k-ApprovalVeto: A Spectrum of Voting Rules Balancing Metric Distortion and Minority Protection

Fatih Erdem Kizilkaya
University of Southern California
Los Angeles, USA
fatih.erdem.kizilkaya@gmail.com

David Kempe
University of Southern California
Los Angeles, USA
david.m.kempe@gmail.com

ABSTRACT

In the context of single-winner ranked-choice elections between m candidates, we explore the tradeoff between two principles that are essential to constitutional democracies: the *majority principle* (maximizing the social welfare) and the *minority principle* (safeguarding minority groups from overly bad outcomes). To measure the social welfare, we use the well-established framework of *metric distortion* subject to various objectives: *utilitarian* (i.e., total cost), *α-percentile* (e.g., median cost for $\alpha = 1/2$), and *egalitarian* (i.e., max cost). To measure the protection of minorities, we introduce the *k-Droop minority criterion*, which requires that if a sufficiently large (parametrized by k) coalition T of voters ranks all candidates in S at the bottom (in any order), then none of the candidates in S should win. The parameter k allows the criterion to interpolate between the minimal requirement that the winner must not be ranked *last* by a strict majority (when $k = 1$) and the strongest protection from bottom choices (when $k = m - 1$). The highest k for which the criterion is satisfied provides a well-defined measure of *minority protection* (ranging from 0 to $m - 1$).

Our main contribution is the analysis of a recently proposed class of voting rules called k-ApprovalVeto, offering a comprehensive range of trade-offs between the two principles. This class spans between PluralityVeto (for $k = 1$) — a simple rule achieving optimal metric distortion — and VoteByVeto (for $k = m$) which picks a candidate from the proportional veto core. We show that k-ApprovalVeto has minority protection at least $k - 1$, and thus, it accommodates any desired level of minority protection via the parameter k. However, this comes at the price of lower social welfare. For the utilitarian objective, the metric distortion becomes $2 \cdot \min(k + 1, m) - 1$, i.e., increases linearly in k. For the α-percentile objective, the metric distortion is the optimal value of 5 for $\alpha \geq k/(k + 1)$ and unbounded for $\alpha < k/(k + 1)$, i.e., the range of α for which the rule achieves optimal distortion becomes smaller. For the egalitarian objective, the metric distortion is the optimal value of 3 for all values of k.

KEYWORDS

Metric Distortion; Minority Protection; Veto Core; Plurality Veto

ACM Reference Format:
Fatih Erdem Kizilkaya and David Kempe. 2025. k-ApprovalVeto: A Spectrum of Voting Rules Balancing Metric Distortion and Minority Protection.

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

In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Detroit, Michigan, USA, May 19 – 23, 2025*, IFAAMAS, 9 pages.

> *All, too, will bear in mind this sacred principle, that though the will of the majority is in all cases to prevail, that will to be rightful must be reasonable; that the minority possess their equal rights, which equal law must protect, and to violate would be oppression.*
> Thomas Jefferson (in his First Inaugural Address)

> *If a majority be united by a common interest, the rights of the minority will be insecure.*
> James Madison (Federalist No. 51)

1 INTRODUCTION

The foundation of democracy is the *majority principle* — the ideal that a decision should be based on the opinion of the majority, ensuring that the outcome is as good as possible for as many individuals as possible. However, following this principle in and of itself poses an immediate risk, commonly known as the majority tyranny, in which the majority pursues exclusively its own objectives at the expense of the interests of the minority factions. Thus, constitutional democracies also incorporate the *minority principle* — the ideal that the authority of the majority should be limited to protect individuals or groups from overly bad outcomes. One of the main challenges in building and maintaining a democracy is to find the right balance between the contradictory factors of the majority and minority principles, as pointed out by the third and fourth presidents of the United States at the opening.

In this paper, we contribute to the crucial effort towards finding a balance between the two principles by presenting an in-depth analysis of k-ApprovalVeto [38] — a spectrum of simple voting rules providing different trade-offs between the two principles. k-ApprovalVeto delegates k approval and veto votes to each of the n voters; k is an integer parameter between 1 and the number of candidates m. These votes are processed as follows:

Approval Votes First, each voter approves (their most favorite) k candidates. As a result, each candidate c starts with a score equal to their *k-approval score* — the number of voters who have c among their top k choices.

Veto Votes Then, the nk veto votes are processed one by one in an arbitrary order. A veto vote of voter v starts from v's bottom choice (i.e., the candidate ranked lowest by v) among not-yet-eliminated candidates and eliminates those whose score is 0. When the veto vote arrives at a candidate c with positive score, it decrements c's score by 1 and terminates, i.e., c is not eliminated even if the score is now 0.

k-ApprovalVeto declares candidates who are not eliminated until the end as (tied) winners. Note that the exact choice of winner(s) might differ based on the order in which veto votes are processed.[1] We are interested in the set of all candidates who can emerge as winners for some veto order. We refer to these winners as the *k-approval veto core* due to a characterization by Kizilkaya and Kempe [38] (given in Section 2) reminiscent of definitions of "core" throughout social choice and game theory. In fact, the possible winners of m-ApprovalVeto are exactly those in the *proportional veto core* [44], discussed below.

As discussed above, we are interested in how k-ApprovalVeto trades off between the majority and minority principles as the parameter k is varied. As primary motivation, we consider guarantees that can be obtained at the extremes $k = m$ and $k = 1$.

- The m-approval veto core is the same as the proportional veto core (defined in Section 2), which consists of all candidates that are not "disproportionately bad" for any coalition of voters large enough to be considered a minority group. Specifically, if an α-fraction of voters prefer a $1-\alpha$ fraction of candidates over candidate c, then c is excluded from the proportional veto core. This definition coincides with the strongest possible guarantee for protection of minorities from worst outcomes (bottom choices), namely, that the winner not be ranked last by strictly more than n/m voters. This is the best possible guarantee as the rankings might be divided into n/m groups with distinct bottom choices.
- At the other extreme, 1-ApprovalVeto equals PluralityVeto[2] [37] which emerged from a line of research towards designing a voting rule with optimal *metric distortion* [3, 29, 30, 47]. The key assumption in metric distortion is that voters and candidates are jointly embedded in an unknown metric space such that candidate c is closer to voter v than candidate c' if and only if v prefers c to c'.[3] Among all such metric spaces, the worst-case ratio between the total distance of voters to a candidate c and that of an optimal candidate is referred to as the metric distortion of c (defined formally in Section 4). PluralityVeto always returns a candidate with metric distortion at most 3, which is the best possible guarantee. Given that the objective is the sum of all distances, PluralityVeto is prone to *majority tyranny* as alluded to above, i.e., it sacrifices minorities for the benefit of the majority, and thus represents an extreme case of ignoring the minority principle.

These two extreme cases achieve the best possible guarantees, respectively, for minorities and the majority, albeit under two very different frameworks. The former follows an *axiomatic* approach while the latter follows a *welfarist* approach. In order to understand the tradeoff as k is varied, we consider generalizations of both types of guarantees as follows.

- As a parameterized relaxation of the minority protection guaranteed by the proportional veto core, we introduce the notion of *Droop*[4] *minority protection*. Recall that the proportional veto core offers the strongest possible minority protection from bottom choices: the winner cannot be ranked last by strictly more than n/m voters. We generalize this requirement by not only considering bottom choices but the full rankings. Our definition is inspired by the classical notion of *solid coalitions* [23]. We say that a coalition T of voters *solidly vetoes* a subset S of candidates if all voters in T prefer all candidates not in S over all candidates in S. We then define the *k-Droop minority criterion* as follows: if a subset S of candidates is solidly vetoed by a coalition T with size exceeding the *k-Droop quota of S* (i.e., $|T|/n > |S|/(k+1)$), then *no* candidate in S should win (because otherwise the outcome would be disproportionately bad for voters in T). The parameter k determines the strictness of the requirement. For $k = m-1$, the criterion implies the strongest possible minority protection from bottom choices, similar to the proportional veto core. On the other hand, for $k = 1$, it only imposes the minimal requirement that the winner must not be ranked last by more than half of the voters. We refer to the largest k for which the criterion is satisfied as the Droop minority protection. (See Section 3 for the formal definitions.)
- PluralityVeto, or any other rule with optimal metric distortion, is prone to majority tyranny due to the choice of the *social cost function*.[5] The framework can be adapted to use other social cost functions; indeed, Anshelevich et al. [3] also considered the *median social cost* (i.e., the median of voters' distances to the candidate) for this very reason — focusing on the median voter reduces the impact of outliers, i.e., voters with very high or very low costs. More generally, Anshelevich et al. [3] consider the *α-percentile social cost*: for any given $\alpha \in [0, 1)$, the distance of the $\lfloor \alpha \cdot n + 1 \rfloor^{\text{th}}$ closest voter to the candidate, which captures the implicit goal of protecting minorities comprising a $1 - \alpha$ fraction of the population. The most protection to minorities — namely, every single individual — is provided by the *egalitarian social cost* [5], i.e., the maximum distance of all voters to the candidate. (See Section 4 for the formal definitions.)

