove a . It follows that in any group of more than $(1 - s)n$ voters who are s -flexible, at least one voter in the group approves a . \square

While the rule $R_{s\text{-threshold}}$ in the proof of Theorem 2.2 achieves the optimal FVR guarantee for the corresponding flexibility threshold s , it is tailored to only one threshold s , which leads to unwanted effects—for instance, the rule completely ignores voters who are almost s -flexible. As such, it would be more desirable to have a rule that yields good guarantees for several thresholds s simultaneously.

2.3 Approval Voting Rule

We next analyze the classic *approval voting rule*, which selects a candidate with the highest number of approvals, breaking ties arbitrarily. For this rule, we determine the tight FVR guarantee for every s .

Theorem 2.3. *Let R_{approval} be the approval voting rule. For each $s \in (0, 1)$, we have $\text{FVR}(R_{\text{approval}}, s) = \frac{1}{1+s}$.*

Proof. Fix $s \in (0, 1)$. We first show the upper bound $\text{FVR}(R_{\text{approval}}, s) \leq \frac{1}{1+s}$. Take an arbitrary instance, and

⁴For example, we can let the voters distribute their approvals one voter after another. In each voter’s turn, the voter approves $\lceil sm \rceil$ candidates that have received the fewest approvals so far, breaking ties arbitrarily.

consider a group B of more than $\frac{1}{1+s} \cdot n$ voters, all of whom are s -flexible. The total number of approvals made by the voters in B is greater than $\frac{1}{1+s} \cdot n \cdot sm$. Since there are m candidates, there exists a candidate a approved by more than $\frac{1}{1+s} \cdot n \cdot s = \frac{s}{1+s} \cdot n$ voters in B . On the other hand, a candidate that is not approved by any voter in B has an approval score of less than $(1 - \frac{1}{1+s})n = \frac{s}{1+s} \cdot n$. In particular, such a candidate is approved by fewer voters than a . Hence, the candidate chosen by the approval voting rule must be approved by some voter in B .

Next, we show that $\text{FVR}(R_{\text{approval}}, s) \geq \frac{1}{1+s}$. Take any $r < \frac{1}{1+s}$. It suffices to show that there exists an instance where a group of at least rn voters are all s -flexible but do not approve the candidate chosen by the approval voting rule. Consider an instance with sufficiently large n and m (to be made more precise later). Let C be a group of $\lceil rn \rceil$ voters, and assume that all voters in $N \setminus C$ approve a common candidate a (and no other candidate). Moreover, assume that each voter in C approves exactly $\lceil sm \rceil$ candidates in $M \setminus \{a\}$, and these voters distribute their approvals among these candidates as equally as possible. Note that a receives $n - \lceil rn \rceil = \lfloor (1-r)n \rfloor$ approvals, while every other candidate receives at most $\left\lceil \frac{\lceil rn \rceil \cdot \lceil sm \rceil}{m-1} \right\rceil$ approvals. Since $r < \frac{1}{1+s}$, we have $1-r > \frac{s}{1+s}$, and so $\lfloor (1-r)n \rfloor > \frac{s}{1+s} \cdot n$ when n is sufficiently large. Hence, a receives more than $\frac{s}{1+s} \cdot n$ approvals for sufficiently large n . On the other hand, note that

$$\begin{aligned} \left\lceil \frac{\lceil rn \rceil \cdot \lceil sm \rceil}{m-1} \right\rceil &< \frac{(rn+1)(sm+1)}{m-1} + 1 \\ &= (rn+1) \cdot \frac{sm+1}{m-1} + 1. \end{aligned}$$

As m grows, $\frac{sm+1}{m-1}$ converges to s , so $(rn+1) \cdot \frac{sm+1}{m-1} + 1$ converges to $(rn+1)s + 1 = rns + s + 1$. Since $r < \frac{1}{1+s}$, we have $rns + s + 1 < \frac{1}{1+s} \cdot ns = \frac{s}{1+s} \cdot n$ for sufficiently large n . This means that for sufficiently large n and m , each candidate in $M \setminus \{a\}$ receives at most $\frac{s}{1+s} \cdot n$ approvals. Consequently, the approval voting rule chooses the candidate a , which is approved by none of the voters in C . \square

Comparing Theorems 2.2 and 2.3, we find that the FVR guarantee of the approval voting rule is essentially optimal for s close to 0, but gets further from optimal as s increases.

