An Axiomatic Characterization of the Borda Mean Rule

Florian Brandl · **Dominik Peters**

May 11, 2017

Abstract Given a preference profile, a social dichotomy function partitions the set of alternatives into a set of approved alternatives and a set of disapproved alternatives. The Borda mean rule approves all alternatives with above-average Borda score, and disapproves alternatives with below-average Borda score. We show that the Borda mean rule is the unique social dichotomy function satisfying neutrality, reinforcement, faithfulness, and the quasi-Condorcet property.

1 Introduction

The objective of social choice is typically to choose the best alternatives from a set of feasible alternatives based on the preferences of various voters. Functions that describe how this choice is made are called *social choice functions*. Hence, a social choice function partitions the set of alternatives into winning alternatives and non-winning (or losing) alternatives. Suppose instead that the goal is to split the alternatives into good alternatives and bad alternatives with the *separation* between both sets being as large as possible. Similarly, one might ask both sets to be as homogeneous as possible. Duddy et al. (2014) argued that social choice functions are not the right tool for this task. Let us give an example to illustrate why this is the case. Consider a class of students that is to be divided into beginners and advanced learners based on how they are ranked by teachers. Presumably, the goal should be to form two groups of students such that the differences in skill level within each group are as small as possible. If all teachers agree on their top-ranked student, any reasonable social choice function

F. Brandl
Department of Informatics
Technical University of Munich
E-mail: brandlfl@in.tum.de

D. Peters
Department of Computer Science
University of Oxford
E-mail: dominik.peters@cs.ox.ac.uk

would uniquely choose the unanimously top-ranked student. Hence, the group of advanced learners would consist of only this one student; all other students are put into the beginners group. In our example this is likely to be an undesired result, since the differences in skill within the beginners group would be barely reduced compared to the entire class.

Thus, we need to drop some of the properties that seem appealing for social choice functions. A more suitable tool for our task are *social dichotomy functions* (Duddy et al., 2014), which yield ordered 2-partitions of the alternatives. We interpret ordered 2-partitions as having the approved alternatives in the first set and the disapproved alternatives in the second set. In contrast to selecting the best alternatives, there is inherent symmetry in the problem of finding a good separation; in particular, rules should usually satisfy *reversal symmetry*: if all of the input preferences are reversed, then the output will also be reversed, so that approved and disapproved alternatives swap place.

The social dichotomy function that we consider in this paper is the *Borda mean rule* which outputs all dichotomous weak orders in which all alternatives with above-average Borda score are approved, and all alternatives with below-average Borda score are disapproved. If there are alternatives with precisely average Borda score, then the rule returns several orders with all ways of breaking the ties. This rule was introduced by Duddy et al. (2014) and further discussed by Duddy et al. (2016) and Zwicker (2016). Notice that the Borda mean rule satisfies reversal symmetry.

Reversal symmetry (similarly defined) is also a natural property for *social preference functions*, which return a set of *linear* orders of the alternatives based on the voters' preferences. We will see that social dichotomy functions are more closely related to social preference functions than to social choice functions. Kemeny's rule (Kemeny, 1959) is an example of a social preference function that has been very influential in social choice theory and it is widely seen as an attractive rule with desirable properties (e.g., Young, 1995). Given a preference profile over an alternative set A, the rule assigns to each possible preference relation $\succcurlyeq$ a *Kemeny score*: the order gets a point for each voter and each pair of alternatives $a,b \in A$ such that $a \succ b$ and the voter agrees with this choice; if the voter disagrees, the order loses a point. Preference relations with maximum score are called *Kemeny rankings*. Kemeny's rule returns exactly the set of all *linear* orders that are Kemeny rankings.

Zwicker (2016) introduced the idea of using Kemeny scores to define aggregation rules for other output types. For example, if we consider the domain of preference relations that have a unique most-preferred element and that are indifferent between all other alternatives, then the rule selecting the orderings from this domain of highest Kemeny score outputs the relations with the winners of *Borda's rule* as most-preferred alternatives. In his paper, Zwicker (2016) proposed the *k-Kemeny rule* which returns the k-chotomous weak order of highest Kemeny score; a preference relation $\succcurlyeq$ is called *k-chotomous* if its induced indifference relation $\sim$ partitions A into at most k indifference classes: thus, they define an ordered k-partition. In particular, 2-chotomous orders are usually called *dichotomous*; these are the orders that partition the alternatives into a set of *approved* and a set of *disapproved* alternatives. Hence, the 2-Kemeny rule is a social dichotomy function. Duddy et al. (2014) showed that the 2-Kemeny rule is identical to the *Borda mean rule*. This equivalent definition of the Borda mean rule suggests that it is a good tool for finding dichotomies that maximize the separation between the set of approved and the set of disapproved alternatives.

