Simple Characterizations of Approval Voting

Florian Brandl
Stanford University
brandl.ffx@gmail.com

Dominik Peters
Carnegie Mellon University
dominikp@cs.cmu.edu

Approval voting allows every voter to cast a ballot of approved alternatives and chooses the alternatives with the largest number of approvals. Due to its simplicity and superior theoretical properties it is a serious contender for use in real-world elections. We support this claim by giving seven characterizations of approval voting based on normatively appealing axioms. All our results involve the reinforcement axiom, which requires choices to be consistent across different electorates. In addition, we consider strategyproofness, consistency with majority opinions, consistency under cloning alternatives, and invariance under removing inferior alternatives. We prove our results by reducing them to a single base theorem, for which we give a simple and intuitive proof.

1. Introduction

Around the world, when electing a leader or a representative, plurality is by far the most common voting system: each voter casts a vote for a single candidate, and the candidate with the most votes is elected. In pioneering work by Brams and Fishburn (1983), they proposed an alternative system: approval voting. Here, each voter may cast votes for an arbitrary number of candidates, and can thus choose whether to approve or disapprove each candidate. The election is won by the candidate who is approved by the highest number of voters. Approval voting allows voters to be more expressive of their preferences, and it can avoid problems such as vote splitting which are endemic to plurality voting. Together with its elegance and simplicity, this has made approval voting a favorite among voting theorists (Laslier, 2011), and has led to a large research literature (Laslier and Sanver, 2010).

Approval voting combines two ideas: a simple yet expressive ballot format, and an aggregation method for deciding on a winner given the submitted ballots. The two components are often not considered separately: once we have decided to use approval ballots, the aggregation method is usually taken to be obvious. While the standard method (electing

Thm. 1			reinf.[2]	faithfulness[16]	disjoint equality[12]
Thm. 3			reinf.[2]	faithfulness[9,17]	cancellation[12,13]
Thm. 4	anon.[1]	neutr.[3,4,15]	reinf.[2]	non-trivial[10]	strategyproofness[9,5]
Thm. 5		neutr.[14]	reinf.[6]	continuity[7]	avoid Condorcet losers[5]
Thm. 6		neutr.[16]	reinf.[8]	continuity[7]	majority consistency[5]
Thm. 7	anon.[1]		reinf.[2]	faithfulness[9,10]	clone consistency[5]
Thm. 8	anon.[1]	neutr.[3]	reinf.[2]	faithfulness[9,10]	independence of losers[5]
Thm. 9	anon.[1]	neutr.[3]	reinf.[2]		independence of dominated alt.[5]
Thm. 10	anon.[1]		reinf.[8]	reversal symm.[5]	independence of never-approved alt.[11]

Table 1: List of our results, with superscripts indicating the labels of examples in Section 5 showing that the specified axiom cannot be dropped.

the candidate who was approved on the highest number of ballots) is certainly natural, there are many other conceivable ways of counting approval ballots. For example, we could use a type of cumulative voting, where each voter has a unit weight which is split uniformly among the approved alternatives. Or we might impose a maximum on the number of candidates that can be approved by a voter, counting ballots that approve too many candidates as invalid. Or we could declare all Pareto optimal candidates as tied winners.

However, we claim that all alternative aggregation methods fail some of the properties that make approval voting unique, such as its robustness to strategic misrepresentation, its clone-proofness, or its consistent behavior when merging election results of different districts. We provide exhaustive support for this claim, by proving a sequence of axiomatic characterizations. Each row of Table 1 corresponds to a result showing that approval voting is the unique aggregation function satisfying the axioms in the row. Taken together, these results provide axiomatic support for the common intuition that approval voting is the uniquely best way to aggregate approval ballots in single-winner elections.

Our results follow a long line of papers that have axiomatically characterized approval voting, starting with early work of Fishburn (1978, 1979). Many of these characterizations depend on technical axioms that have limited normative appeal. For example, Fishburn (1979) uses “cancellation”, which requires that if every candidate is approved by the same number of voters, then the rule should declare a tie between all candidates. Fishburn (1978) uses “disjoint equality”, which prescribes that if there are only two voters with disjoint sets of approved alternatives, the voting rule should declare a tie between all candidates in the union of their approval sets. In place of these somewhat artificial axioms, our characterizations use properties like clone-proofness, strategyproofness, or avoidance of Condorcet losers, that we consider to have more normative appeal.

Technically, the common basis for all our characterizations is the reinforcement axiom, which requires that the voting rule makes consistent choices across different sub-electorates. Imagine, for example, that a nation is split into several states, and suppose that there exists a candidate who wins in every state (when counting only the ballots cast in that state). Reinforcement requires that, when counting all ballots nation-wide, the voting rule elects exactly those candidates who win in every state individually. This axiom applies to voting rules defined for varying numbers of voters, and thus we operate in a framework with variable electorates. Reinforcement is known to be the driving force in many characterizations of scoring-based rules in social choice theory (see, e.g., Young, 1975; Young and Levenglick, 1978; Myerson, 1995). In contrast to many characterization results based on reinforcement, our proofs use only elementary mathematics. In particular, we do not require separating hyperplane theorems. The appeal of direct and elementary proofs is not merely aesthetic; this property allows our characterizations to be used to *explain* the election outcome to voters: given a specific profile of approval ballots, one can automatically produce a short (polynomial-length) proof showing that the axioms imply that exactly the winners of approval voting need to be elected in the given profile.[1]

We will not prove our characterization results from scratch each time. Instead, we will prove them by reduction to a single base theorem, which characterizes approval voting as the unique rule satisfying reinforcement, faithfulness, and disjoint equality. This base theorem strengthens the result of Fishburn (1978, 1979) by avoiding the use of any symmetry arguments based on neutrality. (The simple proof is very short and may be of interest for teaching purposes.) We then prove the remaining characterizations by showing that a voting rule satisfying the axioms will also satisfy faithfulness and disjoint equality, and hence we are done by invoking the base theorem. We take care to ensure that all of our results are axiomatically tight, and in Section 5 we construct 16 example rules which show that none of the axioms can be dropped or significantly weakened.

Our first result using normative axioms characterizes approval voting using its well-known property of being strategyproof, assuming voters' preferences are dichotomous (i.e., they are indifferent between all approved candidates, and between all disapproved candidates). Our characterization uses a weak version of strategyproofness (following Kelly, 1977). We also identify the strongest notion of strategyproofness that approval voting satisfies, which is significantly stronger than Kelly's. We then turn to majoritarian axioms that require a rule to be consistent with the will of a majority of the voters, and show that approval voting can be characterized either using the fact that it never elects a Condorcet loser, or that it only elects candidates with majority support in cases where more than half the voters submit the same approval ballot. Finally, we characterize approval voting by its resistance to the spoiler effect, familiar from plurality voting, that the presence of a weak candidate can change the winner by 'splitting the vote'. We formalize resistance to the spoiler effect in four different ways – independence axioms and Tideman's (1987) cloning consistency – and show that each characterizes approval voting.

[1] The result of Theorem 4 is explainable in this way when strengthening non-triviality to faithfulness. The results of Theorem 5 and 6 use a simple limit argument, and so would require a stronger logic than in other cases. Previously, Cailloux and Endriss (2016) showed that the Borda rule can be similarly explained in terms of the axioms of Young's (1974) characterization.

In Section 6, we discuss other works on characterizing approval voting. Of particular note is Fishburn's (1979) paper, which shows that neutrality and reinforcement characterizes a class of scoring rules. Fishburn then shows that the only scoring rule satisfying disjoint equality or strategyproofness is approval voting. We obtain these results in a more direct manner, without reasoning about scoring rules. In Appendix A, we state omitted proofs.

2. The Model

Let X be a finite set alternatives. The ballot set $\mathcal{A}$ consists of all non-empty subsets of X, which will be called ballots. A ballot profile P is a function from the ballot set $\mathcal{A}$ to the non-negative integers such that $\sum_{A\in\mathcal{A}} P(A) > 0$. We interpret $P(A)$ as the number of voters whose ballot is A. The approval score $P[a]$ of an alternative a is the number of voters whose ballot includes a, so $P[a] = \sum_{A\in\mathcal{A}: a\in A} P(A)$. Often, it will be useful to identify elements of $\mathcal{A}$ with single-voter ballot profiles. For example, $P + A$ is the profile resulting from P by adding one voter with ballot A; similarly, the profile $P + kA$ is obtained by adding k voters with ballot A to P. For a permutation π on X, the profile $\pi(P)$ has $P(A)$ voters with ballot $\pi(A)$ for every ballot A.