2.4 Power Scoring Rules

To achieve better guarantees for larger s , an enticing idea is to give more weight to voters with higher flexibility. However, unlike the rule $R_{s\text{-threshold}}$ in the proof of Theorem 2.2 where this weight is binary (according to whether a voter is s -flexible), we shall use a smoother weighting scheme.

To this end, for each real number $p \geq 0$, we define the p -power approval score of candidate a as $\text{SC}_p(a) := \sum_{i \in N_a} f_i^p$ (recall that N_a denotes the set of voters who approve a). We consider the p -power scoring rule, which chooses a candidate that maximizes the p -power approval score, breaking ties arbitrarily. Note that if $p = 0$, then $\sum_{i \in N_a} f_i^p = |N_a|$, so the rule reduces to the approval voting rule which we already analyzed in Theorem 2.3. As p increases, the rule places more




Figure 1: FVR guarantees of p -power scoring rules for flexibility thresholds $s \in (0, 1)$, compared to the optimal guarantees indicated by the thick line. For $p = 1, 2$, the guarantees match the optimal ones at $s = 1/2, 2/3$, respectively.

importance on voters with high flexibility. For each p , we derive tight FVR guarantees for the corresponding rule.

Theorem 2.4. For each $p > 0$, let $R_{p\text{-power}}$ be the p -power scoring rule. For each $s \in (0, 1)$, we have $\text{FVR}(R_{p\text{-power}}, s) = \frac{1}{1 + \frac{(s(1+p))^{1+p}}{p^p}}$.

The proof of Theorem 2.4, along with all other missing proofs, can be found in the full version of our paper [Halpern et al., 2025]. Comparing Theorem 2.4 with Theorem 2.2, we find that the FVR guarantee of the p -power scoring rule is optimal only at $s = p/(1+p)$. Indeed, one can verify that

$$\frac{1}{1 + \frac{(s(1+p))^{1+p}}{p^p}} = 1 - s \iff s^p(1-s) = \frac{p^p}{(1+p)^{1+p}}.$$

Within the range $x \in [0, 1]$, basic calculus shows that the function $x^p(1-x)$ is maximized at $x = \frac{p}{1+p}$, and the resulting maximum is $\frac{p^p}{(1+p)^{1+p}}$. Hence, the only s for which the FVR guarantee of the p -power scoring rule is optimal is $s = \frac{p}{1+p}$, and the corresponding guarantee is $\frac{1}{1+p}$.

The FVR guarantees of p -power scoring rules for $p = 0$ (i.e., approval voting rule), 1, and 2, as well as the optimal guarantees, are illustrated in Figure 1.

2.5 Optimal Scoring Rule

Our discussion thus far raises an obvious question: is there an inherent trade-off between optimizing for different values of s , or does there exist a rule that achieves the optimal FVR guarantee for all s simultaneously (that is, a rule whose guarantees exactly match the thick line in Figure 1)? Perhaps surprisingly, we show that such a rule in fact exists, and moreover, comes from a simple class called *scoring rules* [Fishburn, 1979].⁵

⁵In the original definition, these are called *simple scoring rules*, but for conciseness, we drop the word ‘simple’.

To define scoring rules, let a *weight function* be a function $w : \mathbb{Q} \cap (0, 1) \rightarrow \mathbb{R}_{\geq 0}$ mapping each flexibility to a weight. From this, we define the *w-score* of a candidate a to be $\text{sc}^w(a) = \sum_{i \in N_a} w(f_i)$. We assume without loss of generality that $w(1) = 0$; if a voter i has $f_i = 1$, then she approves all candidates, so modifying $w(1)$ changes the scores of all candidates by a constant. We call a weight function w *nontrivial* if $w(f) > 0$ for some $f \in (0, 1)$.⁶ We say that a rule is *parameterized by w* if it always picks a candidate maximizing sc^w , and we call it a *scoring rule* if it is parameterized by some nontrivial w . Note that all of the rules we have considered so far are scoring rules: $R_{s\text{-threshold}}$ is parameterized by $w(f) = \mathbb{I}[f \geq s]$, R_{approval} is parameterized by $w(f) = 1$, and $R_{p\text{-power}}$ is parameterized by $w(f) = f^p$.