Our Contributions

Our main contribution is an analysis of k-ApprovalVeto, for all values of k, following both the axiomatic and welfarist approaches, respectively, via the Droop minority protection notion and the metric distortion framework.

The main result of our *axiomatic* analysis (presented in Section 3) is that every candidate in the k-approval veto core (i.e., every possible winner of k-ApprovalVeto) has Droop minority protection at least $k - 1$. This confirms that as k increases, the protection of minorities under k-ApprovalVeto increase gradually.

We complement the axiomatic analysis with a welfarist analysis (presented in Section 4) using the metric distortion framework with respect to various objectives (social cost functions) discussed earlier. Our main results here can be summarized as follows:

[1] As we briefly discuss in Section 5, one can also process the veto votes simultaneously to avoid the arbitrariness due to the choice of veto order.

[2] To be precise, 1-ApprovalVeto is equivalent to an extension of PluralityVeto allowing for tied outcomes, introduced in [38]; the only difference is that the original rule immediately eliminates all candidates whose score reaches 0, and thus, it insists on picking a single winner by declaring the last eliminated candidate as the winner, even for elections with obvious ties.

[3] The motivation here is that the distance between a voter and a candidate represents how much their opinions/positions on key issues differ, or simply, the *cost* incurred by the voter if the candidate is elected. This generalizes the classical notion of *single-peaked* preferences [12, 21, 43], which considers embeddings specifically on a line instead of a general metric space.

[4] The naming is in analogy to the notion of the *Droop proportionality criterion* [22] which is closely related to our notion, as shown in Section 3.

[5] That is, the total distance of voters to a given candidate, which is broadly referred to as the *utilitarian social cost* function in economics.

Theorem 1. *For every $k \in \{1, \ldots, m\}$, the following results hold:*

1. *Every candidate in the k-approval veto core has metric distortion at most $2\min(k+1, m) - 1$ (with respect to the utilitarian social cost, i.e., the sum of distances), and this bound is tight.*
2. *For every $\alpha \geq \frac{k}{k+1}$, every candidate in the k-approval veto core has α-percentile metric distortion at most 5, and this bound is tight.*
3. *For every $\alpha < \frac{k}{k+1}$, there exist instances in which a candidate in the k-approval veto core has unbounded α-percentile metric distortion.*
4. *Every candidate in the k-approval veto core has egalitarian metric distortion at most 3, and this bound is tight.*

The first result above shows that as k increases, the k-approval veto core sacrifices the welfare of the majority in return for higher minority protection, as captured by the axiomatic analysis. Confirming this intuition, the second and third results show that as k increases, the k-approval veto core only optimizes for increasingly "egalitarian" objectives, aiming to protect minorities comprising at most a $\frac{1}{k+1}$ fraction of the population, while offering no approximation guarantees for larger groups. To match the upper bound of 5 for the case $\alpha \geq \frac{k}{k+1}$, we prove that for every $\alpha \in [1/2, 1)$, and *every deterministic voting rule*, there exist inputs on which the α-percentile distortion of the winner is arbitrarily close to 5. This lower bound had been previously shown by Anshelevich et al. [3] only for the range $\alpha \in [1/2, 2/3)$; we thus close one of the remaining gaps for (deterministic) metric distortion. Lastly, the fourth result shows that one can do better in the egalitarian setting, and the k-approval veto core achieves the optimal distortion of 3, for all k.

A noteworthy corollary of Theorem 1 (combined with our improved lower bound given in Theorem 7) is that PluralityVeto has optimal metric distortion for the utilitarian, median and egalitarian objectives, as well as for anything in between, i.e., for the α-percentile objective for any $\alpha \in [1/2, 1)$.

Related Work. The notion of distortion was initially studied in the setting of (normalized) *utilities*, i.e., each voter has non-negative utilities for candidates, adding up to 1 [13, 16, 50, 51]. However, without further assumptions on the structure of the utilities, the distortion of all voting rules can be very high [15]. One very natural and fruitful type of restriction was proposed by Anshelevich et al. [3, 4], who imposed a *metric* structure. This metric structure is much more naturally understood when the utilities are negative, or equivalently, when we interpret the voters as having *costs* for different candidates which are characterized by their distances. The fact that distances must obey the triangle inequality restricts the structure of costs, and allows different voting rules to exhibit a rich range of different distortion values. Among the key results of Anshelevich et al. [3, 4] was a (tight) bound of 5 on the metric distortion of the Copeland rule, both with respect to utilitarian and α-percentile objectives, for any $\alpha \in [1/2, 1)$. These were (nearly) matched by a lower bound of 3 for the utilitarian objective, and a lower bound of 5 for the α-percentile objective for $\alpha \in [1/2, 2/3)$.[6] For larger α, the established lower bound was 3.[7]

[6] For $\alpha < 1/2$, the α-percentile distortion of any deterministic voting rule can be easily seen to be unbounded.

[7] For $\alpha \geq \frac{m-1}{m}$, Anshelevich et al. [3] also gave an upper bound of 3 by showing that plurality voting achieves this bound.

The gap between the upper bound of 5 and the lower bound of 3 on the utilitarian metric distortion of deterministic voting rules inspired a significant thread of research work. Initially, the ranked pairs rule was conjectured to achieve distortion 3, which was disproved by Goel et al. [30] who gave a lower bound of 5, and also by Kempe [35] who strengthened the lower bound to $\Omega(\sqrt{m})$. The first improvement was due to Munagala and Wang [47], who achieved distortion $2 + \sqrt{5} \approx 4.23$ using a novel asymmetric variant of the Copeland rule. Building on the work of Munagala and Wang [47] and Kempe [35], the gap was finally closed by Gkatzelis et al. [29], who showed that the voting rule PluralityMatching achieves distortion 3. The voting rule is an exhaustive search for a candidate whose so-called "domination graph" (see Section 2) has a perfect matching. Subsequently, Kizilkaya and Kempe [37] showed that a much simpler voting rule, called PluralityVeto, achieves the same guarantee of 3; the rule implicitly constructs a perfect matching in the domination graph witnessing the distortion guarantee, leading to a much shorter proof. Recall that, up to some subtle tie breaking issues, PluralityVeto equals 1-ApprovalVeto; the only difference is that the former immediately eliminates all candidates whose score reaches 0, and thus, insists on picking a single winner.

A more general class of voting rules, based on selecting winners who have weighted bipartite matchings in a more general domination graph, was already considered by Gkatzelis et al. [29], and also by Kizilkaya and Kempe [37]. The connection to the proportional veto core (Definition 3) was observed by Peters [48] and explored in more depth by Kizilkaya and Kempe [38], who established — for general weights — the equivalence between a general version of k-ApprovalVeto, matchings in generalized domination graphs, and a general definition of the (p, q)-veto core with weights p, q (Definition 4). The equivalence to matchings was observed for the case of the proportional veto core by Ianovski and Kondratev [32] who use this equivalence to compute the proportional veto core in polynomial time. The utilitarian metric distortion of candidates in the general (p, q)-veto core was investigated in more depth by Berger et al. [11], motivated in part by the goal of selecting candidates with lower distortion when an oracle provides advice. We draw significantly on the results of Berger et al. [11] for the analysis of utilitarian metric distortion of candidates in the k-approval veto core; we also utilize one of their lemmas for the α-percentile objective.