An (approval-based) voting rule f maps each profile P to a set of winning alternatives $f(P) \in \mathcal{A}$. Typically, $f(P)$ will be a singleton, but f may sometimes declare several alternatives to be tied. Our definition of ballot profiles entails that voting rules are anonymous, since they cannot distinguish between voters submitting the same ballot.[2]

We recall a number of axioms for voting rules from the literature. A voting rule satisfies each of the following axioms if the corresponding property holds for all profiles P, P', ballots A, B, and permutations π.

$$f(P) \cap f(P') = f(P + P') \qquad \text{whenever } f(P) \cap f(P') \neq \emptyset \qquad \text{(reinforcement)}$$
$$f(A) = A \qquad \text{(faithfulness)}$$
$$f(\pi(P)) = \pi(f(P)) \qquad \text{(neutrality)}$$
$$f(A + B) = A + B \qquad \text{whenever } A \cap B = \emptyset \qquad \text{(disjoint equality)}$$
$$f(P) = X \qquad \text{whenever } P[a] = P[b] \text{ for all } a, b \in A \qquad \text{(cancellation)}$$

We are interested in the voting rule called *approval voting* (AV), which chooses all alternatives with maximal approval score. It is elementary to check that AV satisfies all of the axioms above. We will also refer to the trivial function $TRIV$ selecting all alternatives in all profiles, and the function $-AV$ selecting all alternatives with minimal approval score. A voting rule is non-trivial if it is not $TRIV$.

[2] One can verify that the characterizations in Sections 3 and 4.2 continue to hold when allowing non-anonymous voting rules. In other results, anonymity is a necessary assumption (see Section 5).

3. Base Theorems

We begin by proving our base theorem, which we will use to obtain all other characterizations in this paper. A related theorem of Fishburn (1978, 1979) uses neutrality instead of faithfulness, and as we show below, Fishburn's result is a direct corollary of our result.

Theorem 1. *AV is the only voting rule satisfying reinforcement, disjoint equality, and faithfulness.*

Proof. Let P be a profile. If some alternative is approved by all voters, i.e., $\bigcap_{A\in\mathcal{A}:P(A)\geqslant 1} A \neq \emptyset$, then faithfulness and reinforcement imply that $f(P) = \bigcap_{A\in\mathcal{A}:P(A)\geqslant 1} A = AV(P)$ and we are done. In this case, we call P a consensus profile.

Now consider the case that P is not a consensus profile. Let $a \in AV(P)$ and $b \in f(P)$. We will show that $a \in f(P)$ and $b \in AV(P)$. If $a = b$ this is obvious, so assume that they are distinct. Let $P[a/b]$ be the number of voters in P who approve a (and possibly other alternatives) but not b, and let $P[b/a]$ be the number of voters in P who approve b (and possibly other alternatives) but not a. Moreover, let $P[\cdot/ab]$ be the number of voters in P who approve neither a nor b. Since P is not a consensus profile, at least one voter has to disapprove a, and so $P[a/b] + P[b/a] + P[\cdot/ab] \geqslant P[b/a] + P[\cdot/ab] > 0$. Let P' be the profile on $P[a/b] + P[b/a] + P[\cdot/ab]$ voters such that

$$P'(\{a\}) = P[b/a], \qquad P'(\{b\}) = P[a/b], \quad \text{and} \quad P'(\{a,b\}) = P[\cdot/ab].$$

In the following, we will consider the profile $P + P'$, and decompose it in two ways.

In the first decomposition, we pair each voter in P (except those approving both a and b) with a voter in P' who approves a disjoint set of candidates:

$$P+P' = \sum_{\substack{A\in\mathcal{A}\\ a\in A, b\notin A}} P(A)\cdot(A+\{b\}) + \sum_{\substack{A\in\mathcal{A}\\ b\in A, a\notin A}} P(A)\cdot(A+\{a\}) + \sum_{\substack{A\in\mathcal{A}\\ a,b\notin A}} P(A)\cdot(A+\{a,b\}) + \sum_{\substack{A\in\mathcal{A}\\ a,b\in A}} P(A)\cdot A.$$

This pairing allows us to apply disjoint equality to each term of the first three sums, and we see that f elects both a and b in each of them. By faithfulness, we obtain the same conclusion for the terms of the fourth sum. Reinforcement implies that $a, b \in f(P + P')$.

In the second decomposition, we pair each $\{a\}$-voter in P' with a $\{b\}$-voter in P'. Since $a \in AV(P)$, we have $P[a/b] \geqslant P[b/a]$, so each $\{a\}$-voter can be matched:

$$P + P' = P + P[b/a]\cdot(\{a\} + \{b\}) + (P[a/b] - P[b/a])\cdot\{b\} + P[\cdot/ab]\cdot\{a, b\}.$$

Considering each term of the sum on the right-hand side separately, we see that f elects b in each of them: $b \in f(P)$ by assumption, $f(\{a\} + \{b\}) = \{a, b\}$ by disjoint equality, and $f(\{b\}) = \{b\}$ and $f(\{a, b\}) = \{a, b\}$ by faithfulness.

If $a \notin f(P)$ then reinforcement applied to the second decomposition implies that $a \notin f(P + P')$, a contradiction to $a, b \in f(P + P')$. If $b \notin AV(P)$, then $P[a] > P[b]$ and thus $P[a/b] - P[b/a] > 0$, so that the third term in the sum does not vanish. Hence by reinforcement $f(P+P') = \{b\}$, again contradicting $a, b \in f(P+P')$. So $a \in AV(P)$ implies $a \in f(P)$, and $b \in f(P)$ implies $b \in AV(P)$. Hence $f(P) = AV(P)$. □

Remark 1. Theorem 1 also holds for the weakening of disjoint equality that only requires $f(A+B) \supseteq A \cup B$ for all disjoint A and B. The proof can be copied almost verbatim. □

Theorem 1 allows us to obtain short proofs of the results of Fishburn (1978, 1979) and Alós-Ferrer (2006).

Theorem 2 (Fishburn, 1978)**.** *AV is the only voting rule satisfying reinforcement, disjoint equality, and neutrality if* $|X| \geqslant 3$.

Proof. We show that reinforcement, disjoint equality, and neutrality imply faithfulness. The result then follows from Theorem 1.

Let A be a ballot. Neutrality implies that $f(A) \in \{X, A, X \setminus A\}$. If $A = X$, the first and the second option coincide and the third option is impossible, since choice sets have to be non-empty. Otherwise, assume for contradiction that $f(A) \in \{X, X \setminus A\}$ and let $b \in f(A) \setminus A$. If $A = \{a\}$ neutrality implies that $f(\{b\})$ is either $X \setminus \{b\}$ or X. Since $|X| \geqslant 3$, $f(\{a\}) \cap f(\{b\}) \neq \emptyset$ and thus reinforcement implies $f(\{a\}) \cap f(\{b\}) = f(\{a\} + \{b\})$, which is either $X \setminus \{a, b\}$ or X, both of which contradict disjoint equality.

For all remaining A, we have $f(A) \cap f(\{b\}) \neq \emptyset$, since $f(\{b\}) = \{b\}$ by the previous case. Then reinforcement implies $\{b\} = f(A) \cap f(\{b\}) = f(A + \{b\})$, which again contradicts disjoint equality. □

Theorem 3 (Alós-Ferrer, 2006; Fishburn, 1979)**.** *AV is the only voting rule satisfying reinforcement, faithfulness, and cancellation.*

Proof. By Theorem 1 and Remark 1, it suffices to show that $f(A + B) \supseteq A \cup B$ for all disjoint ballots A, B. Let $C = X \setminus (A \cup B)$. Cancellation implies that $f(A \cup B + C) = X$ and $f(A + B + C) = X$. So we have

$$\begin{aligned} f(A+B) &= f(A+B) \cap f(A \cup B + C) \\ &= f(A + B + A \cup B + C) \\ &= f(A \cup B) \cap f(A + B + C) = A \cup B, \end{aligned}$$

where the second and third equality follow from reinforcement and the last equality follows from faithfulness. □

Note that cancellation was only used for two and three voters with disjoint ballots in the proof of Theorem 3. Ninjbat (2012) showed that replacing faithfulness by neutrality in Theorem 3 characterizes AV, $-AV$, and $TRIV$.

4. Normative Characterizations

The characterizations in Section 3 rely on axioms such as disjoint equality or cancellation, which we view as technical axioms without obvious normative appeal. The goal of the results in this section is to replace disjoint equality and cancellation by axioms that have clear normative content. Concretely, we will consider strategyproofness, majoritarian properties (such as never electing Condorcet losers), and properties that ensure consistency across different agendas (such as independence of unchosen alternatives). The base theorems stated in Section 3 are useful throughout, since we will prove the following characterizations by reduction to these basic results.

4.1. Strategyproofness

Strategyproofness rules out that a voter can obtain a more preferred outcome by stating an untruthful ballot. Given the limited action space of the voters, we cannot infer much about their preferences over possible choice sets based on their ballots and will thus be conservative in our assumptions about when a voter prefers a choice set to another. Conservative assumptions make the strategyproofness axiom weaker, and the characterization result stronger.