For our optimal rule, unlike the power rules we have seen, we will choose the weight function $w(f) = \frac{1}{1-f}$. Specifically, let R_{OPT} be the rule that assigns a score of $\sum_{i \in N_a} \frac{1}{1-f_i}$ to each candidate a and selects a candidate with the highest score, breaking ties arbitrarily.⁷ We say that a rule R is *FVR-optimal* if $\text{FVR}(R, s) = 1 - s$ for all $s \in (0, 1)$.

Theorem 2.5. *The rule R_{OPT} is FVR-optimal.*

Proof. Observe that an equivalent formulation of R_{OPT} is that each voter i distributes $\frac{1}{1-f_i}$ points to each candidate that the voter *disapproves*, and R_{OPT} chooses a candidate with the *lowest* total points. In this formulation, each voter i distributes a total of $\frac{1}{1-f_i} \cdot (1 - f_i)m = m$ points,⁸ so at most nm points are distributed overall. This means that the candidate a chosen by R_{OPT} receives at most n points.

Next, fix any $s \in (0, 1)$, and consider a group of s -flexible voters who all disapprove a . Since each voter in the group gives at least $\frac{1}{1-s}$ points to a , the size of the group is at most $(1 - s)n$. It follows that whenever a group of more than $(1 - s)n$ voters are all s -flexible, at least one voter in the group approves a . \square

2.6 General Weight Function Analysis

Theorems 2.3 to 2.5 provide analyses for specific scoring rules. What can we say about scoring rules more broadly? In the following theorem, we analyze the FVR guarantees of arbitrary scoring rules. Given a weight function w , let R^w denote a scoring rule parameterized by w .

Theorem 2.6. *Let w be a nontrivial weight function. Then,*

$$\text{FVR}(R^w, s) = \frac{\rho}{\rho + \varphi_s},$$

where $\rho := \sup_f (1 - f) \cdot w(f)$ and $\varphi_s := \inf_{f \geq s} f \cdot w(f)$; if $\rho = \infty$, the ratio is defined to be 1. Further, if w is nondecreasing, then $\varphi_s = s \cdot w(s)$ and the bound simplifies to

$$\text{FVR}(R^w, s) = \frac{\rho}{\rho + s \cdot w(s)}.$$

⁶If $w(f) = 0$ for all f , then the w -scores of all candidates will be equal.

⁷If a voter approves all candidates, we simply ignore that voter.

⁸Again, we ignore voters who approve all candidates. Note that voters who approve none of the candidates also distribute points.

Proof. Fix w , R^w , and $s \in (0, 1)$. We first bound $\text{FVR}(R^w, s)$ from above. If $\rho = \infty$ then the bound holds trivially, so assume that $\rho < \infty$. Fix an instance and let a^* be the candidate chosen by R^w . Let B be a set of s -flexible voters who disapprove a^* , and let $\bar{B} = N \setminus B$ be the set of remaining voters. Note that $\text{sc}^w(a^*) \leq \sum_{i \in \bar{B}} w(f_i)$, because no voter in B can contribute to the score of a^* . Consider now the average score of all candidates, which is equal to

$$\begin{aligned} \frac{1}{m} \sum_{a \in M} \text{sc}^w(a) &= \frac{1}{m} \sum_{a \in M} \sum_{i \in N_a} w(f_i) \\ &= \frac{1}{m} \sum_{i \in N} |A_i| \cdot w(f_i) = \sum_{i \in N} f_i \cdot w(f_i), \end{aligned}$$

where the second equality holds because each voter i contributes $w(f_i)$ points to each of exactly $|A_i|$ candidates. Since a^* has the largest w -score, it must have w -score at least the average. This implies that

$$\sum_{i \in \bar{B}} w(f_i) - \sum_{i \in N} f_i \cdot w(f_i) \geq \text{sc}^w(a^*) - \frac{1}{m} \sum_{a \in M} \text{sc}^w(a) \geq 0.$$

On the other hand,

$$\begin{aligned} \sum_{i \in \bar{B}} w(f_i) - \sum_{i \in N} f_i \cdot w(f_i) &= \sum_{i \in \bar{B}} w(f_i) - \left(\sum_{i \in B} f_i \cdot w(f_i) + \sum_{i \in \bar{B}} f_i \cdot w(f_i) \right) \\ &= \sum_{i \in \bar{B}} (1 - f_i) \cdot w(f_i) - \sum_{i \in B} f_i \cdot w(f_i) \\ &\leq (n - |B|)\rho - |B|\varphi_s, \end{aligned}$$

where the inequality holds by the definitions of ρ and φ_s , as $f_i \geq s$ for $i \in B$. Finally, observe that

$$\begin{aligned} (n - |B|)\rho - |B|\varphi_s \geq 0 &\iff n\rho \geq |B|(\rho + \varphi_s) \\ &\iff \frac{|B|}{n} \leq \frac{\rho}{\rho + \varphi_s}. \end{aligned}$$

Note that since w is nontrivial, we have $\rho > 0$, and so $\rho + \varphi_s > 0$ and the last transition is valid.