Various other considerations have played a role in the analysis of metric distortion. The use of randomization in the selection of a winner can significantly improve the metric distortion: an upper bound of $3 - o(1)$ and a lower bound of 2 had been known from the early work on distortion [8, 36]. In recent breakthrough results, both the upper and lower bounds have been improved [17, 18]. Several works have achieved better bounds on metric distortion when voters can communicate additional information beyond a ranking, such as (limited) information about the *strengths* of their preferences [1, 2, 5]. More generally, the tradeoff between communication and distortion in voting rules has been considered [10, 26, 36, 42, 49]. We refer to the surveys by Anshelevich et al. [6, 7] for further discussion on distortion.

A sequential veto-based mechanism, called *Vote by Veto*, was first studied formally by Mueller [46]. Moulin [44] generalized this mechanism from individuals to coalitions to study the core of the resulting cooperative game, and thus introduced the proportional veto core. Subsequently, Moulin [45] proposed a rule electing from the proportional veto core (equivalent to m-ApprovalVeto up to ties), and studied the strategic behavior of voters under this rule. Ianovski and Kondratev [32] (see also the expanded version [33]) showed how to compute the proportional veto core in polynomial time and introduced an anonymous and neutral rule electing from the proportional veto core. More recently, Kondratev and Ianovski [39] also studied the axiomatic properties of voting rules picking a candidate from the proportional veto core, which also have found applications in windfarm location [28], nuclear fuel disposal [20], and federated learning [19]. To some degree, the veto core is also related to the notion of proportional fairness, which measures whether the influence of cohesive groups of voters on the outcome is proportional to the group size [25].

The trade-off between majority and minority principles has also been a central topic of interest in other works. Kondratev and Nesterov [40] generalized the majority criterion and the majority-loser criterion for this purpose, and analyzed various rules under these axioms. However, these criteria do not allow for a comparison of voting rules that satisfy the mutual majority criterion; in fact, their analysis concludes that all such rules are "best" under both principles, thus failing to capture any trade-off. Faliszewski et al. [27] follow a more similar approach to our work and explore a continuous spectrum between k-Borda and Chamberlin-Courant voting rules in the multi-winner setting. These two voting rules, respectively, represent the two extremes of the majority and minority principles in the multi-winner world. For a detailed overview of multi-winner elections, see [41].

2 PRELIMINARIES

An *election* $\mathcal{E} = (V, C, \overrightarrow{\succ})$ consists of a set of n *voters* V, a set of m *candidates* C and *rankings* $\overrightarrow{\succ} = (\succ_v)_{v \in V}$. In this notation, $\succ_v$ is the *ranking of voter* v, i.e., a total order over C which represents the preferences of v. We write $a \succ_v b$ to express that voter v prefers candidate a over candidate b; we write $a \succcurlyeq_v b$ if $a = b$ or $a \succ_v b$, and say that v *weakly prefers* a over b. We also extend this notation to *coalitions*, i.e., non-empty subsets of voters. By $a \succ_T b$, we denote that every voter in the coalition T prefers a over b; we write $A \succ_T B$ if $a \succ_T b$ for all $a \in A$ and $b \in B$. The complement symbol is always used with respect to the "obvious" ground set, i.e., $\overline{T} = V \setminus T$ if $T \subseteq V$, and $\overline{S} = C \setminus S$ if $S \subseteq C$.

The *top k choices* of voter v, denoted by $\text{top}_k(v)$, are the set of k candidates that v prefers over all other candidates; we also use $\text{top}(v)$ to denote the *top choice* of v, i.e., $\text{top}_1(v) = \{\text{top}(v)\}$. The *plurality score* of a candidate c, denoted by $\text{plu}(c)$, is the number of voters whose top choice is c. The *k-approval score* of a candidate c, denoted by $k\text{-apv}(c)$, is the number of voters who have c among their *top k choices*, i.e., $k\text{-apv}(c) = |\{v \in V \mid c \in \text{top}_k(c)\}|$ and $\text{plu}(c) = 1\text{-apv}(c)$. Extending this notation to sets via addition, we define $k\text{-apv}(S) = \sum_{c \in S} k\text{-apv}(c)$ for $S \subseteq C$. Similarly, we write $\text{bot}_A(v)$ for the least preferred candidate of voter v among a subset of candidates A.

A voting rule $\mathcal{A}$ is an algorithm that, given an election $\mathcal{E}$ as input, returns a non-empty set of candidates $\mathcal{A}(\mathcal{E}) \subseteq C$. We refer to $\mathcal{A}(\mathcal{E})$ as the *(tied) winners* of $\mathcal{E}$ under $\mathcal{A}$, or simply as the winners under $\mathcal{A}$, when $\mathcal{E}$ is clear from the context.

k-Approval Veto Core

In game theory, the term "core" is generally used to refer to a set of outcomes that are not "blocked" by any coalition. In general games, a coalition is said to block an outcome if the members jointly prefer to deviate to another outcome. Collective decisions can also be thought of as outcomes (of a game played between voters) that can be blocked by coalitions that are sufficiently large, which provides a game-theoretical perspective on social choice [24].

We refer to the set of all possible winners of k-ApprovalVeto (outlined in the introduction and specified precisely in Algorithm 1) as the *k-approval veto core*, because they can be characterized via a core definition (Definition 1 below) by a result of Kizilkaya and Kempe [38]. Note that the choice of winner(s) for k-ApprovalVeto depends on the given *veto order* — a repeated sequence of voters $\sigma = (v_1, \ldots, v_{nk})$ in which each voter occurs exactly k times.

Algorithm 1 k-ApprovalVeto with veto order $(v_1, \ldots, v_{nk})$.

```
initialize the set of eligible winners W = C
initialize score(c) = k-apv(c) for all c ∈ C
for i = 1, ..., nk do
    while score(bot_W(v_i)) = 0 do
        remove bot_W(v_i) from W
    decrement score(bot_W(v_i)) by 1
return W
```

As mentioned above, the winners under k-ApprovalVeto can be characterized by a definition of a notion of core. This definition is a natural generalization of the *proportional veto core* of Moulin [44] as discussed below.

Definition 1 (k-Approval Veto Core [38])**.** A coalition $T \subseteq V$ of voters *k-blocks* candidate w with *witness set* $S \subseteq C$ if and only if $S \succ_T w$ and

$$\frac{|T|}{n} > 1 - \frac{k\text{-apv}(S)}{nk}.$$

The *k-approval veto core* of an election $\mathcal{E}$ is the set of all candidates not k-blocked by any coalition, which we denote by $\mathbb{AVC}_k(\mathcal{E})$. (We drop $\mathcal{E}$ from the notation when it is clear from the context.)

Another characterization of the possible winners can be obtained via *k-domination graphs*[8] which are the main technical tool we use to analyze the metric distortion of k-ApprovalVeto in Section 4.

Definition 2 (k-Domination Graphs [29])**.** The *k-domination graph* of candidate w is a bipartite graph $G_k(w)$ defined between k copies of each voter v and $k\text{-apv}(c)$ copies of each candidate c. The graph $G_k(w)$ contains an edge between (each copy of) v and (each copy of) c if and only if $w \succcurlyeq_v c$.

[8]A more general definition then k-domination graphs was first given in [29]. The definition crystallized earlier similar definitions [35, 47] into a more concise form. 1-domination graphs were the key tool in showing that PluralityMatching [29] and PluralityVeto [37] achieve the optimal utilitarian metric distortion of 3.

Definition 1 and Definition 2 characterize the possible winners of k-ApprovalVeto by means of Theorem 2 [38].

Theorem 2 (Kizilkaya and Kempe [38], Theorems 1–3)**.** *For every k and candidate w, the following three statements are equivalent:*

(1) w is a winner of k-ApprovalVeto for some veto order.
(2) No coalition k-blocks w, i.e., $w \in \mathbb{AVC}_k$.
(3) $G_k(w)$ has a perfect matching.