Let P be a profile, where we allow $P = 0$. We say that a voter with truthful ballot A can *manipulate* in the profile $P + A$ by reporting a ballot B if either

(i) $f(P + A) \cap A = \emptyset$ but $f(P + B) \cap A \neq \emptyset$; or

(ii) $f(P + A) \not\subseteq A$ but $f(P + B) \subseteq A$.

Thus, in case (i), the manipulator has some approved alternatives in the winning set rather than only disapproved alternatives. In case (ii), the manipulator ensures that all winning alternatives are approved instead of having some disapproved alternatives in the winning set. We say that f is *strategyproof* if no manipulation is possible in any profile.

This definition of strategyproofness coincides with that proposed by Kelly (1977) when ballots are viewed as *dichotomous preferences*: a voter is indifferent between all approved alternatives, prefers every approved alternative to every disapproved alternative, and is indifferent between all disapproved alternatives. Kelly's rationale for this definition of strategyproofness is that, if violated, some voter can "clearly manipulate", because no matter which tie-breaking mechanism is invoked to select a final outcome from choice sets, manipulating is always at least as good as truth-telling and strictly better for some tie-breaking mechanism. Another interpretation goes as follows: suppose ties are broken according to some lottery with positive probability for each chosen alternative, and this lottery is unknown to the voters. If a voter can manipulate, they can thereby either avoid certainty of receiving a disapproved alternative or guarantee an approved alternative. Brandt et al. (2018a) elaborate on this interpretation in more detail.

It is not hard to see that AV is strategyproof in this sense (see Proposition 1 below). In contrast, with non-dichotomous preferences, Kelly strategyproofness is incompatible with Condorcet-consistency (cf. Brandt, 2015). We can characterize AV using strategyproofness. Based on his discussion of scoring rules, Fishburn (1979, Theorem 10) obtains a similar characterization using a stronger notion of strategyproofness, though his proof does not require the extra strength.

Theorem 4. *AV is the only non-trivial voting rule satisfying reinforcement, neutrality, and strategyproofness.*

Proof. We prove that any such f satisfies faithfulness and disjoint equality. Theorem 1 then implies $f = AV$.

First we show faithfulness. Neutrality implies that $f(A) \in \{X, A, X \setminus A\}$ for all ballots A. If $f(A) = A$ for all A, then f satisfies faithfulness and there is nothing left to show. If $f(A) = X$ for all A, reinforcement implies that $f(P) = X$ for all profiles P, i.e., $f = TRIV$, which is contrary to the assumption that f is non-trivial. If there is a ballot A such that $f(A) = X \setminus A$, then the voter in the single voter profile A has a case (i) manipulation by reporting X, since $f(X) = X$. In the remaining case, there are ballots A, B such that $f(A) = A$ and $f(B) = X$. If $|B| > |A|$, let $B' \subseteq B$ be a ballot with $|B'| = |A|$. By neutrality, $f(A) = A$ implies $f(B') = B'$. So the voter in the single voter profile B has a case (ii) manipulation by reporting B' instead of B, which contradicts strategyproofness. Thus, by neutrality, there is $k \in \{2, \ldots, m-1\}$ such that $f(A) = A$ for all A with $|A| \geqslant k$ and $f(A) = X$ for all A with $|A| \leqslant k-1$. Let A be a ballot such that $|A| = k-1 \leqslant m-2$ and $a, b \in X \setminus A$ be two distinct alternatives. Then, $f(A) = X$, $f(A \cup \{a\}) = A \cup \{a\}$, and $f(A \cup \{b\}) = A \cup \{b\}$. Hence, by reinforcement, $f(A \cup \{a\} + A) = A \cup \{a\}$ and $f(A \cup \{a\} + A \cup \{b\}) = A$. Thus, in the profile $A \cup \{a\} + A$, a voter with approval set A has a case (ii) manipulation by reporting $A \cup \{b\}$.

Second we show disjoint equality. To this end, let A, B be two disjoint ballots. We first show that $f(A + B) \subseteq A \cup B$. Assume for contradiction that this is not the case, i.e., $f(A + B) \setminus (A \cup B) \neq \emptyset$. Let $b \in B$, $c \in f(A + B) \setminus (A \cup B)$, and $C = B \setminus \{b\} \cup \{c\}$. By faithfulness, we have $f(C) = C$. Then, reinforcement implies that $f(A + B + C) = f(A + B) \cap C$. Hence, $b \notin f(A + B + C)$ and $c \in f(A + B + C)$. Since $|B| = |C|$ and $A \cap B = A \cap C = \emptyset$, this contradicts neutrality.

So neutrality implies $f(A+B) \in \{A \cup B, A, B\}$. Assume for contradiction that $f(A+B) = A$. (The case $f(A + B) = B$ is analogous.) Let $a \in A$ and $b \in B$. Neutrality and the fact that $f(A+\{b\}) \subseteq A \cup \{b\}$ imply that $f(A+\{b\}) \in \{A \cup \{b\}, A, \{b\})$. Since $f(A+B) \cap B = \emptyset$, strategyproofness implies $f(A + \{b\}) \cap B = \emptyset$, as otherwise the voter with ballot B has a case (i) manipulation in the profile $A + B$. Hence, $f(A + \{b\}) = A$. Faithfulness implies that $f(\{a, b\}) = \{a, b\}$. Thus, by reinforcement, $f(A + \{b\} + \{a, b\}) = A \cap \{a, b\} = \{a\}$. Then strategyproofness implies that $f(\{a\} + \{b\} + \{a, b\}) = \{a\}$, as otherwise a voter with ballot $\{a\}$ has a case (ii) manipulation in the profile $\{a\} + \{b\} + \{a, b\}$. This contradicts neutrality.

In summary, we have that f satisfies reinforcement, disjoint equality, and faithfulness, and so $f = AV$ by Theorem 1. □

Our definition of strategyproofness only considers unilateral deviations. If we strengthen it to allow for deviations by groups of voters, AV turns out to be manipulable. For example, in the profile $P = \{a\} + \{b\} + 2\{c\}$, AV chooses $\{c\}$. If the voters with ballots $\{a\}$ and $\{b\}$ report $\{a, b\}$ instead, we obtain the profile $P' = 2\{a, b\} + 2\{c\}$ for which AV returns $\{a, b, c\}$.

This constitutes a successful case (i) manipulation for those two voters. Examples involving case (ii) manipulations can be constructed likewise. Thus, from Theorem 4, we see that no non-trivial voting rule satisfies reinforcement, neutrality, and group strategyproofness.[3]

We used a weak definition of strategyproofness in Theorem 4, requiring that lying can never make the output unambiguously better. For example, this notion does not rule out that lying might add approved alternatives to the output while removing others. It turns out that AV satisfies a significantly stronger version, requiring that every lie makes the output weakly worse. We say that a voting rule satisfies *strong strategyproofness* if for all profiles P (allowing $P = 0$) and all ballots A and B, either

(i) $f(P + A) \subseteq A$, or

(ii) $f(P + B) \cap A = \emptyset$, or

(iii) $f(P + A) \setminus f(P + B) \subseteq A$ and $(f(P + B) \setminus f(P + A)) \cap A = \emptyset$.

Thus, either truth-telling yields only approved alternatives (case (i)), lying yields only disapproved alternatives (case (ii)), or, if neither of those holds, lying can only remove approved alternatives and add disapproved alternatives.

Strong strategyproofness is stronger in that it requires truth-telling to be a dominant, rather than just undominated, strategy. It is also stronger in that it is based on a finer method of comparing sets of tied outcomes, which Fishburn (1972) motivated as arising as the preferences of an expected utility maximizing voter when ties between alternatives are broken by an even-chance lottery. For non-dichotomous preferences, such a strong notion is impossible to satisfy.[4] However, using dichotomous preferences, AV satisfies it. This can be deduced from the theory developed by Brams and Fishburn (1978, Theorem 4); here, we give a direct proof.

Proposition 1. *AV satisfies strong strategyproofness.*

Proof. Consider a profile P and two ballots A, B. Assume that neither (i) nor (ii) holds, i.e., $AV(P + A) \not\subseteq A$ and $AV(P + B) \cap A \neq \emptyset$. For $a \in AV(P + A) \cap (X \setminus A)$ and $b \in AV(P + B) \cap A$, we have

$$(P + A)[a] \leqslant (P + B)[a] \leqslant (P + B)[b] \leqslant (P + A)[b] \leqslant (P + A)[a].$$

Hence, $(P+A)[a] = (P+B)[b]$ for all $a \in AV(P+A)$ and $b \in AV(P+B)$, meaning that the approval score of approval winners is the same in both profiles. So $AV(P + A) \cap (X \setminus A) \subseteq AV(P + B)$ and $AV(P + B) \cap A \subseteq AV(P + A)$, which is equivalent to (iii). □

[3] When dropping reinforcement, Brandt et al. (2019, Remark 2) show that the voting rule returning all Pareto undominated alternatives is group strategyproof.