The remainder of the proof, which bounds $\text{FVR}(R^w, s)$ from below, is left to the full paper [Halpern *et al.*, 2025]. \square

Theorem 2.6 is a powerful tool that allows us to analyze arbitrary scoring rules. Indeed, Theorems 2.3 to 2.5 could be proven as consequences (although we keep them separate for exposition). In addition, we can derive other nontrivial consequences, including the following fact that R_{OPT} is essentially the unique scoring rule that ensures FVR-optimality.

Theorem 2.7. *Suppose that R^w is FVR-optimal for some nontrivial weight function w . Then, there exists a positive real number c such that $w(f) = \frac{c}{1-f}$ for all f .*

3 Multi-Winner Voting

3.1 Preliminaries

In this section, we turn to the more general setting of multi-winner voting. Given an instance, our rules R now choose a

subset of winning candidates of a given size k . We refer to a subset of k candidates as a k -committee (or just *committee* if k is clear from the context) and refer to a rule that selects a k -committee as a k -committee rule. Assume without loss of generality that $k < m$, since if $k = m$, there is a single k -committee which also makes all voters maximally happy.

To generalize the FVR condition to the multi-winner setting, we need to generalize the notion of a voter “approving the chosen candidate” to “approving the chosen committee.” Two natural generalizations immediately come to mind. Fix a committee W . A quite lenient notion we could consider is to say that a voter i approves W if she approves *at least one* of the committee members (i.e., $W \cap A_i \neq \emptyset$), while a much more stringent requirement is that i approves W only if she approves *all* of the committee members (i.e., $W \subseteq A_i$). We will generalize both of these notions at once by introducing an additional (positive integer) parameter $t \leq k$, and saying that a voter t -approves a committee W if she approves at least t of the committee members (i.e., $|W \cap A_i| \geq t$). If a voter does not t -approve a committee (i.e., $|W \cap A_i| < t$), we will say that she t -disapproves it. The earlier notions correspond to the special cases of $t = 1$ and $t = k$, respectively.

With this generalization in hand, we can now extend the notion of FVR to the multi-winner setting. Unlike in the single-winner setting, it will be more convenient to state our guarantees including their dependency on m (rather than in the worst case over all m), although we will primarily be interested in the behavior of these bounds as m grows large. This is because in some cases, unlike in the single-winner setting, as m grows, the worst case may occur for a fixed finite m rather than in the limit. With this in mind, for a k -committee rule R , we define $\text{FVR}(R, s, t, m)$ as the largest possible fraction of voters that are all s -flexible, yet t -disapprove the chosen committee when choosing from m candidates using R .

3.2 Lower Bound

We begin by deriving a lower bound analogous to Theorem 2.2 from the single-winner setting. Let $h(P, K, \ell; \cdot)$ and $H(P, K, \ell; \cdot)$ be the probability mass function and cumulative distribution function of a hypergeometric random variable with population size P , population success size K , and draw size ℓ , respectively.⁹ In other words, this is the distribution over the number of “good” elements sampled when taking a random subset of size ℓ out of P elements where K of the elements are good.

Theorem 3.1. *For any k -committee rule R , any $s \in (0, 1)$, and any $t \leq k$, we have*

$$\text{FVR}(R, s, t, m) \geq H(m, \lceil sm \rceil, k; t - 1).$$

Additionally, there exists a k -committee rule R (tailored to this (s, t) combination) for which this is tight.

The proof of this theorem (along with other results about multi-winner voting) relies on the following correspondence

⁹Recall that $h(P, K, \ell; t) = \frac{\binom{K}{t} \binom{P-K}{\ell-t}}{\binom{P}{\ell}}$ for $t \in \{0, 1, \dots, \ell\}$, and $H(P, K, \ell; t) = \sum_{t'=0}^t h(P, K, \ell; t')$ for the same range of t .

between the number of candidates approved and the number of k -committees which are t -approved.