We remark here that $\mathbb{AVC}_{m-1} \subseteq \mathbb{AVC}_m$, because any perfect matching in $G_{m-1}(w)$ can be extended to a perfect matching in $G_m(w)$ by matching the m^{th} copy of voter v to v's bottom choice. This construction does not extend to other k, because the edge from v to the k^{th} ranked choice of v may not exist in $G_k(w)$ for $k < m$. We also remark that in general, the inclusion $\mathbb{AVC}_{m-1} \subset \mathbb{AVC}_m$ can be strict. For example, for 3 candidates $\{a, b, c\}$ and 12 voters of whom 7 rank $a \succ b \succ c$ and 5 rank $b \succ a \succ c$, we have that $\mathbb{AVC}_2 = \{a\}$ while $\mathbb{AVC}_3 = \{a, b\}$.

Related Notions. The classical notion of *proportional veto core* by Moulin [44] (given in Definition 3) is equivalent to the m-approval veto core because $m\text{-apv}(S) = n \cdot |S|$ for all $S \subseteq C$.

Definition 3 (Proportional Veto Core[9] [44])**.** A coalition $T \subseteq V$ of voters *blocks* candidate c with *witness set* $S \subseteq C$ if $S \succ_T c$ and

$$\frac{|T|}{n} > 1 - \frac{|S|}{m}.$$

The *proportional veto core* is the set of all candidates not blocked by any coalition.

The k-ApprovalVeto voting rule, the notion of the k-approval veto core, and the notion of k-domination graphs can all be generalized further, to fractional weights and non-uniform weights not only for candidates, but also for voters. A general definition and proof of a general version of Theorem 2 were given by Kizilkaya and Kempe [38]. The key definition is the following, for any normalized vectors $\boldsymbol{p}, \boldsymbol{q}$ of weights over voters and candidates, respectively.

Definition 4 (($\boldsymbol{p}, \boldsymbol{q}$)-Veto Core [38])**.** A coalition $T \subseteq V$ of voters $(\boldsymbol{p}, \boldsymbol{q})$*-blocks* candidate c with *witness set* $S \subseteq C$ if $S \succ_T c$ and $p(T) > 1 - q(S)$. The $(\boldsymbol{p}, \boldsymbol{q})$ *veto core* is the set of all candidates not $(\boldsymbol{p}, \boldsymbol{q})$-blocked by any coalition.

The k-approval veto core equals the $(\boldsymbol{p}, \boldsymbol{q})$-veto core with $p(v) = \frac{1}{n}$ for all $v \in V$ and $q(c) = \frac{k\text{-apv}(c)}{nk}$ for all $c \in C$.

3 MINORITY PROTECTION

In this section, we introduce the k-Droop minority criterion, which we then use to measure the opposition of minorities to a particular candidate being chosen as the winner. Our main result in this section is that k-ApprovalVeto provides better protection to preferences of minorities as k increases. We will model minorities as solid coalitions by deriving inspiration from the well-known notion of *proportionality for solid coalitions*

Proportionality for Solid Coalitions

First, consider a single-winner election in which each voter can only communicate their top choice. The widely accepted *majority criterion* requires that if candidate c receives a strict majority of votes, i.e., $\text{plu}(c)/n > 1/2$, then c must win. This criterion can be naturally generalized to the election of (a committee of) k winners, by requiring that any candidate c with $\text{plu}(c)/n > 1/(k+1)$ must be selected as a winner. Notice that there can be at most k such candidates.[10] The number $n/(k+1)$ is known as the *Droop quota*, and the requirement that any candidate with a number of votes exceeding the Droop quota be included in the committee (of size k) is known as the *Droop proportionality criterion* [22].

When full rankings are available, stronger requirements can be formulated. For example, if many voters have the same top 2 choices (possibly in different orders), then both candidates should be part of the committee. This idea can be captured more generally via the notion of *solid coalitions*, popularized by Dummett [23].[11]

Definition 5 (Solid Support [23])**.** A coalition $T \subseteq V$ *solidly supports* a subset of candidates $S \subseteq C$ if $S \succ_T \overline{S}$, i.e., all voters in T have S as their top choices (though possibly in different orders).

Note that each singleton coalition $\{v\}$ solidly supports $\text{top}(v)$; thus, the relation of T solidly supporting S is an extension of the relation of c being the top choice of v.

Definition 6 (Weak[12] Droop Proportionality Criterion [9, 34, 48])**.** A committee $K \subseteq C$ of size k satisfies the *weak Droop proportionality criterion (for solid coalitions)* if, for all subsets S of candidates solidly supported by a coalition T with size exceeding the Droop quota of S (i.e., $|T|/n > |S|/(k+1)$), all candidates in S are in K (i.e., $S \subseteq K$).

This criterion is closely related to the criterion we propose for minority protection. Indeed, we will show in Proposition 1 that they are equivalent when the rankings of all voters are reversed; but first, we provide the intuition behind our criterion, and formally define it in Definition 7.

k-Droop Minority Criterion

Recall that any candidate ranked last by more than n/m voters cannot be in the proportional veto core. This guarantee protects minority groups (whose size exceeds n/m) from the worst possible outcome for the group. Analogously to the majority criterion, this is the best guarantee one can hope for when each voter can only communicate their bottom choice (instead of top choices) because there might be n/m distinct groups with different bottom choices. As for the weak Droop proportionality criterion, we can generalize this basic requirement via solid coalitions when full rankings are available (rather than just bottom choices); e.g., if many voters have the same bottom two choices, then they both should not be elected.

However, this strong minority protection comes at a high cost in terms of the metric distortion, a good stand-in for social welfare.

[9] We slightly rearrange the definition in [44] here for clarity.

[10] Of course, it is possible that no candidate satisfies this criterion, just as it is possible that no candidate achieves a majority in a single-winner election.

[11] Solid coalitions are used extensively in social choice theory, most notably in the well-known *mutual majority criterion*; indeed, the mutual majority criterion is exactly the single-winner case of the weak Droop proportionality criterion (Definition 6).

[12] The standard (non-weak) version [9] requires that for all subsets S of candidates solidly supported by a coalition T with size $|T|/n > \ell/(k+1)$ where $\ell \leq |S|$, at least ℓ candidates in S should be in the committee.

This motivates us to propose a parameterized definition of minority protection which interpolates smoothly between the strong requirements imposed by the proportional veto core and essentially no protection of the minority. Such a definition will then allow us to study the tradeoffs between protection of minorities and welfare of the majority, quantifying and analyzing the tradeoffs outlined in the initial quotes of Jefferson and Madison. In the same spirit as Definition 5 and Definition 6, we define solidly vetoing coalitions and the k-Droop minority criterion, respectively, as follows:

Definition 7 (Solid Veto and Droop Minority Criterion)**.** Given an election $\mathcal{E} = (V, C, \overrightarrow{\succ})$:

(1) A coalition $T \subseteq V$ *solidly vetoes* $S \subseteq C$ if $\overline{S} \succ_T S$.
(2) A candidate w satisfies the *k-Droop minority criterion (for solid coalitions)* if the following holds for all subsets S of candidates: if S is solidly vetoed by a coalition T with size exceeding the k-Droop quota of S (i.e., $|T|/n > |S|/(k+1)$), then $w \notin S$.

We write $\mathbb{DMC}_k(\mathcal{E})$ to denote the set of all candidates satisfying the k-Droop minority criterion. (We will drop $\mathcal{E}$ from the notation when it is clear from the context.)

While the parameter k is originally inspired by the size of the committee, in Definition 7, it is used as a measure of how strictly minorities should be protected. The relationship to the committee size, and between the k-Droop minority criterion and the Droop proportionality criterion, are captured by the following proposition. The proof is deferred to the full version of this paper.