[4] For example, it is incompatible with anonymity and Pareto optimality (Brandt et al., 2018a).

4.2. Majoritarian Properties

In many democratic contexts, an important goal of voting is to uncover the will of a majority. If the preferences of a majority of voters share a certain feature, then this should be reflected in the collective decision. A classic example of this desideratum is the Condorcet criterion of choosing a Condorcet winner whenever it exists.

In this section, we assume that voter preferences are dichotomous and coincide with the approval sets reported by the voters. In this case, the majority relation is transitive (Inada, 1969) and coincides with the relation obtained by ordering alternatives according to their approval score. The maximal elements of the majority relation are the approval winners and preferred to every other alternative by a weak majority of voters (so they are weak Condorcet winners). If there is a unique approval winner, it is a Condorcet winner. Thus, AV is Condorcet consistent. So AV is characterized by Condorcet consistency except for the case of ties at the top of the majority relation. From the proof of Theorem 6 of Fishburn (1979), it follows that AV is characterized by neutrality, reinforcement, continuity, and Condorcet consistency. In the context of voting with linear orders rather than approval ballots, these axioms are incompatible (Young and Levenglick, 1978, Theorem 2).

In a profile P, an alternative a is a *Condorcet loser* if for every other alternative b, more voters prefer b to a than prefer a to b. A voting rule *avoids Condorcet losers* if it never chooses a Condorcet loser: $a \notin f(P)$. This is weaker than Condorcet consistency, and hence AV satisfies it. Some refinements of AV satisfy it as well (see Example 3 in Section 5), and thus AV is not characterized by reinforcement, neutrality, and avoidance of Condorcet losers. We can pin down AV uniquely by adding the *continuity axiom*, which requires that if a is chosen uniquely in a profile P but not in P', then a is again chosen uniquely once we add enough copies of P to P'. Formally, f is *continuous* if for any two profiles P, P' with $f(P) = \{a\}$, there is an integer k such that $f(P' + kP) = \{a\}$.[5] Myerson (1995) calls this property the *overwhelming majority* axiom. Together, the four axioms characterize AV. In the context of voting with linear orders, the same axioms characterize Borda's rule (combining results of Smith, 1973 and Young, 1975).

Theorem 5. *AV is the only voting rule that satisfies reinforcement, neutrality, continuity, and avoids Condorcet losers.*

Proof. We show that any such voting rule satisfies cancellation and then apply Theorem 3. Assume for contradiction that f does not satisfy cancellation. So there is a profile P such that $P[a] = P[b]$ for all alternatives a, b and $f(P) \neq X$. Let $a \in f(P)$. We now construct a profile P_a such that all alternatives have the same approval score in P_a and $f(P_a) = \{a\}$. To this end, let $\Pi_a = \{\pi \in \Pi(X) \colon \pi(a) = a\}$ be the set of all permutations on X that hold a fixed and $P_a = \sum_{\pi \in \Pi_a} \pi(P)$ the profile obtained by summing up all permutations of P

[5] In the presence of reinforcement, this definition is equivalent to requiring that $a \in f(P' + kP)$ for some k. This is because if $a \in f(P' + kP)$, then $f(P' + (k+1)P) = f(P' + kP) \cap f(P) = \{a\}$ by reinforcement.

for permutations in Π_a. By neutrality, $a \in f(\pi(P))$ for all π and so

$$f(P_a) = \bigcap_{\pi \in \Pi_a} f(\pi(P)) = \bigcap_{\pi \in \Pi_a} \pi(f(P)) = \{a\},$$

where the first equality follows from repeated application of reinforcement, the second equality from neutrality of f, and the third equality from the assumption that $f(P) \neq X$. Then, continuity implies that there is $k \in \mathbb{N}$ such that $f(X \setminus \{a\} + kP_a) = f(P_a) = \{a\}$. However, a is a Condorcet loser in the profile $X \setminus \{a\} + kP_a$, which contradicts the assumption that f avoids Condorcet losers. □

Majority consistency requires that if more than half of the voters in a profile P submit the same ballot A, then at least one alternative from A is a winner (possibly among other alternatives not from A). Notice that every alternative not in A is disapproved by a majority of voters. In such situations, AV will return a subset of A. For voting with linear orders, similar axioms characterize the plurality scoring rule (Lepelley, 1992, see also Sanver, 2002).

Theorem 6. *AV is the only non-trivial voting rule satisfying reinforcement, neutrality, continuity, and majority consistency.*

4.3. Consistency Across Variable Agendas

Reinforcement imposes consistent behavior across situations where the set of voters varies. Similarly, we might want to let the set of alternatives (the *agenda*) vary, and impose consistency conditions for varying agendas. Such conditions can ensure that collective choices are rationalizable, or they can prevent manipulations by someone with agenda-setting power, who could introduce inferior alternatives or copies of already existing alternatives.

To formalize such properties, we require a more general set-up. We re-interpret X as the set of all potential alternatives. The set of *agendas* $\mathcal{P}(X)$ consists of all non-empty subsets of X. Given an agenda $Y \in \mathcal{P}(X)$, the corresponding ballot set $\mathcal{A}_Y$ contains all non-empty subsets of Y. As before, a profile P on an agenda Y is a function from $\mathcal{A}_Y$ to the non-negative integers such that $\sum_{A \in \mathcal{A}_Y} P(A) > 0$. The restriction of a profile P on agenda Y to agenda $Z \subseteq Y$, written P_Z, is defined by $P_Z(B) = \sum_{A \in \mathcal{A}_Y :\, A \cap Z = B} P(A)$ for all $Z \in \mathcal{A}_Z$. A voting rule f_Y *on an agenda* Y maps a profile on Y to a subset of Y. A voting rule $f = (f_Y)_{Y \in \mathcal{P}(X)}$ specifies a voting rule for each agenda. All axioms defined earlier hold if they hold within each agenda. We will sometimes abuse notation by applying f_Z to a profile P on a larger agenda $Y \supseteq Z$. In such cases, P is to be understood as its restriction P_Z to Z.

In the context of linear orders, Tideman (1987) noticed that under many common voting rules, the winner can change when some candidate is 'cloned' by introducing new candidates that each voter ranks in an adjacent position to the original candidate. Few voting rules can avoid this behavior. In the context of approval ballots, AV does avoid it. To define an appropriate version of Tideman's axiom for approvals, we say that two alternatives a, b are *clones* of each other in a profile P if every voter either approves both a and b or disapproves both. A voting rule f satisfies *clone consistency* if, whenever a and b are clones in P, then $f_{Y\setminus\{b\}}(P) = f_Y(P) \cap (Y \setminus \{b\})$ and $b \in f_Y(P)$ if and only if $a \in f_{Y\setminus\{b\}}(P)$. So adding a clone b of an alternative a to a profile has no effect on whether other alternatives are chosen or not, and b is chosen if and only if a was chosen in the original profile.

We show that reinforcement, clone consistency, and faithfulness characterize AV. In fact, using a more technical argument (omitted here), we can show that the only rules that satisfy reinforcement and clone consistency are AV, $-AV$, and $TRIV$. Thus, faithfulness is only required to rule out $-AV$ and $TRIV$.

Theorem 7. *AV is the only voting rule satisfying reinforcement, clone consistency, and faithfulness if $|X| \geqslant 4$.*

Another choice-theoretic axiom for variable agendas prescribes that removing unchosen alternatives from the agenda should not change the choice set (see, e.g., Chernoff, 1954; Aizerman and Aleskerov, 1995; Brandt and Harrenstein, 2011). This axiom prevents a losing alternative from spoiling the election. Formally, a voting rule f satisfies *independence of losers* if $f_Y(P) = f_Z(P)$ for all profiles P and agendas $Z \subseteq Y$ with $f_Y(P) \subseteq Z$. Because removing losers does not change the approval scores of other candidates, AV satisfies this property, and can be characterized by reinforcement, neutrality, faithfulness, and independence of losers. When dropping faithfulness, at the expense of a more technical proof, one can show that AV, $-AV$, and $TRIV$ are the only rules that satisfy these axioms.

Theorem 8. *AV is the only voting rule satisfying reinforcement, neutrality, faithfulness, and independence of losers.*

Proof. Take any such voting rule f and some agenda $Y \in \mathcal{P}(X)$. We will omit the subscript Y for f within this proof. We show that f satisfies disjoint equality and apply Theorem 1.