Lemma 3.2. *If a voter approves exactly ℓ candidates, then she t -approves a $(1 - H(m, \ell, k; t - 1))$ -fraction of all k -committees.*

Proof. This follows from the definition of the hypergeometric distribution. Define the population to be the m candidates and a “success” to be the ℓ candidates approved by the voter. Then, if we pick a k -committee uniformly at random, the probability that the voter approves at most $t - 1$ candidates on the committee is $H(m, \ell, k; t - 1)$. Therefore, the voter t -approves the committee with probability $1 - H(m, \ell, k; t - 1)$. \square

While the bound in Theorem 3.1 may not be the most natural-looking, note that as m grows large, the hypergeometric distribution with parameters $m, \lceil sm \rceil, k$ approaches a binomial $\text{Bin}(k, s)$ distribution. Hence, the FVR bound approaches

$$\sum_{t'=0}^{t-1} \binom{k}{t'} s^{t'} (1-s)^{k-t'}.$$

Although still unwieldy, the binomial limit gives us a pathway to gain more intuition on this bound. By standard concentration inequalities, if $s \ll t/k$, then this bound is exponentially close to 1, whereas if $s \gg t/k$, then it is exponentially close to 0. Further, we can see that it simplifies in some special cases. Indeed, when $k = 1$ (so $t = 1$), the hypergeometric bound simplifies to $1 - \frac{\lceil sm \rceil}{m}$. This is always at most $1 - s$ and approaches $1 - s$ as m grows; hence, it coincides with the single-winner bound (Theorem 2.2). When we fix $t = 1$ but let k be arbitrary, the limiting binomial bound simplifies to $(1 - s)^k$. If we fix $t = k$, it simplifies to $1 - s^k$. For finite m , the hypergeometric distribution is slightly more concentrated than the binomial distribution—for $t \ll sk$, the binomial bound is an overestimate, while for $t \gg sk$, it is an underestimate. Specifically, for $t = 1$ the hypergeometric bound is always less than the binomial limit, and for $t = k$ it is always more. This aligns with the fact that when sampling without replacement, getting 0 success is less likely than with replacement. Similarly, getting k (out of k) successes is less likely without replacement than with replacement.

3.3 Simultaneously-Optimal Rules

We will say that a k -committee rule R is *FVR-optimal* for (s, t) if for all m , $\text{FVR}(R, s, t, m) = H(m, \lceil sm \rceil, k; t - 1)$. As with the single-winner setting, Theorem 3.1 implies that for any fixed (s, t) , we can find an FVR-optimal rule for this combination. However, what if we wish to attain optimality for multiple choices of (s, t) simultaneously? We first show that for a fixed t , we can achieve optimality over all s via a simple algorithm. Given k and t , we construct what we call a (k, t) -expanded instance as follows. The candidates in this expanded instance are all $\binom{m}{k}$ possible committees of size k . Each voter approves a committee W exactly when she t -approves W in the original instance. We have the following.

Theorem 3.3. Fix k and t . Let $R_{\text{expanded}}^{k,t}$ be the rule that runs R_{OPT} on the (k, t) -expanded instance and selects the winning committee. Then, for all $s \in (0, 1)$, $R_{\text{expanded}}^{k,t}$ is FVR-optimal for (s, t) .

This theorem shows that we can directly generalize our single-winner rule to be optimal (in a certain sense) in the multi-winner setting. However, a downside of $R_{\text{expanded}}^{k,t}$ is that, in its current implementation, it runs in time exponential in k , as we must construct an instance with $\binom{m}{k}$ candidates. Hence, a natural question is whether we can obtain the same guarantees with a polynomial-time algorithm. We show next that this is indeed possible using a greedy algorithm, which effectively applies the method of conditional expectations, presented as Algorithm 1.

Algorithm 1 (k, t) Sequential Algorithm

```

 $C \leftarrow \emptyset$ 
for  $j = 1, \dots, k$  do
  for  $i = 1, \dots, n$  such that  $A_i \setminus C \neq \emptyset$  do
     $w_i \leftarrow \frac{h(m-j-1, |A_i \setminus C|-1, k-j; t-1-|A_i \cap C|)}{H(m, |A_i|, k; t-1)}$ 
    Let  $a_j \in M \setminus C$  be a candidate maximizing  $\sum_{i \in N_a} w_i$ ,
    breaking ties arbitrarily.
     $C \leftarrow C \cup \{a_j\}$ 
return  $C$ 

```

At each step, the algorithm chooses a candidate that maximizes a certain score among the remaining candidates, and adds it to the committee. However, unlike in our previous scoring rules, the weight given to each voter i now depends not only on i 's flexibility, but also on the size of the current committee and the number of candidates in that committee approved by i . The algorithm yields the following guarantee.