Proposition 1. *A committee K of size k satisfies the weak Droop proportionality criterion if and only if all candidates in $\overline{K}$ satisfy the k-Droop minority criterion in the election with reversed rankings.*

In Definition 7, if T exceeds the k-Droop quota of S, then it also exceeds the ℓ-Droop quota of S for all $\ell > k$. Thus, $\mathbb{DMC}_k \subseteq \mathbb{DMC}_{k-1}$ for all k, i.e., the k-Droop minority criterion gets harder to satisfy for larger k. In particular, $\mathbb{DMC}_0 = C$ contains all candidates, and $\mathbb{DMC}_m = \emptyset$ (by considering $T = V$ and $S = C$). As a result, for every candidate c, there exists a unique $k \in \{0, \dots, m-1\}$ such that $c \in \mathbb{DMC}_k \setminus \mathbb{DMC}_{k+1}$; this gives rise to a well-defined measure, which we refer to as the *Droop minority protection* of candidate c.

Our main result in this section is the following theorem, showing that higher values of k provide higher Droop minority protection for candidates in the k-approval veto core.

Theorem 3. *The Droop minority protection of every candidate in the k-approval veto core is at least $k-1$, i.e., $\mathbb{AVC}_k \subseteq \mathbb{DMC}_{k-1}$.*

Proof. We show that $\overline{\mathbb{DMC}_{k-1}} \subseteq \overline{\mathbb{AVC}_k}$. Let $c \notin \mathbb{DMC}_{k-1}$. Then, there exists a coalition T solidly vetoing a subset of candidates $S \ni c$ such that $\frac{|T|}{n} > \frac{|S|}{(k-1)+1} = \frac{|S|}{k}$. Now, considering $\overline{S}$, we observe that $\overline{S} \succ_T c$ and $k\text{-apv}(\overline{S}) = nk - k\text{-apv}(S) \geq nk - n|S|$; hence, $\frac{|S|}{k} \geq 1 - \frac{k\text{-apv}(\overline{S})}{nk}$. Substituting this inequality, we obtain that $\frac{|T|}{n} > 1 - \frac{k\text{-apv}(\overline{S})}{nk}$. By Definition 1, this means that T k-blocks c with witness set $\overline{S}$. Hence, $c \notin \mathbb{AVC}_k$. □

The above lower bound of $k-1$ is not tight. Indeed, the Droop minority protection of every candidate in the 1-approval veto core is at least 1. This is because $\mathbb{DMC}_1$ consists of exactly the candidates that are ranked last by at most $n/2$ voters, and candidates that are ranked last by strictly more than $n/2$ voters cannot possibly win under 1-ApprovalVeto = PluralityVeto.

4 METRIC DISTORTION

In this section, we present a complete analysis of the metric distortion of the k-approval veto core. We begin by reviewing the relevant definitions.

Framework

A *metric* over a set S is a function $d : S \times S \to \mathbb{R}_{\geq 0}$ which satisfies the following three conditions for all $a, b, c \in S$: (1) Positive Definiteness: $d(a,b) = 0$ if and only[13] if $a = b$; (2) Symmetry: $d(a,b) = d(b,a)$; (3) Triangle inequality: $d(a,b) + d(b,c) \geq d(a,c)$. Given an election $\mathcal{E} = (V, C, \overrightarrow{\succ})$, we say that a metric d over $V \cup C$ is *consistent* with the rankings $\overrightarrow{\succ}$, and write $d \sim \overrightarrow{\succ}$, if $d(v,c) \leq d(v,c')$ for all $v \in V$ and $c, c' \in C$ such that $c \succ_v c'$.

The metric distortion framework of Anshelevich et al. [3] characterizes the quality of a candidate w (chosen as the *winner*) based on the distances between voters and w. Specifically, we study the following three notions of social cost. Given a candidate w and a metric $d \sim \overrightarrow{\succ}$, (1) the *utilitarian social cost* of w is defined as $\text{cost}_d^+(w) = \sum_{v \in V} d(v,w)$; (2) the *$\alpha$-percentile social cost* of w, for a given $\alpha \in [0,1)$, is defined as $\text{cost}_d^\alpha(w) = d(v_d^\alpha(w), w)$ where $v_d^\alpha(w)$ denotes the $\lfloor \alpha n + 1 \rfloor^{\text{th}}$ closest voter to w under d; and (3) the *egalitarian social cost* of w is defined as $\text{cost}_d^1(w) = \max_{v \in V} d(v,w)$.[14] The metric distortion of a candidate under utilitarian, α-percentile, or egalitarian social cost is defined as follows.

Definition 8 (Metric Distortion)**.** Under the social cost objective $* \in \{+, \alpha, 1\}$, the *(metric) distortion* of a candidate c in an election $\mathcal{E}$ is the largest possible ratio between the social cost of c and that of an optimal candidate c_d^* under any metric $d \sim \overrightarrow{\succ}$. That is,

$$\text{dist}_{\mathcal{E}}^*(c) = \sup_{d \sim \overrightarrow{\succ}} \frac{\text{cost}_d^*(c)}{\text{cost}_d^*(c_d^*)}.$$

We refer to $\text{dist}_{\mathcal{E}}^+(c)$, $\text{dist}_{\mathcal{E}}^\alpha(c)$ and $\text{dist}_{\mathcal{E}}^1(c)$, respectively, as the *utilitarian*, the *α-percentile* and the *egalitarian* (metric) distortion of candidate c. Extending the notion to voting rules, we say that a voting rule $\mathcal{A}$ has distortion (at most) θ if $\text{dist}_{\mathcal{E}}^*(w) \leq \theta$ for every election $\mathcal{E}$ and for every winner w of $\mathcal{E}$ under $\mathcal{A}$, i.e., $\text{dist}^*(\mathcal{A}) = \max_{\mathcal{E}} \max_{w \in \mathcal{A}(\mathcal{E})} \text{dist}_{\mathcal{E}}^*(w)$.

Utilitarian Metric Distortion

We begin by pinning down the most frequently studied notion of distortion, namely, utilitarian metric distortion, of candidates in the k-approval veto core, for all k. The upper bound for $k < m$ follows from a recent bound by Berger et al. [11] for the $(\boldsymbol{p}, \boldsymbol{q})$-veto core (see Definition 4). Our main contribution is therefore an improved bound for $k = m$ (i.e., the proportional veto core, see Definition 3) and a matching lower bound for all k.

[13] Our proofs do not require the "only if" condition, so technically, all our results hold for pseudo-metrics, not just metrics.
[14] Note that the egalitarian social cost is not subsumed by the α-percentile social cost as the latter is not well-defined for $\alpha = 1$.

Theorem 4. *The utilitarian metric distortion of every candidate in the k-approval veto core is at most $2\min(k+1,m)-1$ and this bound is tight, i.e., for all k, there is an election $\mathcal{E}$ with $\text{dist}^+_{\mathcal{E}}(c) = 2\min(k+1,m)-1$ for some $c \in \mathbb{AVC}_k(\mathcal{E})$.*

We begin by proving the upper bound. For $k < m$, as mentioned above, we can use the following result by Berger et al. [11]:

Lemma 1 (Corollary 3.12 of Berger et al. [11])**.** *The distortion of every candidate in the $(\boldsymbol{p}, \boldsymbol{q})$-veto core is at most*

$$1 + \frac{2\max_v p(v)}{\min_c \frac{q(c)}{\text{plu}(c)}}.$$

As discussed previously, the k-approval veto core is the special case $p(v) = \frac{1}{n}$ for all $v \in V$ and $q(c) = \frac{k\text{-apv}(c)}{nk}$ for all $c \in C$; here, the bound of Lemma 1 reduces to $1 + 2k\max_c \frac{\text{plu}(c)}{k\text{-apv}(c)}$. Since $\text{plu}(c) \leq k\text{-apv}(c)$ for all $c \in C$, this implies an upper bound of $2k+1$ for the k-approval veto core, matching the claimed upper bound in Theorem 4 for $k < m$. While the proof follows easily from Lemma 1, the latter has a rather involved proof. (In the full version, we give a simpler proof using the flow technique of Kempe [35].)