Let A, B be two disjoint ballots and $a \in A$ and $b \in B$. Neutrality implies that $f(A+B) \in \{X, A \cup B, A, B, X \setminus B, X \setminus A\}$. First assume for contradiction that $f(A+B) \not\subseteq A \cup B$ and let $c \in f(A+B) \setminus (A \cup B)$. By faithfulness, we have $f(\{c\}) = \{c\}$. Hence, by reinforcement, $f(A+B+\{c\}) = \{c\}$. Independence of losers implies that $f(A+B+\{c\}) = f_{\{a,b,c\}}(\{a\} + \{b\} + \{c\}) = \{c\}$, which contradicts neutrality. So we have $f(A+B) \in \{A \cup B, A, B\}$.

Second, consider the case $f(A+B) = A$. From the previous case and neutrality, we know that $f(\{a\} + \{b\}) = \{a, b\}$. Hence, by reinforcement, $f(A+B+\{a\}+\{b\}) = \{a\}$. Independence of losers implies that $f(A+B+\{a\}+\{b\}) = f_{\{a,b\}}(\{a\}+\{b\}+\{a\}+\{b\}) = \{a\}$, which again contradicts neutrality. Similarly, we get a contradiction if $f(A+B) = B$. Hence, $f(A+B) = A \cup B$ is the only possibility, and thus f satisfies disjoint equality. □

To decide whether removal of an alternative is allowed to change the outcome, independence of losers references the choice set of the voting rule under consideration. Alternatively, we can look for a more objective approach to identify inferior alternatives whose removal should not change the election outcome. We will consider two such notions. The first requires that the removal of a Pareto dominated alternative does not change the set of winners. The second only requires that removing an alternative which is not approved by any voter does not change the outcome.

An alternative y is Pareto dominated in a profile P if there exists an alternative x such that every voter in P who approves y also approves x, and there is a voter in P who approves x but not y. A voting rule satisfies *independence of Pareto dominated alternatives* if $f_Z(P) = f_Y(P)$ for all agendas Y, Z and profiles P on Y such that $Z \subseteq Y$ and all alternatives in $Y \setminus Z$ are Pareto dominated in P. In conjunction with reinforcement and neutrality this characterizes AV. In the context of linear orders, these axioms characterize the plurality scoring rule (Richelson, 1978; Ching, 1996; Öztürk, 2017).

Theorem 9. *AV is the only voting rule satisfying reinforcement, neutrality, and independence of Pareto dominated alternatives.*

A weakening of the independence axiom of Theorem 9 is *independence of never-approved alternatives*, which prescribes that $f_Y(P) = f_Z(P)$ for all agendas Y, Z and profiles P on Y such that $Z \subseteq Y$ and $(Y \setminus Z) \cap A = \emptyset$ for all ballots A with $P(A) > 0$. As example Example 5 in Section 5 shows, AV is not the only rule satisfying this weaker independence condition together with the other axioms. However, AV is characterized when adding an axiom often called reversal symmetry: if all voters switch to approving the complement of their ballot, then all chosen alternatives should become unchosen alternatives (unless all alternatives were chosen in the original profile). For a ballot $A \neq Y$, we denote by A^c the complement of A in Y, i.e., $A^c = Y \setminus A$; for $A = Y$, $A^c = Y$. Similarly, P^c is the profile where $P^c(A^c) = P(A)$ for all ballots $A \in \mathcal{A}_Y$. Then *reversal symmetry* requires that $f(P) \cap f(P^c) = \emptyset$ whenever $f(P) \neq X$. In the context of linear orders, these axioms characterize Borda's scoring rule (Morkelyunas, 1982, see also Saari and Barney, 2003).

Theorem 10. *AV is the only voting rule satisfying reinforcement, reversal symmetry, and independence of never-approved alternatives.*

Proof. Fix some agenda Y. We show that f_Y satisfies faithfulness and disjoint equality and then invoke Theorem 1 to conclude $f_Y = AV_Y$.

For faithfulness, fix some ballot A. Then independence of never approved alternatives implies $f(A) = f_A(A) = A$. For agenda A, $A = A^c$, and so $f_A(A) = f_A(A^c)$. Then reversal symmetry implies $f_A(A) = A$, which shows faithfulness.

For disjoint equality, let A, B be two disjoint ballots and $P = A + B$. Independence of never-approved alternatives implies $f(P) = f_{A\cup B}(P)$. For agenda $A \cup B$, $P = P^c$, and so $f_{A\cup B}(P) = f_{A\cup B}(P^c)$. Then reversal symmetry implies $f_{A\cup B}(P) = A \cup B$, which shows disjoint equality. □

5. Independence of Axioms

In each of our characterizations, the axioms used are independent: when dropping any one of the axioms, other rules satisfy the remaining axioms. To show this, we give a lengthy list of example rules below that satisfy particular combinations of axioms. Most of these examples are technical in nature, and not in themselves interesting. To see which example can be used to prove independence in a specific result, refer to Table 1.

Example 1. The rule that is like AV but counts voter 1 double.

This rule is not anonymous, but it satisfies neutrality, reinforcement, faithfulness, strategyproofness, clone consistency, reversal symmetry, independence of losers, of dominated alternatives, and of never-approved alternatives. That the rule satisfies these axioms is easily deduced from the fact that AV satisfies them.

Example 2. The rule PO selecting all Pareto optimal alternatives.

This rule fails reinforcement, since $\mathrm{PO}(a+c) = \{a, c\}$ and $\mathrm{PO}(b+c) = \{b, c\}$, but $\mathrm{PO}(a+b+c+c) = \{a, b, c\}$. However, PO satisfies anonymity, neutrality, disjoint equality, cancellation, strategyproofness (Brandt et al., 2018b, Thm. 3.1), clone consistency, independence of losers, and independence of dominated alternatives.

Example 3. The rule AV_{lex} selecting the lexicographically first approval winner.

This rule fails neutrality, but it satisfies anonymity, reinforcement (if a is the lexicographically first approval winner in both P and P', then $AV(P)$ and $AV(P')$ intersect, and so $AV(P+P') = AV(P) \cap AV(P')$ because AV satisfies reinforcement, and a is the lexicographically first element of the latter set), strategyproofness (since AV satisfies strong strategyproofness and lexicographic tie-braking is consistent across choice sets), independence of losers and of dominated alternatives (since AV satisfies these properties).

Example 4. The constant rule.

This rule fails neutrality, but satisfies anonymity, reinforcement, and strategyproofness.

Example 5. The plurality rule ignoring all non-singleton ballots.

This rule fails strategyproofness since $f(\{a\} + \{b, c\}) = \{a\}$ where the second voter can manipulate to $f(\{a\} + \{b\}) = \{a, b\}$. It fails majority consistency and fails to avoid Condorcet losers since $f(\{a\} + \{b, c\} + \{b, c\}) = \{a\}$. It fails clone consistency since $f(\{a\} + \{b\}) = \{a, b\}$ but $f(\{a\} + \{b, c\}) = \{a\}$. It fails independence of losers and of dominated alternatives, since $f(\{a\} + \{a, b\} + \{c\}) = \{a, c\}$ but $f(\{a\} + \{a\} + \{c\}) = \{a\}$. It fails reversal symmetry since $f(\{a\} + \{b, c\}) = \{a\}$ but $f(\{b, c\} + \{a\}) = \{a\}$. However, it is anonymous, neutral, and satisfies reinforcement and continuity because it is a scoring rule. It also satisfies independence of never-approved alternatives, since if we delete a never-approved alternative, the set of voters with a singleton ballot does not change.

Example 6. The rule CNL selecting all alternatives that are not Condorcet losers.

This rule fails reinforcement, since $f(\{a\} + \{b\}) = \{a, b\}$ and $f(\{a\} + \{c\}) = \{a, c\}$ but $f(\{a\} + \{a\} + \{b\} + \{c\}) = \{a, b, c\}$. The rule is neutral, it avoids Condorcet losers by definition, and it satisfies continuity (since if a is the Condorcet loser in profile P, then a is the Condorcet loser in $P' + kP$ for sufficiently large k, so $a \notin f(P' + kP)$).

Example 7. The rule selecting the approval winners with highest plurality score.

This rule fails continuity, since if $P = \{b\}$ and $P' = \{a\} + \{a, b\} + \{b, c\}$, then $f(P) = \{a\}$, but $f(P + kP') = \{b\}$ for all k. The rule is anonymous, neutral, it satisfies reinforcement (since it is a composite scoring rule), and it avoids Condorcet losers and is majority consistent (since it returns a subset of the approval winners).

Example 8. The rule selecting the approval winners, considering only those ballots that occur most frequently in the profile.