Social choice theory abounds with different proposals for voting rules; which of them should we choose to use? Axiomatic characterizations provide some of the strongest reasons in favor of using certain rules. For example, Kemeny's rule is largely seen as a very attractive social preference function because of its characterization by Young and Levenglick (1978) (though there are other reasons as well). In this paper, we present an axiomatic characterization of the Borda mean rule, using the same axioms as the characterization of Kemeny's rule by Young and Levenglick (1978), showing that the above argument in favor of Kemeny's rule applies just as well to the Borda mean rule, hopefully establishing its place as a very natural social dichotomy function. In formal terms, our result is that the Borda mean rule is the unique social dichotomy function satisfying neutrality, reinforcement, faithfulness, and the quasi-Condorcet property.[1] Our proof follows a similar structure as Young's (1974a) characterization of Borda's rule. In particular, we use linear algebra and exploit the orthogonal decomposition of weighted tournaments popularized by Zwicker (1991), but we do not need any convex separation theorems.

Most of our axioms are commonly used, including the uncontroversial axioms of neutrality (requiring that all alternatives are treated equally) and faithfulness (requiring sensible behavior in single-voter situations). Reinforcement (often known as *consistency*) is the workhorse of many axiomatic characterizations in social choice. It is a variable-electorate axiom which requires that if the same dichotomy is selected in two disjoint profiles, then it is still selected if we merge the two profiles into one. Reinforcement is typically satisfied by rules which maximize a sum of the "scores" that each voter assigns to a potential output. The most specialized axiom in our collection is the quasi-Condorcet property, introduced by Young and Levenglick (1978) and also used by Barthélemy and Janowitz (1991). It requires that a "dummy alternative" (one that is tied with every other alternative in a majority comparison) can move around freely within the output relation. (We give a formal definition below.) Since the quasi-Condorcet property implies the cancellation axiom, any rule satisfying it and reinforcement can only depend on the weighted majority relation. The axioms in our collection are independent, in the sense that our result does not hold if any of them is dropped.

2 Related Work

The Borda mean rule was introduced and studied by Duddy et al. (2016) in the special case of *binary aggregation*. In their setting, every voter holds a binary evaluation of the alternatives or, equivalently, a dichotomous preference relation. A *binary aggregation function* maps the voters preferences to an ordered tripartition of approved, tied, and disapproved alternatives. In this setting, the Borda mean rule approves all alternatives with above-average approval score, disapproves all alternatives with below-average

[1] For expository purposes, Young and Levenglick (1978) introduce what they call the "Condorcet axiom" which is a strengthening of faithfulness and quasi-Condorcet property. However, as they note, in their proof this strengthening is not required.

approval score, and alternatives with average approval score are tied. Duddy et al. (2016) show that the Borda mean rule is the only binary aggregation function satisfying faithfulness, consistency, cancellation, and neutrality. Their notion of consistency is a version of Smith's (1973) axiom of *separability*: if an alternative is approved by one electorate and either approved or ranked as tied by another electorate, then it is approved by the union of both electorates (and analogously for disapproved alternatives).

Since social dichotomy functions can be viewed as returning a set of multiple winners, the recent literature on *multiwinner voting rules* is related (for a survey, see Faliszewski et al., 2017). Voting rules in that setting return a committee of k alternatives, where k is fixed. Examples include the k-Borda rule (which returns the k alternatives with highest Borda score, see Debord, 1992), as well as Chamberlin and Courant's (1983) rule and Monroe's (1995) rule which aim for committees providing proportional representation. Note that, in contrast, the definition of a social dichotomy function does not impose any cardinality constraint on the set of approved candidates. Axiomatic characterizations of multiwinner rules using consistency-type axioms are provided by Skowron et al. (2016) for linear order preferences and by Lackner and Skowron (2017) for approval preferences. The k-Borda rule was characterized by Debord (1992); his result is close to ours. The k-Borda rule can be equivalently defined as the rule that returns the Kemeny score-optimal dichotomous orders with exactly k approved candidates.

Many characterizations of Borda's rule as a social choice function, and of scoring rules more generally, are available (for a survey, see Chebotarev and Shamis, 1998). Young (1974a) gave the first characterization of Borda's rule using reinforcement. Hansson and Sahlquist (1976) gave an alternative proof that does not use linear algebra. Young (1975) characterized the class of all scoring rules, and identified Borda among them by adding an additional axiom. Smith (1973) independently found a characterization of scoring rules as social welfare functions; Young (1974b) gave an alternative proof of that result. Fishburn (1978) characterized approval voting using reinforcement, which is just Borda's rule restricted to profiles of dichotomous preferences.

The Borda mean rule is also related to Nanson's rule, which, in order to determine a winner, repeatedly eliminates all alternatives with below-average Borda score (Niou, 1987). The Borda mean rule is just the result of stopping Nanson's procedure after its first round.