To give an improved upper bound of $2m-1$ for the case $k = m$, i.e., for the proportional veto core, we utilize the following lemma of Anshelevich et al. [4]:

Lemma 2 (Anshelevich et al. [4], Lemma 6)**.** *For every pair of candidates $w \neq c^*$ and for all metrics d consistent with the rankings,*

$$\frac{\text{cost}^+_d(w)}{\text{cost}^+_d(c^*)} \leq \frac{2n}{|\{v \mid w \succ_v c^*\}|} - 1.$$

Using this lemma, we prove the bound by contrapositive. For any candidate c such that $|\{v \mid c \succ_v c^*\}| \geq n/m$, Lemma 2 directly implies a distortion of at most $2m-1$. Therefore, consider a candidate c such that fewer than n/m voters prefer c over c^*, so strictly more than $n - (n/m)$ voters prefer c^* over c. Since $m\text{-apv}(c^*) = n$, this means that the coalition of all voters v with $c^* \succ_v c$ m-blocks c with witness set $\{c^*\}$. Hence, c is not in $\mathbb{AVC}_m$.

We complete the proof of Theorem 4 by giving matching lower bounds for all k in Lemma 3; the proof is deferred to the full version.

Lemma 3. *For all k and $\epsilon > 0$, there is an election $\mathcal{E}$ with $\text{dist}^+_{\mathcal{E}}(c) \geq 2\min(k+1,m) - 1 - \epsilon$ for some $c \in \mathbb{AVC}_k(\mathcal{E})$.*

α-Percentile Metric Distortion

In this section, we show that the α-percentile distortion of every candidate in the k-approval veto core is at most 5 for $\alpha \geq k/(k+1)$ (Theorem 5), and unbounded for $\alpha < k/(k+1)$ (Theorem 6). In particular, PluralityVeto has α-percentile distortion at most 5 for all $\alpha \geq 1/2$, which we show to be the best possible bound for any (deterministic) voting rule (Theorem 7). Previously, Anshelevich et al. [3] had shown that the α-percentile distortion of every voting rule is unbounded for $\alpha \in [0, 1/2)$, at least 5 for $\alpha \in [1/2, 2/3)$, and at least 3 for $\alpha \in [2/3, 1)$. Thus, we improve their lower bound from 3 to 5 for $\alpha \in [2/3, 1)$. This establishes that PluralityVeto is an optimal voting rule in terms of α-percentile distortion for all α. The only other rule that is known to achieve constant α-percentile distortion is the Copeland rule, which enjoys the same bound of 5 for all $\alpha \geq 1/2$. Our analysis uses the following two lemmas:

Lemma 4 (Anshelevich et al. [3], Lemma 29)**.** *For every $\alpha \in [0,1)$, and pair of candidates c and c', $\text{cost}^\alpha_d(c) \leq \text{cost}^\alpha_d(c') + d(c,c')$.*

Lemma 5 (Berger et al. [11], Lemma 3.9)**.** *For every pair of voters v and v', and every pair of candidates c and c', if $c \succcurlyeq_v \text{top}(v')$, then for every metric consistent with the rankings,*

$$d(v,c') + d(v',c') \geq \frac{d(c,c')}{2}.$$

Theorem 5. *The α-percentile distortion of every candidate in the k-approval veto core is at most 5 for all $\alpha \geq \frac{k}{k+1}$.*

Proof. Fix some $k \in \{1,\ldots,m\}$ and $\alpha \geq k/(k+1)$. Let $\mathcal{E} = (V, C, \overrightarrow{\succ})$ be an election, and let $d \sim \overrightarrow{\succ}$ be a metric consistent with the rankings. Let $w \in \mathbb{AVC}_k$, and let c^* be an optimal candidate under d. We distinguish two cases, based on whether w and c^* are "close" to each other (compared to the cost of w) or not.

If $d(w,c^*) \leq \frac{4}{5} \cdot \text{cost}^\alpha_d(w)$, then using Lemma 4, we obtain

$$\begin{aligned}\text{cost}^\alpha_d(c^*) &\geq \text{cost}^\alpha_d(w) - d(w,c^*)\\ &\geq \text{cost}^\alpha_d(w) - \frac{4}{5}\cdot\text{cost}^\alpha_d(w) = \frac{1}{5}\cdot\text{cost}^\alpha_d(w).\end{aligned}$$

If $\text{cost}^\alpha_d(w) < \frac{5}{4}\cdot d(w,c^*)$, we show that there are at least $n/(k+1)$ voters v with distance $d(v,c^*) \geq d(w,c^*)/4$, which implies that $\text{cost}^\alpha_d(c^*) \geq d(w,c^*)/4$. Assume for contradiction that this is not the case, i.e., the set $V' = \{v \in V \mid d(v,c^*) < d(w,c^*)/4\}$ contains strictly more than $nk/(k+1)$ voters.

Consider the set of candidates $C' = \{\text{top}(v) \mid v \in V'\}$ which are the top choice of at least one voter in V'. Then, the total plurality score $P := \sum_{c\in C'} \text{plu}(c) \geq |V'| > nk/(k+1)$, and the bipartite graph $G_k(w)$ contains at least P copies of candidates in C' (possibly more, by considering lower rankings when $k > 1$). Furthermore, for any $v, v' \in V'$, we have that $d(v,c^*) + d(v',c^*) < d(w,c^*)/2$. Then, Lemma 5 implies that no voter $v \in V'$ prefers w over $\text{top}(v')$ for any $v' \in V'$. As a result, $G_k(w)$ cannot contain any edges from copies of $v \in V'$ to any copies of candidates $c \in C'$. In other words, the only edges to copies of candidates $c \in C'$ can come from copies of voters $v \notin V'$; and since there are strictly fewer than $n/(k+1)$ voters not in V', there are strictly fewer than $nk/(k+1)$ such copies. On the other hand, we argued above that $G_k(w)$ contains at least $P > nk/(k+1)$ copies of candidates in C'. Thus, we have exhibited a set of P nodes in $G_k(w)$ whose neighborhood contains strictly fewer than P nodes. By Hall's marriage theorem, $G_k(w)$ does not have a perfect matching, contradicting the assumption of the theorem that $w \in \mathbb{AVC}_k$. ◻

We next show matching lower bounds. First, we show (in the full version) that for $\alpha < \frac{k}{k+1}$, the α-percentile distortion of candidates in the k-approval veto core may be unbounded.

Theorem 6. *For all $\alpha < \frac{k}{k+1}$, there exists an election $\mathcal{E}$ such that $\text{dist}^\alpha_{\mathcal{E}}(w) = \infty$ for all candidates $w \in \mathbb{AVC}_k(\mathcal{E})$.*

Next, we show (also in the full version) that no deterministic voting rule can achieve α-percentile distortion smaller than 5, for any $\alpha < 1$. The proof is based on a straightforward extension of the construction of Anshelevich et al. [3].

Theorem 7. *For every (deterministic) voting rule $\mathcal{A}$, and constants $\alpha \in [1/2, 1)$ and $\epsilon > 0$, there exists an election $\mathcal{E}$ and candidate $w \in \mathcal{A}(\mathcal{E})$ such that $\text{dist}^\alpha_{\mathcal{E}}(w) = 5 - \epsilon$.*

Egalitarian Metric Distortion

Here, we show that all candidates in the k-approval veto core, for all k, have egalitarian (metric) distortion at most 3, which is known to be the best possible guarantee [5]. This is based on the key observation that very minimal conditions are enough to ensure egalitarian distortion at most 3; in particular, all Pareto efficient candidates have egalitarian distortion at most 3. We begin by recalling the definition of Pareto domination:

Definition 9 (Pareto domination). A candidate c is *Pareto dominated* by a candidate c' if $c' \succ_v c$ for every voter v. If a candidate c is not Pareto dominated by any candidate c', then c is *Pareto efficient.*

We first adapt (in the full version) the proof of Theorem 30 of Anshelevich et al. [3], and show that if candidate c is not Pareto dominated by candidate c', the egalitarian distortion of c cannot be much higher than that of c'.