This rule fails reinforcement since $f(\{a, b\} + \{a, c\}) = \{a\}$ and $f(\{a, b\}) = \{a, b\}$, but $f(\{a, b\}, \{a, b\}, \{a, c\}) = \{a, b\}$. It is anonymous and neutral. It satisfies majority consistency, since a ballot that is reported by a majority is a (uniquely) most-frequent ballot. It satisfies continuity since for any profiles P and P', we have for large enough k that $f(P' + kP) = f(P)$, since the most-frequent ballots in P become the most-frequent ballots in $P' + kP$ for large enough k. It satisfies independence of never-approved alternatives, since removing a never-approved alternative does not change any approval sets, and does not change which ballots are most-frequent.

Example 9. The rule $-AV$ returning the alternatives with lowest approval score.

Example 10. The rule $TRIV$ returning all alternatives.

Example 11. The rule like AV, but counting plurality and veto ballots double.

This rule fails independence of never-approved alternatives, since $f_{\{a,b,c,d\}}(\{a, b\}, \{c\}) = \{c\}$ but $f_{\{a,b,c\}}(\{a, b\}, \{c\}) = \{a, b, c\}$. It is anonymous, satisfies reinforcement (since it is a scoring rule), and satisfies reversal symmetry (since AV satisfies reversal symmetry, and reversing a profile preserves voter weights).

Example 12. The rule like AV, but counting veto ballots double.

This rule fails disjoint equality and cancellation in 2-voter profile, since $f(\{a\} + X \setminus \{a\}) = X \setminus \{a\}$. It satisfies reinforcement (being a scoring rule), and faithfulness. It also satisfies cancellation in 3-voter profiles since if A, B, C are disjoint ballots, then none of A, B, C can be a veto ballot (we disallow empty ballots here).

Example 13. The scoring rule where a ballot A gives $|A|(|X|-|A|)$ points to each approved alternative.

This rule fails cancellation on 3-voter profiles because for $|X| \geqslant 4$, we have $f(\{a\} + \{b\} + X \setminus \{a, b\}) = X \setminus \{a, b\}$ because a and b have $|X| - 1$ points each and other alternatives have $(|X| - 2) \cdot 2 > |X| - 1$ points each. However, this rule satisfies reinforcement (since it is a scoring rule), faithfulness, and cancellation on 2-voter profiles (since if $A \cup B = X$ and $A \cap B = \emptyset$, then $|A|(|X| - |A|) = |B|(|X| - |B|)$, and so A and B give the same number of points to each approved alternative, and hence $f(A + B) = X$).

Example 14. The rule that is like AV, except that the ballot $\{a\}$ gives 1 point to a (as usual), and 0.5 points to b and c each.

This rule fails neutrality. Assume $|X| \geqslant 4$. The rule satisfies reinforcement and continuity (since it is a scoring rule). The rule also avoids Condorcet losers. Note that each alternative's score under f is at least that alternative's approval score, and they are equal except possibly for b and c. Let P be a profile with a Condorcet loser ℓ, where $\ell \neq b$ and $\ell \neq c$. Then ℓ's score is ℓ's approval score. Take $d \in X \setminus \{\ell, b, c\}$. Then d's score is d's approval score, so d's score is strictly higher than ℓ's score, since ℓ is the unique Condorcet loser. Hence $\ell \notin f(P)$. Let P be a profile where b is the Condorcet loser (in particular, c majority-beats b), yet $b \in f(P)$. Delete all ballots $\{a\}$ from P to obtain P'. In P', c still majority-beats b. Since f behaves like AV on P', c has strictly more points than b in P'. But since the $\{a\}$-ballots give the same number of points to b and c, c must have had strictly more points than b in P, contradicting that $b \in f(P)$; so $b \notin f(P)$. Symmetrically, we have $c \notin f(P)$ whenever c is the Condorcet loser in P.

Example 15. The rule where each voter gives $1 + \epsilon$ points to approved alternatives, and 1 point to a if a is not approved (where ϵ is infinitesimally small).

An alternative description of this rule is: if the ballots of all voters intersect (a consensus profile), then return that intersection. Otherwise, return $\{a\}$. This rule is not neutral, but it satisfies faithfulness. Being a composite scoring rule, it satisfies reinforcement. It is also strategyproof: in a consensus profile, for each voter, a subset of approved alternatives is selected, so there are no manipulations. If $\{a\}$ is returned, then a manipulator can only change this by approving alternatives that all other voters approve, but this is not a successful manipulation.

Example 16. The rule that is like AV, except that the ballot approving a gives 2 points to a and 1 point to every other approved alternative, and ballots disapproving a give -1 points to a and 1 point to each approved alternative.

This rule fails faithfulness since $f(\{a, b\}) = \{a\}$, and it fails neutrality by the same example. The rule is anonymous, satisfies reinforcement and continuity (since it is a scoring rule), disjoint equality (since for disjoint A and B, either neither approves a so $f(A+B) = A \cup B$, or exactly one of them approves a, whence a gets $2 + -1 = 1$ points, and every other alternative in $A \cup B$ gets 1 point as well, so $f(A + B) = A \cup B$), and majority consistency (suppose in profile P, ballot A is reported by more than half of the voters. If $a \notin A$, then each alternative in A has score at least $P(A)$ and all alternatives not in A have a lower score. If $a \in A$, then each alternative in A other than a has score at least $P(A)$, alternative a has score $2P(A) - \sum_{B \in \mathcal{A} \colon a \notin B} P(B) > P(A)$ because a minority disapproves a, and alternatives outside A have score less than $P(A)$. In either case $f(P) \subseteq A$).

Example 17. The rule where each ballot assigns $+1$ points to approved alternatives and -1 points to disapproved alternatives, except that ballots approving a but disapproving b give $+2$ points to a and -2 points to b, and that ballots approving b but disapproving a give $+2$ points to b and -2 points to a.

This rule fails faithfulness and fails neutrality, since $f(\{a, c\}) = \{a\}$. Since it is a scoring rule, it satisfies anonymity, reinforcement, and continuity. It satisfies cancellation, since in a profile in which all approval scores are equal, we have $P[a/b] = P[b/a]$ (in the notation of the proof of Theorem 1), and so under this rule all alternatives get 0 points in total.

6. Related Characterizations of Approval Voting

The literature on axiomatic characterizations of approval voting and variants thereof goes well beyond what we have discussed above.

Using variants of the axioms in Theorem 1, Sertel (1988) shows that AV is the only voting rule that satisfies anonymity, weak unanimity (faithfulness), weak consistency (reinforcement where one of the profiles is a single-voter profile), and strong disjoint equality. The latter property requires that if a ballot A contains none of the winners for a profile P, then the winners for the profile $P + A$ are the winners for P and alternatives in A whose approval score in P is one less than the maximal approval score.

Fishburn (1979) uses neutrality, continuity, and reinforcement to characterize the class of *scoring rules* on the domain of approval ballots. A scoring rule is specified by a vector $(s_1, \dots, s_m) \in \mathbb{R}^m$ assigning a score to each ballot cardinality. Given a profile P, the score of $x \in X$ is $\sum_{A \in \mathcal{A}\colon x \in A} s_{|A|} \cdot P(A)$. The scoring rule returns the set of alternatives with the highest score. For example, AV is the scoring rule $(1, \dots, 1)$; cumulative voting is the scoring rule $(1, \frac{1}{2}, \frac{1}{3}, \dots, \frac{1}{m})$; and plurality voting is the scoring rule $(1, 0, \dots, 0)$. Without the continuity axiom, we obtain the class of composite scoring rules, where each score s_i is specified by a vector in $\mathbb{R}^m$ (rather than a number in $\mathbb{R}$), and we compare vectors in $\mathbb{R}^m$ lexicographically. An example of such a rule is approval voting, with ties broken in favor of candidates who have the highest plurality score. Fishburn's result is closely related to Young's (1975) characterization of scoring rules when ballots are linear orders over alternatives, and the characterizations of scoring rules by Myerson (1995) and Pivato (2013) in an abstract setting. Theorem 10 of Fishburn (1979) shows that AV is the only rule for which truthfully reporting ones ballot is the only undominated strategy, where dominance is defined similarly as in Section 4.1. A related result by Vorsatz (2008) shows that AV is the only strategyproof scoring rule on approval ballots for a notion of strategyproofness that is slightly weaker than ours. Alcalde-Unzu and Vorsatz (2009) characterize the class of *size* AV rules, which are scoring rules where the score of a ballot weakly decreases in the number of approved alternatives. In addition to anonymity, neutrality, reinforcement, and a continuity axiom, they assume congruity (adding a voter who disapproves a losing alternative does not make it a winner) and contraction (removing alternatives from a voter's ballot does not add new winners unless all winners are removed).