The quasi-Condorcet property, a key axiom in our characterization, was introduced by Young and Levenglick (1978) for characterizing Kemeny's rule. The axiom also proved useful in the literature about the *median procedure* for aggregating other kinds of data structures, such as for median semilattices (Barthélemy and Janowitz, 1991) and median graphs (McMorris et al., 2000).

3 Definitions

Let $\mathbb{N} = \{1, 2, \dots\}$ be a set of voters with preferences over a finite set A of alternatives, where $|A| = m$. The preferences of an agent $i \in \mathbb{N}$ are given by a binary relation

$\succcurlyeq_i \subseteq A \times A$ which is complete and transitive; such a relation is called a *preference relation*. We will write $a \succ_i b$ if $a \succcurlyeq_i b$ but $b \not\succcurlyeq_i a$, and $a \sim_i b$ if both $a \succcurlyeq_i b$ and $b \succcurlyeq_i a$.

A preference relation $\succcurlyeq$ is called a *linear order* if it is antisymmetric, so that $a \sim b$ only if $a = b$. A preference relation $\succcurlyeq$ is *dichotomous* if there is a partition (A_1, A_2) of A into two subsets such that $a \succ b$ if and only if $a \in A_1$ and $b \in A_2$. Note that one of A_1 and A_2 may be empty, in which case $\succcurlyeq = A \times A$. Equivalently, an order is dichotomous if and only if there are no three alternatives $a, b, c \in A$ with $a \succ b \succ c$. We will write $\mathscr{R}(A)$ for the set of all preference relations over A and $\mathscr{D}(A)$ for the set of dichotomous orders. When the set A is clear from the context, we write $\mathscr{R}$ and $\mathscr{D}$, respectively.

An *electorate* N is a finite and non-empty subset of $\mathbb{N}$. The set of all electorates is denoted by $\mathscr{F}(\mathbb{N})$. A *preference profile* $P \in \mathscr{R}^N$ on electorate N is a function assigning a preference relation to each voter in N. The preferences of voter i in profile P are then denoted by $\succcurlyeq_i$. A *social dichotomy function (SDF)* f is a map from the set of all profiles to non-empty subsets of $\mathscr{D}$, so that $f(P) \subseteq \mathscr{D}$ for all profiles P.

The *reverse* $\overleftarrow{\succcurlyeq}$ of a preference relation $\succcurlyeq$ is defined by $(a,b) \in \overleftarrow{\succcurlyeq}$ if and only if $(b,a) \in \succcurlyeq$. We extend this concept to *sets* of orders in the natural way, so that, for example, $\overleftarrow{f(P)} = \{\overleftarrow{\succcurlyeq} : \succcurlyeq \in f(P)\}$. If σ is a permutation of A, we can also naturally define the relation $\sigma(\succcurlyeq) = \{(\sigma(a), \sigma(b)) \colon (a,b) \in \succcurlyeq\}$, and extend this definition to sets and profiles of preference relations.

Given a profile P over N, and two alternatives $a, b \in A$, let us write

$$n_{ab} := |\{i \in N \colon a \succ_i b\}|$$

for the number of voters in P who strictly prefer a to b. The *majority margin* of a over b is then given by $m_{ab} := n_{ab} - n_{ba}$; if $m_{ab} > 0$ then a majority of voters prefers a to b. Note that the majority margins form a skew-symmetric $m \times m$ matrix with zeros on the main diagonal (since $m_{ab} = -m_{ba}$).[2] We can interpret this matrix as a *weighted tournament* T whose vertices are given by the alternatives; there is an arc from a to b if and only if $m_{ab} > 0$, and the arc is labelled by m_{ab}.

The (symmetric) *Borda score* $\beta(a)$ of an alternative $a \in A$ is given by

$$\beta(a) := \sum_{b \in A \setminus \{a\}} m_{ab},$$

essentially the net weighted out-degree of a in the weighted tournament induced by P. It is easy to see that β, thus defined, is a positive affine transformation of the Borda scores as defined through the usual scoring vector $(m-1, m-2, \ldots, 1, 0)$; indeed the scoring-based Borda score of a is $\beta(a)/2 + |N|(m-1)/2$. Thus, for example, the same alternatives are Borda winners for either definition of Borda scores. Note that, because the majority margins are skew-symmetric, we have $\sum_{a \in A} \beta(a) = 0$, and so the average (symmetric) Borda score of the alternatives is always 0, which makes it convenient to deal with symmetric Borda scores.

[2] A matrix $M \in \mathbb{R}^{m \times m}$ is skew-symmetric if $M = -M^T$.