Theorem 3.4. Fix k and t . Let $R_{\text{alg}}^{k,t}$ be the rule induced by Algorithm 1. Then, for all $s \in (0, 1)$, $R_{\text{alg}}^{k,t}$ is FVR-optimal for (s, t) .

To gain some intuition about the algorithm, note that for the (k, t) -expanded instance in Theorem 3.3 (and in the proof of Theorem 2.5), we need not choose a committee maximizing the w -score (for the optimal $w(f) = \frac{1}{1-f}$); instead, it is sufficient to choose one with *above-average* score. In other words, if we pick a committee W uniformly at random, $\mathbb{E}[\text{sc}^w(W)]$ would be sufficient for optimal guarantees. However, a priori, it is possible that only a single committee W^* has $\text{sc}(W^*) \geq \mathbb{E}[\text{sc}(W)]$ while all others are strictly below, where we drop the superscript w for convenience. To get around this and “round” the random committee into a deterministic one, we greedily construct $C = \{a_1, \dots, a_k\}$ such that at each step, $\mathbb{E}[\text{sc}(W) \mid \{a_1, \dots, a_j\} \subseteq W] \geq \mathbb{E}[\text{sc}(W)]$. The weighting rule we use essentially tells us the marginal gain to the score of permanently adding each possible candidate a . On average, adding each does not change the expectation, so adding one with the largest marginal gain can only improve the conditional expectation, thereby giving the desired bound.

Next, we show that unlike the positive results we found for simultaneous optimality over s in Theorems 3.3 and 3.4, the

same does not hold for t .

Theorem 3.5. For any $k \geq 2$ and $t \leq \lfloor k/2 \rfloor$, no rule is simultaneously FVR-optimal for $(1/2, t)$ and $(1/2, k)$.

3.4 Compatibility with Other Notions of Representation

Finally, we address the compatibility between FVR and other notions of representation. While several such notions have been proposed in the literature [Aziz et al., 2017; Sánchez-Fernández et al., 2017; Peters et al., 2021; Brill and Peters, 2023], one of the weakest is *justified representation* (JR). We show that FVR-optimality is incompatible even with JR; this implies a similar incompatibility with stronger representation notions (see the discussion in Section 1.2).

Recall that a committee W of size k is said to satisfy JR if there is no *blocking coalition* T of voters which fulfills the following properties:

1. *Large*: $|T| \geq n/k$;
2. *Cohesive*: $\bigcap_{i \in T} A_i \neq \emptyset$;
3. *Unrepresented*: $W \cap A_i = \emptyset$ for all $i \in T$.

We establish the following incompatibility.

Theorem 3.6. For each $k > 1$, there exists s such that no k -committee rule is FVR-optimal for $(s, 1)$ and satisfies JR at the same time.

4 Discussion

In this paper, we have introduced the notion of voter flexibility and a corresponding objective, *flexible-voter representation* (FVR), under approval voting. In the single-winner setting, we present a simple rule that is simultaneously optimal for all flexibility thresholds. In the multi-winner setting, while some impossibilities exist, we still find a polynomial-time rule which makes similar guarantees, once the definition of “approving a committee” (i.e., the parameter t) is fixed.

One may wonder what the practical implications of these objectives are. That is, should we care more about voters who are flexible? The answer certainly depends on the context. For example, consider an election in which Party A runs 20 candidates whereas Party B runs only 2. In this setting, Party A voters are not necessarily “more flexible” than Party B voters, and it may be preferable to simply take the approval winner. On the other hand, consider a setting where a community is deciding on one of several projects to fund. In this case, we may wish to reward voters who are more flexible. Indeed, the flexibility of such voters may indicate that they find their disapproved projects strongly unacceptable. Furthermore, encouraging flexibility can help achieve consensus, an inherently difficult task if few voters are flexible and no candidate has a reasonably-sized base.