Vorsatz (2007) considers voting rules for variable sets of voters and alternatives. The choices for variable sets of alternatives have to be rationalizable by a transitive relation (cf. Sen, 1977) and indifferent voters do not influence the choice from two-alternatives sets. He shows that AV is the only such voting rule that is anonymous, neutral, strategyproof (as in Vorsatz (2008)), and strictly monotonic, where the latter requires that ties on two alternative sets are resolved if one alternative gains additional support. In the same framework, Massó and Vorsatz (2008) characterize *weighted* AV, where the approval score of an alternative is multiplied by its weight, which is exogenously given and fixed across ballot profiles. Their result requires anonymity, reinforcement, and two properties that put bounds on the ratios of the weights.

Baigent and Xu (1991) consider rules that aggregate choice functions into a collective choice function, i.e., a function that specifies the collective choice from each subset of alternatives. Here the approval score of an alternative corresponds to the number of voters who choose it. If the aggregation rule yields collective choice functions satisfying neutrality, positive responsiveness (additional support for a chosen alternative makes it the unique choice), and independence of symmetric substitutions (the collective choice only depends on the vector of approval scores), then it has to be approval voting.

7. Discussion

We have provided seven characterizations of approval voting based on the reinforcement axiom in combination with other appealing properties. These results are strong arguments for using approval voting once we accept the premise that voters submit approval ballots. Crucially, all our results hinge on reinforcement. In the context of political elections, making consistent choices across sub-electorates prevents severe instances of gerrymandering and is thus particularly important.

On a more technical note, many reinforcement-based results in social choice theory reason over an unbounded number of voters, e.g., when employing convex separation theorems on the set of fractional preference profiles. By contrast, some of our proofs allow us to derive bounds on the number of voters. For example, if a voting rule satisfies all assumptions of Theorem 1 on a set of n voters and all its subsets, then it has to be approval voting on subsets of up to $n/2$ voters. The proofs of Theorems 5 and 6 do not yield such bounds, since continuity requires arbitrarily large electorates.

The examples in Section 5 show that no axiom can be dropped from any of our characterizations. For Theorems 8 and 9, we use the voting rule choosing the lexicographically first approval winner, AV_{lex}, as an example that neutrality cannot be dropped. However, AV_{lex} is very similar to AV, since it is a refinement. We leave as an open question whether there is a voting rule satisfying all axioms of Theorems 8 or 9 except for neutrality, and that is not a subset of approval voting.

Lastly, in this paper we have considered the case of choosing a set of winning alternatives. However, approval-based rules are also natural candidates for choosing rankings, committees, or lotteries over alternatives. For example, Lackner and Skowron (2018) study approval-based rules for electing committees. Further study of other output types seems promising.

References

M. Aizerman and F. Aleskerov. *Theory of Choice*, volume 38 of *Studies in Mathematical and Managerial Economics*. North-Holland, 1995.

J. Alcalde-Unzu and M. Vorsatz. Size approval voting. *Journal of Economic Theory*, 144 (3):1187–1210, 2009.

C. Alós-Ferrer. A simple characterization of approval voting. *Social Choice and Welfare*, 27(3):621–625, 2006.

N. Baigent and Y. Xu. Independent necessary and sufficient conditions for approval voting. *Mathematical Social Sciences*, 21(1):21–29, 1991.

S. J. Brams and P. C. Fishburn. Approval voting. *The American Political Science Review*, 72(3):831–847, 1978.

S. J. Brams and P. C. Fishburn. *Approval Voting*. Boston: Birkhäuser, 1983.

F. Brandt. Set-monotonicity implies Kelly-strategyproofness. *Social Choice and Welfare*, 45(4):793–804, 2015.

F. Brandt and P. Harrenstein. Set-rationalizable choice and self-stability. *Journal of Economic Theory*, 146(4):1721–1731, 2011.

F. Brandt, C. Saile, and C. Stricker. Voting with ties: Strong impossibilities via SAT solving. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1285–1293. IFAAMAS, 2018a.

F. Brandt, C. Saile, and C. Stricker. Voting with ties: Strong impossibilities via SAT solving. In *Proceedings of the 7th International Workshop on Computational Social Choice (COMSOC)*, 2018b.

F. Brandt, C. Saile, and C. Stricker. Strategyproof social choice when preferences and outcomes may contain ties. 2019. Working paper.

O. Cailloux and U. Endriss. Arguing about voting rules. In *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 287–295. IFAAMAS, 2016.

H. Chernoff. Rational selection of decision functions. *Econometrica*, 22(4):422–443, 1954.

S. Ching. A simple characterization of plurality rule. *Journal of Economic Theory*, 71(1): 298–302, 1996.

P. C. Fishburn. Lotteries and social choices. *Journal of Economic Theory*, 5(2):189–207, 1972.

P. C. Fishburn. Axioms for approval voting: Direct proof. *Journal of Economic Theory*, 19(1):180–185, 1978.

P. C. Fishburn. Symmetric and consistent aggregation with dichotomous voting. In J. J. Laffont, editor, *Aggregation and Revelation of Preferences*. North-Holland, 1979.

K. Inada. The simple majority decision rule. *Econometrica*, 37:490–506, 1969.

J. S. Kelly. Strategy-proofness and social choice functions without single-valuedness. *Econometrica*, 45(2):439–446, 1977.

M. Lackner and P. Skowron. Consistent approval-based multi-winner rules. In *Proceedings of the 19th ACM Conference on Economics and Computation (ACM-EC)*, pages 47–48, 2018.

J.-F. Laslier. And the loser is... plurality voting. In D. Felsenthal and M. Machover, editors, *Electoral Systems*, Studies in Choice and Welfare, pages 327–351. Springer, 2011.

J.-F. Laslier and M. R. Sanver, editors. *Handbook on Approval Voting*. Studies in Choice and Welfare. Springer-Verlag, 2010.

D. Lepelley. Une caractérisation du vote à la majorité simple. *RAIRO. Recherche opérationnelle*, 26(4):361–365, 1992.

J. Massó and M. Vorsatz. Weighted approval voting. *Economic Theory*, 36(1):129–146, 2008.

A. Morkelyunas. On two choice rules similar to the plurality rule and Borda's rule. *Matematicheskie Metody v Sotsial'nyh Naukah (Mathematical Methods in Social Sciences)*, 15: 27–36, 1982. In Russian. Translated to English in Ben Silver, editor, *Fourteen Papers Translated from the Russian*, American Mathematical Society translations, ser. 2, vol. 136, 1987.

R. B. Myerson. Axiomatic derivation of scoring rules without the ordering assumption. *Social Choice and Welfare*, 12(1):59–74, 1995.

U. Ninjbat. Approval voting without faithfulness. *Czech Economic Review*, 7(1):15–20, 2012.

Z. E. Öztürk. A composition-consistency characterization of the plurality rule. 2017. Working paper.

M. Pivato. Variable-population voting rules. *Journal of Mathematical Economics*, 49(3): 210–221, 2013.

J. Richelson. A characterization result for the plurality rule. *Journal of Economic Theory*, 19(2):548–550, 1978.

D. G. Saari and S. Barney. Consequences of reversing preferences. *Mathematical Intelligencer*, 25:17–31, 2003.

R. Sanver. Scoring rules cannot respect majority in choice and elimination simultaneously. *Mathematical Social Sciences*, 43(2):151–155, 2002.

A. K. Sen. Social choice theory: A re-examination. *Econometrica*, 45(1):53–89, 1977.

M. R. Sertel. Characterizing approval voting. *Journal of Economic Theory*, 45(1):207–211, 1988.

J. H. Smith. Aggregation of preferences with variable electorate. *Econometrica*, 41(6): 1027–1041, 1973.

T. N. Tideman. Independence of clones as a criterion for voting rules. *Social Choice and Welfare*, 4(3):185–206, 1987.

M. Vorsatz. Approval voting on dichotomous preferences. *Social Choice and Welfare*, 28 (1):127–141, 2007.

M. Vorsatz. Scoring rules on dichotomous preferences. *Social Choice and Welfare*, 31(1): 151–162, 2008.

H. P. Young. An axiomatization of Borda's rule. *Journal of Economic Theory*, 9(1):43–52, 1974.

H. P. Young. Social choice scoring functions. *SIAM Journal on Applied Mathematics*, 28 (4):824–838, 1975.

H. P. Young and A. Levenglick. A consistent extension of Condorcet's election principle. *SIAM Journal on Applied Mathematics*, 35(2):285–300, 1978.

APPENDIX

A. Omitted Proofs

Theorem 6. *AV is the only non-trivial voting rule satisfying reinforcement, neutrality, continuity, and majority consistency.*

Proof. We show that any such voting rule satisfies faithfulness and disjoint equality and then apply Theorem 1. First we prove faithfulness. Let A be some ballot. Neutrality and majority consistency imply that $f(A) \in \{X, A\}$. If $f(A) = X$ for all ballots A, then by reinforcement, $f = TRIV$, which is contrary to the assumption. So there is some ballot A such that $f(A) = A$. For $c \in X \setminus A$, we have that $f(\{c\}) \in \{X, \{c\}\}$. If $f(\{c\}) = X$, then reinforcement implies that $f(A + 2\{c\}) = A$, which contradicts majority consistency. So by neutrality, $f(\{c\}) = \{c\}$ for all alternatives c. Now let A be an arbitrary ballot and assume for contradiction that $f(A) = X$. For $c \in X \setminus A$, we have $f(\{c\}) = \{c\}$ and so reinforcement implies $f(2A + \{c\}) = \{c\}$, which contradicts majority consistency. Thus, $f(A) = A$ and f satisfies faithfulness.