4 Borda Mean Rule

As we have mentioned in the introduction, there are several equivalent ways of defining the Borda mean rule. The most straightforward definition uses the average Borda score directly:

$$BM(P) = \Big\{ \succcurlyeq \in \mathscr{D} \colon a \succ b \text{ for all } a, b \in A \text{ with } \beta(a) > 0 \text{ and } \beta(b) < 0 \Big\}.$$

Thus, the Borda mean rule returns all dichotomous preference relations where alternatives with above-average Borda score are placed in the upper indifference class and alternatives with below-average Borda score are placed in the lower indifference class. (Recall that, for symmetric Borda scores, the average Borda score is always 0.) Alternatives with exactly average Borda score are placed once in the upper and once in the lower indifference class (so that multiple rankings are returned).

In the framework of Zwicker (2016), the Borda mean rule is obtained as a special case of Kemeny's rule with dichotomous output. Precisely, the Borda mean rule is the rule returning the dichotomous preference relations of maximum Kemeny score:

$$BM(P) = \arg\max_{\succcurlyeq \in \mathscr{D}} \sum_{x \succ y} m_{xy}.$$

Hence, the Borda mean rule minimizes the aggregate distance of $\succcurlyeq$ to the voters' preferences or, alternatively, maximizes the agreement with the voters' preferences.

It can be observed from the definition that the Borda mean rule only depends on the pairwise majority margins and hence on the weighted tournament induced by a preference profile. Thus, the Borda mean rule is a *C2 rule* in Fishburn's (1977) classification. This property will play an important role in our characterization. An interesting property of the Borda mean rule is that it always approves Condorcet winners and always disapproves Condorcet losers, provided they exist. This can be seen by recalling that if a is the Condorcet winner, then $\beta(a) > 0$ from the definition of β, and similarly for Condorcet losers; alternatively one can note that the Kemeny score of a dichotomy $\succcurlyeq$ strictly improves if we move the Condorcet winner from the lower to the upper indifference class.

To help us understand the Borda mean rule, we discuss its behavior for the weighted tournaments given in Figure 1. In the tournament T, all alternatives have Borda score 0 (later we will say that T is *purely cyclic*), and so we have $BM(T) = \mathscr{D}$. The tournament T' is *purely cocyclic* with Borda scores $3, 0, -3$ for x, y, z, respectively. Hence $BM(T') = \{\{x,y\} \succcurlyeq \{z\}, \{x\} \succcurlyeq \{y,z\}\}$. T'' has a cyclic part and a cocyclic part. The Borda scores are $2, 1, -3$ for x, y, z, which implies that $BM(T'') = \{\{x,y\} \succcurlyeq \{z\}\}$.

5 Axioms

One may think that SDFs are a well-studied object, since every SCF partitions the set of alternatives into a set of winning alternatives and a set of losing alternatives and hence induces an SDF. As discussed by Duddy et al. (2014), this way of constructing SDFs seems to miss the point, since SCFs aim to select a set of "good" alternatives,

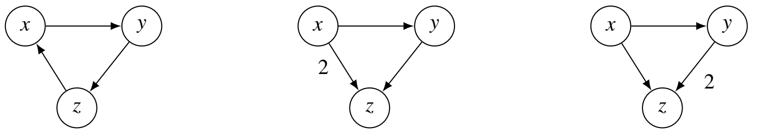

(a) Cyclic tournament T (b) Cocyclic tournament T' (c) Mixed tournament T''

Fig. 1: Examples for the Borda mean rule. The weight of an edge denotes the majority margin between the two adjacent alternatives. Unlabelled edges have weight 1.

typically as small as possible, and not a set that maximizes separation with its complement. *Reversal symmetry* formalizes this crucial difference: if all voters reverse their preferences, then the approved set becomes the disapproved set and *vice versa*. Formally, an SDF satisfies reversal symmetry if

$$f(\overleftarrow{P}) = \overleftarrow{f(P)} \quad \text{for all } P \in \mathscr{R}^N. \qquad \text{(Reversal symmetry)}$$

While the Borda mean rule satisfies reversal symmetry, we do not impose this axiom for our characterization. Instead, we use the same four axioms that also feature in Young and Levenglick's (1978) characterization of Kemeny's rule. First, we require SDFs to satisfy *neutrality*: renaming the alternatives in a preference profile leads to the same renaming in the output relations. Neutrality thus prescribes that an SDF is symmetric with respect to the alternatives and prevents it from being biased towards certain alternatives. Let $\Pi(A)$ denote the set of all permutations on A. Then, an SDF f satisfies neutrality if

$$f(\sigma(P)) = \sigma(f(P)) \quad \text{for all } P \in \mathscr{R}^N, N \in \mathscr{F}(\mathbb{N}), \text{and } \sigma \in \Pi(A). \qquad \text{(Neutrality)}$$