Lemma 6. *Let $\mathcal{E} = (V, C, \vec{\succ})$ be an election and $d \sim \vec{\succ}$ be a metric. If c and c' are candidates such that c' does not Pareto dominate c, then $\text{cost}_d^1(c) \le 3 \cdot \text{cost}_d^1(c')$.*

Because any Pareto efficient candidate c satisfies the condition of Lemma 6 for the *optimal* candidate $c^*(= c')$, we immediately obtain the following corollary:

Corollary 1. *Every Pareto efficient candidate c has egalitarian metric distortion at most 3.*

Every candidate in the proportional veto core is Pareto efficient because the grand coalition (of all voters) m-blocks every Pareto dominated candidate: this shows that the egalitarian distortion of every candidate in the proportional veto core is at most 3.

In extending this result to the k-approval veto core for $k < m$, we face the obstacle that the k-approval veto core can contain Pareto-dominated candidates. This was observed for the case $k = 1$ by Kizilkaya and Kempe [38] who showed that such domination was only possible in a very limited sense: both the dominated and dominating candidate had to have plurality score 0. We extend this insight to the k-approval veto core in the following lemma, whose proof is given in the full version.

Lemma 7. *Let $\mathcal{E}$ be an election, and $c \in \mathbb{AVC}_k$ such that c is Pareto-dominated by c'. Then, $k\text{-apv}(c') = 0$.*

We are now ready to state and prove our main theorem on the egalitarian distortion of the k-approval veto core.

Theorem 8. *For every k, the egalitarian distortion of every candidate in the k-approval veto core is at most 3.*

Proof. Fix an election $\mathcal{E} = (V, C, \vec{\succ})$ and a metric $d \sim \vec{\succ}$, and let w be a candidate in the k-approval veto core. Let c^* be an optimal candidate under d. If c^* does not Pareto-dominate w, then the result follows immediately by applying Lemma 6 to w and c^*.

Otherwise, by Lemma 7, we first obtain that $k\text{-apv}(c^*) = 0$. By Theorem 2, there is a perfect matching M of $G_k(w)$. Let v be a most distant voter from w under d. $G_k(w)$ must contain k copies of v, all of which are matched under M (possibly to multiple copies of the same candidate). Let c be any candidate such that v is matched to at least one copy of c. Because $w \succ_v c$ by definition of the edges of $G_k(w)$ and $\text{cost}_d^1(c) \ge d(c, v)$ by definition of egalitarian cost, we can bound $\text{cost}_d^1(w) = d(w, v) \le d(c, v) \le \text{cost}_d^1(c)$.

Next, we observe that because $G_k(w)$ contained a copy of c, and the number of copies of c is $k\text{-apv}(c)$, we get that $k\text{-apv}(c) > 0$, whereas $k\text{-apv}(c^*) = 0$. Therefore, at least one voter (ranking c in the top k positions) prefers c over c^*; in particular, c^* cannot Pareto-dominate c. By Lemma 6, this implies that $\text{cost}_d^1(c) \le 3 \cdot \text{cost}_d^1(c^*)$, and hence $\text{cost}_d^1(w) \le 3 \cdot \text{cost}_d^1(c^*)$. □

5 CONCLUSION AND FUTURE DIRECTIONS

Our analysis shows that as k increases, k-ApprovalVeto sacrifices welfare gradually to enhance minority protection. Along with its simplicity, this makes k-ApprovalVeto potentially practical for settings where it is desirable to balance the majority and minority principles. Therefore, studying other axiomatic properties of k-ApprovalVeto would be of interest. For instance, as stated, the rule violates the essential axiom of *anonymity* (i.e., the requirement that all voters be treated equally a priori). However, by processing veto votes simultaneously, as shown in [38], both anonymity and *neutrality* (i.e., the counterpart of anonymity for candidates) can be satisfied.

Another practical aspect of k-ApprovalVeto is that the parameter k (i.e., the number of approval votes) provides an intuitive means to adjust the desired level of (Droop) minority protection, albeit at the cost of some social welfare. Perhaps the most significant question about this trade-off is whether k-ApprovalVeto achieves the optimal balance, i.e., the minimal loss in welfare to reach the desired level of minority protection. Specifically, we leave determining the best achievable metric distortion by a voting rule satisfying the k-Droop minority criterion as an open question. Exploring the Droop minority protection of other voting rules could be a valuable step towards this challenging goal. On a related note, it would also be of interest to study Droop minority protection in comparison with other notions such as *normalized* distortion [13, 16, 50, 51], and/or in more general settings such as *randomized* voting [8, 17, 18, 31].

A strong assumption in our k-Droop minority criterion is that minorities are modeled as coalitions solidly vetoing a subset of candidates S, i.e., all members of the coalition rank candidates in S at the bottom (in some order). However, this assumption is often unrealistic in practice, as minorities rarely form *perfectly* solid coalitions. Hence, an important direction for future work is to extend the k-Droop minority criterion to accommodate more robust models of minorities. This direction is in parallel with the work of Brill and Peters [14] exploring robust and verifiable proportionality axioms (such as EJR+ and PJR+) in the multi-winner voting setting. A similar approach could be pursued here, as the k-Droop minority criterion is equivalent to the Droop proportionality criterion when the rankings are reversed. In connection with this work, it would also be of interest to investigate the hardness of verifying whether a given candidate satisfies the k-Droop minority criterion. For $k = m - 1$, the verification problem can be solved in polynomial time, as shown by Ianovski and Kondratev [32], since the criterion precisely recovers the proportional veto core.

ACKNOWLEDGMENTS

We would like to thank anonymous reviewers for very detailed and helpful feedback.

REFERENCES

[1] Ben Abramowitz, Elliot Anshelevich, and Wennan Zhu. 2019. Awareness of Voter Passion Greatly Improves the Distortion of Metric Social Choice. In *Proc. 15th Conference on Web and Internet Economics (WINE)*. 3–16.

[2] Georgios Amanatidis, Georgios Birmpas, Aris Filos-Ratsikas, and Alexandros A. Voudouris. 2021. Peeking behind the ordinal curtain: Improving distortion via cardinal queries. *Artificial Intelligence* 296 (2021), 103488.

[3] Elliot Anshelevich, Onkar Bhardwaj, Edith Elkind, John Postl, and Piotr Skowron. 2018. Approximating optimal social choice under metric preferences. *Artificial Intelligence* 264 (2018), 27–51.

[4] Elliot Anshelevich, Onkar Bhardwaj, and John Postl. 2015. Approximating Optimal Social Choice under Metric Preferences. In *Proc. 29th AAAI Conf. on Artificial Intelligence*. 777–783.

[5] Elliot Anshelevich, Aris Filos-Ratsikas, Christopher Jerrett, and Alexandros A. Voudouris. 2024. Improved Metric Distortion via Threshold Approvals. In *Proc. 38th AAAI Conf. on Artificial Intelligence*, Vol. 38. 9460–9468.

[6] Elliot Anshelevich, Aris Filos-Ratsikas, Nisarg Shah, and Alexandros A. Voudouris. 2021. Distortion in social choice problems: an annotated reading list. *SIGecom Exchanges* 19, 1 (2021), 12–14.

[7] Elliot Anshelevich, Aris Filos-Ratsikas, Nisarg Shah, and Alexandros A. Voudouris. 2021. Distortion in Social Choice Problems: The First 15 Years and Beyond. In *Proc. 30th Intl. Joint Conf. on Artificial Intelligence*. 4294–4301.

[8] Elliot Anshelevich and John Postl. 2016. Randomized Social Choice Functions Under Metric Preferences. In *Proc. 25th Intl. Joint Conf. on Artificial Intelligence*. 46–59.

[9] Haris Aziz and Barton E. Lee. 2020. The Expanding Approvals Rule: Improving Proportional Representation and Monotonicity. *Social Choice and Welfare* 54, 1 (2020), 1–45. https://doi.org/10.1007/s00355-019-01208-3

[10] Matthias Bentert and Piotr Skowron. 2020. Comparing Election Methods Where Each Voter Ranks Only Few Candidates. In *Proc. 34th AAAI Conf. on Artificial Intelligence*. 2218–2225.

[11] Ben Berger, Michal Feldman, Vasilis Gkatzelis, and Xizhi Tan. 2024. Learning-Augmented Metric Distortion via (p, q)-Veto Core. In *Proc. 25th ACM Conf. on Economics and Computation*.