Second we prove disjoint equality. Let A, B be two disjoint ballots. Neutrality implies that $f(A+B) \in \{X, A, B, A\cup B, X\setminus A, X\setminus B, X\setminus(A\cup B)\}$. We will exclude all possibilities except for $f(A + B) = A \cup B$. Assume for contradiction that $f(A + B)$ contains some $c \in X\setminus(A\cup B)$. Moreover, let $a \in A$. Faithfulness implies that $f(A\setminus\{a\}\cup\{c\}) = A\setminus\{a\}\cup\{c\}$. Reinforcement then implies $f(A + B + A \setminus \{a\} \cup \{c\}) = f(A + B) \cap (A \setminus \{a\} \cup \{c\})$. In particular, $c \in f(A+B+A\setminus\{a\}\cup\{c\})$ and $a \notin f(A+B+A\setminus\{a\}\cup\{c\})$. This contradicts neutrality, since the profile $A+B+A\setminus\{a\}\cup\{c\}$ is symmetric with respect to a and c, so that either both of them or neither of them has to be chosen. Thus, $f(A+B) \in \{A\cup B, A, B\}$. The remainder of the proof excludes the latter two possibilities. It proceeds along three increasingly general cases.

Case 1: $A = \{a\}$ and $B = \{b\}$ for two alternatives a, b. By the previous analysis and neutrality, remains $f(\{a\} + \{b\}) = \{a, b\}$ as the only possibility.

Case 2: Arbitrary cardinality A and $B = \{b\}$ for some alternative b. If $f(A+\{b\}) = \{b\}$, continuity implies that $f(A + k(A + \{b\})) = \{b\}$ for some integer k, which contradicts majority consistency. If $f(A + \{b\}) = A$, reinforcement and Case 1 imply that for $a \in A$, $f(A+\{b\}+\{a\}+\{b\}) = \{a\}$. Then, continuity implies that $f(\{b\}+k(A+\{b\}+\{a\}+\{b\})) = \{a\}$ for some integer k, which contradicts majority consistency. Thus, $f(A+\{b\}) = A\cup\{b\}$.

Case 3: Arbitrary cardinality A, B. If $f(A + B) = B$, choose some $b \in B$. Case 2 implies that $f(A + \{b\}) = A \cup \{b\}$. So $f(A + B + A + \{b\}) = \{b\}$. Then continuity implies that $f(A + k(A + B + A + \{b\})) = \{b\}$ for some k, which contradicts majority consistency. The case $f(A + B) = A$ is analogous. Thus, $f(A + B) = A \cup B$ remains as the only possibility. □

Theorem 7. *AV is the only voting rule satisfying reinforcement, clone consistency, and faithfulness if $|X| \geqslant 4$.*

Proof. Take any such voting rule f and some agenda $Y \in \mathcal{P}(X)$. We will omit the subscript Y for f within this proof. We show that f satisfies disjoint equality and then apply Theorem 1.

Let $A, B \in \mathcal{A}_Y$ be two disjoint ballots and $C = Y \setminus (A \cup B)$. In the two-voter profile $A + B$ all alternatives in A are clones of each other and likewise all alternatives in B and in C. So clone consistency implies that for each of these sets, either all alternatives are chosen or none, i.e., $f(A + B) \in \{Y, A \cup B, A, B, C\}$.

We show that in addition, A is chosen if and only if B is chosen. So assume that $A \subseteq f(A+B)$, let $a \in A$, $b \in B$, and $c \in C$, and consider the agenda $\{a, b, c\}$. (If $A \cup B = Y$, omit c in what follows.) Clone consistency implies that $a \in f_{\{a,b,c\}}(A+B) = f(A+B) \cap \{a, b, c\}$. Let $x \in X \setminus \{a, b, c\}$, which exists since $|X| \geqslant 4$, and consider the profile $\{a, x\} + \{b\}$ on the agenda $\{a, b, c, x\}$. Clone consistency implies $x \in f_{\{a,b,c,x\}}(\{a, x\} + \{b\})$. Further applications of clone consistency imply that $x \in f_{\{b,c,x\}}(\{x\} + \{b\})$, $x \in f_{\{a,b,c,x\}}(\{x\} + \{a, b\})$, $x \in f_{\{a,c,x\}}(\{x\} + \{a\})$, $b \in f_{\{a,b,c,x\}}(\{b, x\} + \{a\})$, and $b \in f_{\{a,b,c\}}(\{b\} + \{a\})$. So we get $b \in f_{\{a,b,c\}}(A + B) = f(A + B) \cap \{a, b, c\}$ and thus, $B \subseteq f(A + B)$.

The remaining possibilities are $f(A+B) \in \{Y, A \cup B, C\}$. Assume for contradiction that $f(A + B) \cap C \neq \emptyset$. Faithfulness implies $f(C) = C$. Applying reinforcement to $A + B$ and the one-voter profile C yields $f(A + B + C) = f(A + B) \cap f(C) = C$. Essentially the same line of reasoning as in the previous paragraph shows that $f(A + B + C) = Y$, which is a contradiction. Thus $f(A + B) = A \cup B$, and so f satisfies disjoint equality. □

Theorem 8. *AV is the only voting rule satisfying reinforcement, neutrality, faithfulness, and independence of losers.*

Proof. Take any such voting rule f and some agenda $Y \in \mathcal{P}(X)$. We will omit the subscript Y for f within this proof. We show that f satisfies disjoint equality and then apply Theorem 1.

To this end, let A, B be two disjoint ballots and $a \in A$ and $b \in B$. Neutrality implies that $f(A + B) \in \{X, A \cup B, A, B, X \setminus B, X \setminus A\}$. First assume for contradiction that $f(A+B) \not\subseteq A \cup B$ and let $c \in f(A + B) \setminus (A \cup B)$. By faithfulness, we have $f(\{c\}) = \{c\}$. Hence, by reinforcement, $f(A + B + \{c\}) = \{c\}$. Independence of losers implies that $f(A+B+\{c\}) = f_{\{a,b,c\}}(\{a\} + \{b\} + \{c\}) = \{c\}$, which contradicts neutrality. So we have $f(A + B) \in \{A \cup B, A, B\}$.

Second, consider the case $f(A + B) = A$. From the previous case and neutrality, we know that $f(\{a\} + \{b\}) = \{a, b\}$. Hence, by reinforcement, $f(A + B + \{a\} + \{b\}) = \{a\}$. Independence of losers implies that $f(A+B+\{a\}+\{b\}) = f_{\{a,b\}}(\{a\}+\{b\}+\{a\}+\{b\}) = \{a\}$, which again contradicts neutrality. Similarly, we get a contradiction if $f(A + B) = B$. Hence, $f(A+B) = A \cup B$ remains as the only possibility and so f satisfies disjoint equality as desired. □

Theorem 9. *AV is the only voting rule satisfying reinforcement, neutrality, and independence of Pareto dominated alternatives.*

Proof. Fix some agenda Y. We show that f_Y satisfies disjoint equality and then invoke Theorem 2 to conclude $f_Y = AV_Y$. The subscript Y will be omitted in the rest of the proof. Let A, B be two disjoint ballots. Observe that in the profile $A + B$ all alternatives in $X \setminus (A \cup B)$ are Pareto dominated by alternatives in $A \cup B$. So independence of Pareto dominated alternatives implies $f(A + B) = f_{A \cup B}(A + B) \subseteq A \cup B$. Then it follows from neutrality that $f(A + B) \in \{A, B, A \cup B\}$. Without loss of generality, we may assume for contradiction that $f(A + B) = A$. For $a \in A$ and $b \in B$, neutrality and independence of Pareto dominated alternatives imply $f(\{a\} + \{b\}) = f_{\{a,b\}}(\{a\} + \{b\}) = \{a, b\}$. Then it follows from reinforcement that $f(A+B+\{a\}+\{b\}) = \{a\}$. In the profile $A+B+\{a\}+\{b\}$, all alternatives except a and b are Pareto dominated by either a or b. Thus, by independence of Pareto dominated alternatives, $\{a\} = f(A+B+\{a\}+\{b\}) = f_{\{a,b\}}(\{a\}+\{b\}+\{a\}+\{b\})$, which contradicts neutrality. $\square$