When dealing with variable electorates, it seems reasonable to require that if two disjoint electorates agree on the same ranking, then this ranking should be returned when they are combined into one electorate. This is known as *reinforcement*. An SDF f satisfies reinforcement if

$$f(P) \cap f(P') \neq \varnothing \text{ implies } f(P) \cap f(P') = f(P \cup P') \quad \text{for all } P \in \mathscr{R}^N \text{ and } P' \in \mathscr{R}^{N'}$$
$$\text{with } N \cap N' = \varnothing. \qquad \text{(Reinforcement)}$$

Notice that reinforcement is agnostic about the type of output. It may be defined in the same way for every kind of aggregation function, such as social choice functions (which return a subset of alternatives) or social preference functions (which return a set of linear orders of the alternatives). Reinforcement was introduced by Young (1974a, 1975) (he called it "consistency") to characterize scoring rules; the axiom is related to "separability" introduced by Smith (1973) (now often called consistency) for social welfare functions.

Our axioms so far are completely oblivious of the meaning of preferences. Without an axiom that prescribes some degree of correlation of the voters' preferences with the aggregated preferences, the trivial SDF always returning all dichotomies is not ruled out. An arguably minimal axiom of this nature is *faithfulness*, which requires that whenever the electorate consists of one voter with dichotomous preferences then his preferences should be uniquely returned. Formally, an SDF f satisfies faithfulness if

$$f(P) = \{\succcurlyeq_i\} \quad \text{for all } P \in \mathscr{R}^{\{i\}} \text{ and } i \in \mathbb{N} \text{ with } \succcurlyeq_i \in \mathscr{D}. \qquad \text{(Faithfulness)}$$

Lastly, we consider an axiom that specifies how to deal with "dummy" alternatives that are independent from the others in the sense that they are tied with every other alternative in a pairwise majority comparison. Formally, an alternative $x \in A$ is a *dummy* if $n_{xy} = n_{yx}$ for all $y \in A$. The *quasi-Condorcet property* asserts that dummy alternatives can be placed arbitrarily in the output ranking. To formalize this, let $\succcurlyeq|_{A\setminus\{x\}}$ be the preference relation on $A \setminus \{x\}$ obtained by restricting $\succcurlyeq \in \mathscr{R}$ to alternatives in $A \setminus \{x\}$. Restriction of a preference profile is defined by restricting each voter's preference relation. If $\hat{\succcurlyeq}$ is a dichotomous preference relation on $A \setminus \{x\}$, then $\hat{\succcurlyeq}\langle x\rangle = \{\succcurlyeq \in \mathscr{D}(A) \colon \succcurlyeq|_{A\setminus\{x\}} = \hat{\succcurlyeq}\}$ is the set of dichotomous preference relations on A obtained by adding x to $\hat{\succcurlyeq}$, once as approved and once as disapproved. For a set S of preference relations on $A \setminus \{x\}$, we define this operation by $S\langle x\rangle = \bigcup_{\hat{\succcurlyeq} \in S} \hat{\succcurlyeq}\langle x\rangle$. We say that an SDF f satisfies the quasi-Condorcet property if

$$f(P) = f(P|_{A\setminus\{x\}})\langle x\rangle \quad \text{for all } P \in \mathscr{R}^N, N \in \mathscr{F}(\mathbb{N}), \text{ and } x \in A \text{ with } n_{xy} = n_{yx} \text{ for all } y \in A. \qquad \text{(Quasi-Condorcet property)}$$

The quasi-Condorcet property is a strengthening of the *cancellation* axiom, which requires that all dichotomies are returned whenever all majority margins are zero. Formally, f satisfies cancellation if

$$f(P) = \mathscr{D} \quad \text{for all } P \in \mathscr{R}^N, N \in \mathscr{F}(\mathbb{N}) \text{ with } n_{xy} = n_{yx} \text{ for all } x, y \in A. \qquad \text{(Cancellation)}$$

Within the class of scoring rules, the cancellation axiom (for SCFs) characterizes Borda's rule (Young, 1975).

6 The Linear Algebra of Weighted Tournaments

Let us define a few special weighted tournaments that will be useful later, see Figure 2 for drawings. Given three alternatives $x, y, z \in A$, we write C_{xyz} for the weighted tournament with $m_{xy} = m_{yz} = m_{zx} = 1$ and $m_{yx} = m_{zy} = m_{xz} = -1$, and all non-specified values 0. Thus, C_{xyz} is a 3-cycle. Next, given a set $X \subseteq A$ of alternatives, we write D_X for the weighted tournament with

$$m_{ab} = \begin{cases} 1 & \text{if } a \in X, b \notin X, \\ -1 & \text{if } a \notin X, b \in X, \\ 0 & \text{otherwise.} \end{cases}$$

Thus, D_X is the weighted tournament induced by a profile containing a single dichotomous voter i with $X \succ_i A \setminus X$. Finally, for alternatives $x, y \in A$ we will need the weighted tournament $S^x_y = D_{\{x\}} + D_{A\setminus\{y\}}$ which consists of a single "top" alternative x, a single "bottom" alternative y, and all other alternatives in between.