[12] Duncan Black. 1948. On the Rationale of Group Decision Making. *J. Political Economy* 56 (1948), 23–34.

[13] Craig Boutilier, Ioannis Caragiannis, Simi Haber, Tyler Lu, Ariel D. Procaccia, and Or Sheffet. 2015. Optimal social choice functions: A utilitarian view. *Artificial Intelligence* 227 (2015), 190–213.

[14] Markus Brill and Jannik Peters. 2023. Robust and Verifiable Proportionality Axioms for Multiwinner Voting. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC '23)*. 301.

[15] Ioannis Caragiannis, Swaprava Nath, Ariel D. Procaccia, and Nisarg Shah. 2017. Subset Selection Via Implicit Utilitarian Voting. *Journal of Artificial Intelligence Research* 58 (Jan. 2017), 123–152.

[16] Ioannis Caragiannis and Ariel D. Procaccia. 2011. Voting almost maximizes social welfare despite limited communication. *Artificial Intelligence* 175, 9 (2011), 1655–1671.

[17] Moses Charikar and Prasanna Ramakrishnan. 2022. Metric Distortion Bounds for Randomized Social Choice. In *Proc. 33rd ACM-SIAM Symp. on Discrete Algorithms*. 2986–3004.

[18] Moses Charikar, Prasanna Ramakrishnan, Kangning Wang, and Hongxun Wu. 2024. Breaking the Metric Voting Distortion Barrier. In *Proc. 35th ACM-SIAM Symp. on Discrete Algorithms*. 1621–1640.

[19] Bhaskar Ray Chaudhury, Aniket Murhekar, Zhuowen Yuan, Bo Li, Ruta Mehta, and Ariel D. Procaccia. 2024. Fair Federated Learning via the Proportional Veto Core. In *Forty-first International Conference on Machine Learning*.

[20] François Diaz-Maurin, Jerold Yu, and Rodney C. Ewing. 2021. Socio-technical multi-criteria evaluation of long-term spent nuclear fuel management strategies: A framework and method. *Science of The Total Environment* 777 (2021), 146086.

[21] Anthony Downs. 1957. An economic theory of political action in a democracy. *The Journal of Political Economy* 65, 2 (1957), 135–150.

[22] H. R. Droop. 1881. On Methods of Electing Representatives. *Journal of the Statistical Society of London* 44, 2 (1881), 141–196.

[23] Michael Dummett. 1984. *Voting Procedures*. Oxford University Press UK.

[24] Michael Dummett and Robin Farquharson. 1961. Stability in Voting. *Econometrica* 29, 1 (1961), 33–43.

[25] Soroush Ebadian, Anson Kahng, Dominik Peters, and Nisarg Shah. 2022. Optimized Distortion and Proportional Fairness in Voting. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC '22)*. 563–600.

[26] Brandon Fain, Ashish Goel, Kamesh Munagala, and Nina Prabhu. 2019. Random Dictators with a Random Referee: Constant Sample Complexity Mechanisms for Social Choice. In *Proc. 33rd AAAI Conf. on Artificial Intelligence*. 1893–1900.

[27] Piotr Faliszewski, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. 2017. Multiwinner Rules on Paths From k-Borda to Chamberlin–Courant. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*. 192–198.

[28] Gonzalo Gamboa and Giuseppe Munda. 2007. The problem of windfarm location: A social multi-criteria evaluation framework. *Energy Policy* 35, 3 (2007), 1564–1583.

[29] Vasilis Gkatzelis, Daniel Halpern, and Nisarg Shah. 2020. Resolving the Optimal Metric Distortion Conjecture. In *Proc. 61st IEEE Symp. on Foundations of Computer Science*. 1427–1438.

[30] Ashish Goel, Anilesh Kollagunta Krishnaswamy, and Kamesh Munagala. 2017. Metric Distortion of Social Choice Rules: Lower Bounds and Fairness Properties. In *Proc. 18th ACM Conf. on Economics and Computation*. 287–304.

[31] Stephen Gross, Elliot Anshelevich, and Lirong Xia. 2017. Vote Until Two of You Agree: Mechanisms with Small Distortion and Sample Complexity. In *Proc. 31st AAAI Conf. on Artificial Intelligence*. 544–550.

[32] Egor Ianovski and Aleksei Y. Kondratev. 2021. Computing the Proportional Veto Core. *Proceedings of the AAAI Conference on Artificial Intelligence* 35, 6 (2021), 5489–5496.

[33] Egor Ianovski and Aleksei Y. Kondratev. 2023. Computing the Proportional Veto Core. arXiv:2003.09153 [cs]

[34] Svante Janson. 2016. Phragmén's and Thiele's election methods. *arXiv preprint arXiv:1611.08826* (2016). https://arxiv.org/abs/1611.08826 95 pages. A serious typo corrected in this version.

[35] David Kempe. 2020. An Analysis Framework for Metric Voting based on LP Duality. In *Proc. 34th AAAI Conf. on Artificial Intelligence*. 2079–2086.

[36] David Kempe. 2020. Communication, Distortion, and Randomness in Metric Voting. In *Proc. 34th AAAI Conf. on Artificial Intelligence*. 2087–2094.

[37] Fatih Erdem Kizilkaya and David Kempe. 2022. PluralityVeto: A Simple Voting Rule with Optimal Metric Distortion. In *Proc. 31st Intl. Joint Conf. on Artificial Intelligence*. 349–355.

[38] Fatih Erdem Kizilkaya and David Kempe. 2023. Generalized Veto Core and a Practical Voting Rule with Optimal Metric Distortion. In *Proc. 24th ACM Conf. on Economics and Computation*. Association for Computing Machinery, 913–936.

[39] Aleksei Y. Kondratev and Egor Ianovski. 2024. Veto Core Consistent Preference Aggregation. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS '24)*. International Foundation for Autonomous Agents and Multiagent Systems, 1020–1028.

[40] Aleksei Y. Kondratev and Alexander S. Nesterov. 2020. Measuring Majority Power and Veto Power of Voting Rules. *Public Choice* 183, 1 (April 2020), 187–210.

[41] Martin Lackner and Piotr Skowron. 2023. *Multi-Winner Voting with Approval Preferences*. Springer International Publishing. https://doi.org/10.1007/978-3-031-09016-5

[42] Debmalya Mandal, Ariel D. Procaccia, Nisarg Shah, and David P. Woodruff. 2019. Efficient and Thrifty Voting by Any Means Necessary. In *Proc. 33rd Advances in Neural Information Processing Systems*.

[43] Hervé Moulin. 1980. On Strategy-Proofness and Single Peakedness. *Public Choice* 35 (1980), 437–455.

[44] Hervé Moulin. 1981. The Proportional Veto Principle. *The Review of Economic Studies* 48, 3 (1981), 407–416.

[45] Hervé Moulin. 1982. Voting with Proportional Veto Power. *Econometrica* 50, 1 (1982), 145–162.

[46] Dennis C. Mueller. 1978. Voting by Veto. *Journal of Public Economics* 10, 1 (Aug. 1978), 57–75.

[47] Kamesh Munagala and Kangning Wang. 2019. Improved Metric Distortion for Deterministic Social Choice Rules. In *Proc. 20th ACM Conf. on Economics and Computation*. 245–262.

[48] Jannik Peters. 2023. A Note on Rules Achieving Optimal Metric Distortion. arXiv:2305.08667 [cs]

[49] Grzegorz Pierczyński and Piotr Skowron. 2019. Approval-Based Elections and Distortion of Voting Rules. In *Proc. 28th Intl. Joint Conf. on Artificial Intelligence*. 543–549.

[50] Ariel D. Procaccia. 2010. Can approximation circumvent Gibbard-Satterthwaite?. In *Proc. 24th AAAI Conf. on Artificial Intelligence*. 836–841.

[51] Ariel D. Procaccia and Jeffrey S. Rosenschein. 2006. The distortion of cardinal preferences in voting. In *Proc. 10th Intl. Workshop on Cooperative Inform. Agents X*. 317–331.