More broadly, we believe that flexibility is an interesting foundational metric in its own right. Even when the chosen rules are not FVR-optimal, flexible-voter representation can be used to compare various candidates and weigh them according to their total support. Applied in a way suitable to the context, flexibility measures can help ensure a more desirable outcome in approval-based voting scenarios.

Acknowledgments

This work was partially supported by the National Science Foundation under grants IIS-2147187, IIS-2229881, and CCF-2007080, by the Office of Naval Research under grant N00014-20-1-2488, by the Singapore Ministry of Education under grant number MOE-T2EP20221-0001, and by an NUS Start-up Grant. We thank the anonymous reviewers for their valuable feedback.

References

- [Aziz *et al.*, 2017] Haris Aziz, Markus Brill, Vincent Conitzer, Edith Elkind, Rupert Freeman, and Toby Walsh. Justified representation in approval-based committee voting. *Social Choice and Welfare*, 42(2):461–485, 2017.
- [Boehmer *et al.*, 2024] Niclas Boehmer, Piotr Faliszewski, Łukasz Janeczko, Dominik Peters, Grzegorz Pierczyński, Šimon Schierreich, Piotr Skowron, and Stanisław Szufa. Evaluation of project performance in participatory budgeting. In *Proceedings of the 33rd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2678–2686, 2024.
- [Brams and Fishburn, 2002] Steven J. Brams and Peter C. Fishburn. Voting procedures. In Kenneth J. Arrow, Amartya K. Sen, and Kotaro Suzumura, editors, *Handbook of Social Choice and Welfare*, chapter 4, pages 173–236. North-Holland, 2002.
- [Brams and Fishburn, 2007] Steven J. Brams and Peter C. Fishburn. *Approval Voting*. Springer, 2nd edition, 2007.
- [Brandl and Peters, 2022] Florian Brandl and Dominik Peters. Approval voting under dichotomous preferences: A catalogue of characterizations. *Journal of Economic Theory*, 205:105532, 2022.
- [Brill and Peters, 2023] Markus Brill and Jannik Peters. Robust and verifiable proportionality axioms for multiwinner voting. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC)*, page 301, 2023.
- [Fishburn, 1979] Peter C. Fishburn. Symmetric and consistent aggregation with dichotomous voting. In Jean Jacques Laffont, editor, *Aggregation and Revelation of Preferences*, chapter 11, pages 201–218. North-Holland, 1979.
- [Halpern *et al.*, 2025] Daniel Halpern, Ariel D. Procaccia, and Warut Suksompong. The proportional veto principle for approval ballots. *CoRR*, abs/2505.01395, 2025.
- [Ianovski and Kondratev, 2021] Egor Ianovski and Aleksei Y. Kondratev. Computing the proportional veto core. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, pages 5489–5496, 2021.
- [Kizilkaya and Kempe, 2023] Fatih Erdem Kizilkaya and David Kempe. Generalized veto core and a practical voting rule with optimal metric distortion. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC)*, pages 913–936, 2023.
- [Kondratev and Ianovski, 2024] Aleksei Y. Kondratev and Egor Ianovski. Veto core consistent preference aggregation. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 1020–1028, 2024.
- [Laslier and Sanver, 2010] Jean-François Laslier and M. Remzi Sanver. *Handbook on Approval Voting*. Springer, 1st edition, 2010.
- [Moulin, 1981] Hervé Moulin. The proportional veto principle. *Review of Economic Studies*, 48(3):407–416, 1981.
- [Moulin, 1982] Hervé Moulin. Voting with proportional veto power. *Econometrica*, 50(1):145–162, 1982.
- [Peters *et al.*, 2021] Dominik Peters, Grzegorz Pierczyński, and Piotr Skowron. Proportional participatory budgeting with additive utilities. In *Proceedings of the 34th Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 12726–12737, 2021.
- [Peters, 2023] Jannik Peters. A note on rules achieving optimal metric distortion. *CoRR*, abs/2305.08667, 2023.
- [Sánchez-Fernández *et al.*, 2017] Luis Sánchez-Fernández, Edith Elkind, Martin Lackner, Norberto Fernández, Jesús A. Fisteus, Pablo Basanta-Val, and Piotr Skowron. Proportional justified representation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, pages 670–676, 2017.
- [Zwicker, 2016] William S. Zwicker. Introduction to the theory of voting. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 2, pages 23–56. Cambridge University Press, 2016.