For our characterization, it will be useful to understand the structure of weighted tournaments better, and so we give a brief introduction to their linear algebra. Let V be

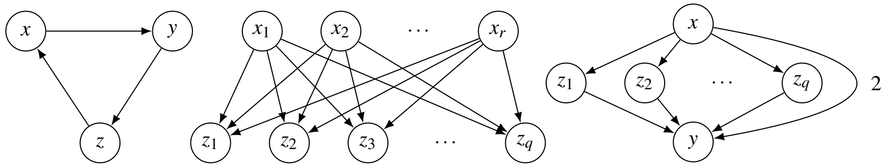

(a) Cycle C_{xyz} (b) Tournament D_X (c) Tournament S^x_y

Fig. 2: Some types of tournaments.

the vector space of rational-valued skew-symmetric $m \times m$ matrices (and, equivalently, of weighted tournaments). Note that $\dim V = \binom{m}{2}$. This vector space can be endowed with the usual inner product, identifying a skew-symmetric matrix with a vector in $\mathbb{Q}^{m(m-1)/2}$. We will be interested in an orthogonal decomposition of V into two subspaces:

$$V = V_{\text{cycle}} \oplus V_{\text{cocycle}},$$

where V_{cycle} is the *cycle space* (of dimension $\binom{m}{2} - (m-1)$) and V_{cocycle} is the *cocycle space* (of dimension $m-1$). The cycle space $V_{\text{cycle}} = \langle C_{xyz} \colon x,y,z \in A \rangle$ is defined as the span of all 3-cycles (equivalently, the span of all simple cycles). The cocycle space $V_{\text{cocycle}} = \langle D_X : X \subseteq A \rangle$ is defined as the span of all tournaments D_X (equivalently, the span of all $D_{\{x\}}$). It can be checked that these two subspaces are orthogonal and jointly span V.

Proposition 1 *The subspaces V_{cycle} and $V_{cocycle}$ are orthogonal and jointly span V, that is, $V = V_{cycle} \oplus V_{cocycle}$.*

With this decomposition, given a weighted tournament T, we can uniquely write $T = T_{\text{cycle}} + T_{\text{cocycle}}$, where $T_{\text{cycle}} \in V_{\text{cycle}}$ is the *cyclic component* of T and $T_{\text{cocycle}} \in V_{\text{cocycle}}$ is the *cocyclic component* of T. We say that T is *purely cocyclic* if $T = T_{\text{cocycle}}$, i.e., if $T_{\text{cycle}} = 0$. Of the examples in Figure 2, C_{xyz} is purely cyclic, and D_X and S^x_y are purely cocyclic; in Figure 1, the tournament T'' has non-zero cyclic and cocyclic components.

Let us next give a convenient characterization of purely cocyclic tournaments.

Lemma 1 **(Zwicker, 2016)** *A weighted tournament T is purely cocyclic if and only if it is* difference generated*, i.e., there exists a function $\gamma \colon A \to \mathbb{R}$ such that $m_{ab} = \gamma(a) - \gamma(b)$ for all $a,b \in A$. In fact, if T is purely cocyclic, then it is difference generated by $\gamma(a) := \beta(a)/m$, i.e., by Borda scores, suitably rescaled.*

Obviously, a difference generated tournament cannot contain a cycle. For example, the tournament D_X is difference generated with $\gamma(x) = 1$ for all $x \in X$ and $\gamma(z) = 0$ for all $z \notin X$. The tournament S^x_y is difference generated with $\gamma(x) = 1$, $\gamma(y) = -1$, and $\gamma(z) = 0$ for $z \in A \setminus \{x,y\}$.

Proof (of Lemma 1) If tournaments T_1 and T_2 are difference generated by γ_1 and γ_2 respectively, then it is easy to see that $\alpha T_1 + \beta T_2$ is difference generated by $\alpha\gamma_1 + \beta\gamma_2$. As we noted above, the tournaments D_X are difference generated. Hence all tournaments in the space V_{cocycle}, which is spanned by the tournaments D_X, are difference generated.

Suppose T is difference generated by γ, and consider the tournament C_{xyz} for some $x,y,z \in A$. Then

$$T \cdot C_{xyz} = (\gamma(x) - \gamma(y)) + (\gamma(y) - \gamma(z)) + (\gamma(z) - \gamma(x)) = 0.$$

Hence T is orthogonal to every C_{xyz}, and thus it is orthogonal to every weighted tournament spanned by 3-cycles. Hence T is orthogonal to V_{cycle}, and hence $T = T_{\text{cocycle}}$, so that T is purely cocyclic. □

7 Characterization

We are now ready to state and prove our main result.

Theorem 1 *An SDF f satisfies neutrality, reinforcement, faithfulness, and the quasi-Condorcet property if and only if f is the Borda mean rule.*

The fact that the Borda mean rule satisfies all four axioms follows readily from the definition. Hence we only prove the "only if" part of Theorem 1. Its proof is split up into six statements.

Our first lemma is also part of Young's (1974a) characterization of Borda's rule. Its conclusion does not depend on the type of output of f, and in particular also holds for social choice functions and social preference functions (for the appropriate definition of cancellation). We include the proof for completeness.

Lemma 2 **(Young, 1974a)** *If an SDF f satisfies reinforcement and cancellation, then f only depends on the majority margins.*

Proof Suppose f satisfies reinforcement and cancellation, and let P_1 and P_2 be profiles that induce the same majority margins. Assume first that P_1 and P_2 are defined on disjoint electorates. Let $Q = \overleftarrow{P_1}$, interpreted as having an electorate disjoint from those of P_1 and P_2. Since $P_1 \cup Q$ and $P_2 \cup Q$ both induce the empty weighted tournament, by cancellation, we have that $f(P_1 \cup Q) = \mathscr{D} = f(P_2 \cup Q)$. Hence, using reinforcement twice,

$$\begin{aligned} f(P_1) = f(P_1) \cap \mathscr{D} &= f(P_1 \cup (P_2 \cup Q)) \\ &= f((P_1 \cup Q) \cup P_2) = \mathscr{D} \cap f(P_2) = f(P_2). \end{aligned}$$

For profiles P_1 and P_2 whose electorates are not disjoint, find a profile P_3 whose electorate is disjoint from both P_1 and P_2, and so that P_3 induces the same majority margins as P_1 and P_2. Using the argument above twice, we have $f(P_1) = f(P_3) = f(P_2)$. □

Since the quasi-Condorcet property implies cancellation, this shows that f only depends on the majority margins induced by a preference profile. Thus, f is a C2 rule in the sense of Fishburn (1977). In particular, this implies that f is *anonymous*, i.e., the outcome is invariant under renaming the voters.

Lemma 3 **(Young, 1974a)** *If an SDF f satisfies reinforcement and cancellation, then f can be uniquely extended to the domain V of* all rational *weighted tournaments, in a way that preserves reinforcement, neutrality, faithfulness, and the quasi-Condorcet property.*

Proof For any $T \in V$ and natural number $n \in \mathbb{N}$, define $f(\frac{1}{n}T) = f(T)$. This is well-defined since if $\frac{1}{n}T = \frac{1}{n'}T'$, then $n'T = nT'$, and then by reinforcement and definition,

$$f(\tfrac{1}{n}T) = f(T) = f(n'T) = f(nT') = f(T') = f(\tfrac{1}{n'}T'). \qquad \square$$

Lemma 3 enables us to change the domain of f from preference profiles to V, the set of rational-valued skew-symmetric matrices (equivalently, weighted tournaments), as f is invariant on the set of profiles that induce a given weighted tournament. As this is more convenient to work with, we will view f as a function with domain V from now on.

The next lemma and the resulting corollary show that the cyclic part of a weighted tournament can be ignored when computing the outcome of f. Since the cocyclic part is completely determined by the Borda scores (by Lemma 1), we see that f only depends on the Borda scores. To achieve this result, we will show that f is trivial on purely cyclic tournaments, in the sense of returning all dichotomies. As a first step, we show this for the building blocks C_{xyz} of the cycle space, using an argument that makes heavy use of neutrality.

Lemma 4 *Suppose $A = \{a,b,c\}$ consists of three alternatives. If f satisfies neutrality, reinforcement, and cancellation, then $f(C_{abc}) = \mathscr{D}$.*

Proof Let $C = C_{abc}$ and $\succcurlyeq \in f(C)$. Let $\sigma = (a\,b\,c)$ and observe that $C = \sigma(C)$. Thus, by neutrality of f, we must have $\sigma(\succcurlyeq) \in f(C)$ and $\sigma^2(\succcurlyeq) \in f(C)$. Hence, either

$$\{a,b,c\} \in f(C) \quad \text{or} \quad \{a\} \succ \{b,c\} \in f(C) \quad \text{or} \quad \{b,c\} \succ \{a\} \in f(C).$$

Now consider the permutation $\hat{\sigma}$ of A that transposes b and c, i.e., $\hat{\sigma} = (a)(b\,c)$. Then, by neutrality, $\hat{\sigma}(f(C)) = f(\hat{\sigma}(C))$. Thus, in each of the three cases,

$$f(C) \cap f(\hat{\sigma}(C)) = f(C) \cap \hat{\sigma}(f(C)) \neq \varnothing.$$

Hence, reinforcement and cancellation imply that

$$f(C) \cap f(\hat{\sigma}(C)) = f(C \cup \hat{\sigma}(C)) = \mathscr{D}.$$

Hence, $f(C) = \mathscr{D}$. $\square$

Next, we lift the result for the C_{xyz} to apply to all tournaments in V_{cycle}.

Corollary 1 *If an SDF f satisfies neutrality, reinforcement, and the quasi-Condorcet property, then f depends only on Borda scores.*

Proof Let T be any weighted tournament, and consider its orthogonal decomposition $T = T_{\text{cycle}} + T_{\text{cocycle}}$. We will show that $f(T) = f(T_{\text{cocycle}})$. Because T_{cocycle} only depends on Borda scores (by Lemma 1), then so does f. Since the space of purely cyclic tournaments is spanned by 3-cycles, we can write $T_{\text{cycle}} = \sum_{x,y,z} \lambda_{xyz} C_{xyz}$, where we may assume $\lambda_{xyz} \geq 0$ for all $x, y, z \in A$ (since we can replace negative values by observing that $C_{xyz} = -C_{zyx}$). By Lemma 4 and the quasi-Condorcet property, we have $f(C_{xyz}) = \mathscr{D}$. Thus, by reinforcement, $f(T_{\text{cycle}}) = \bigcap_{\lambda_{xyz}>0} f(C_{xyz}) = \mathscr{D}$. Thus, again by reinforcement, $f(T) = f(T_{\text{cycle}}) \cap f(T_{\text{cocycle}}) = f(T_{\text{cocycle}})$. □

Remark 1 The conclusion of Corollary 1 can also be proven using the axioms of neutrality, reinforcement, faithfulness (in addition), and cancellation (rather than quasi-Condorcet), by adapting the proofs of Debord (1992, Section 3). □

With the conclusion of Corollary 1 in place, the quasi-Condorcet property becomes a much stronger axiom: while previously it only implied that dummy alternatives (those that are majority-tied with every other alternatives) can be moved around freely, now we see that this is the case for all alternatives with Borda score 0.

Next we observe that f is equivalent to the Borda mean rule for the purely cocyclic tournaments S^x_y shown in Figure 2c. These tournaments have the useful property that the Borda score of all but two alternatives is zero, and, as we will see in the proof of Lemma 6, every purely cocyclic tournament can be decomposed into such tournaments.

Lemma 5 *If an SDF f satisfies neutrality, reinforcement, faithfulness, and the quasi-Condorcet property, then $f(S^x_y) = BM(S^x_y)$ for all $x, y \in A$.*

Proof By Corollary 1, any such f depends only on Borda scores. The weighted tournament S^x_y (see Section 6) is Borda-score equivalent to the weighted tournament $\widehat{S}^x_y$ given by

$$m_{xy} = 2, \quad m_{yx} = -2, \quad \text{and } m_{ab} = 0 \text{ otherwise.}$$

In $\widehat{S}^x_y$, all alternatives except x and y are dummies. By faithfulness, $f(\widehat{S}^x_y|_{\{x,y\}}) = \{x \succ y\}$. By the quasi-Condorcet property, $f(\widehat{S}^x_y) = \{\succcurlyeq \in \mathscr{D}(A) \colon x \succ y\} = BM(S^y_x)$. □

By decomposing purely cocyclic tournaments into tournaments of form S^x_y, we can pin down the output of f for all purely cocyclic tournaments.

Lemma 6 *Let T be a purely cocyclic weighted tournament. If f satisfies neutrality, reinforcement, and the quasi-Condorcet property, then $f(T) = BM(T)$.*

Proof By Lemma 1, T is difference generated by a function $\gamma \colon A \to \mathbb{R}$. We prove the statement by induction on the number of alternatives with non-zero Borda score. If there are no such alternatives, then every alternative has Borda score 0, and so by the quasi-Condorcet property, $f(T) = \mathscr{D} = BM(T)$.

Now assume that some alternative has non-zero Borda score, and so in particular γ is not constant. Let $\bar{x} \in \arg\max_{x \in A} \gamma(x)$ and $\underline{x} \in \arg\min_{x \in A} \gamma(x)$. We may assume without loss of generality that $\sum_{x \in A} \gamma(x) = 0$, since adding a constant function to γ does not change the weighted tournament it generates. This implies that $\gamma(\bar{x}) > 0$ and

$\gamma(\underline{x}) < 0$. Let $\delta = \min\{|\gamma(\bar{x})|, |\gamma(\underline{x})|\} > 0$. Let T' be the tournament that is difference generated by $\gamma' \colon A \to \mathbb{R}$ with $\gamma'(\bar{x}) = \gamma(\bar{x}) - \delta$, $\gamma'(\underline{x}) = \gamma(\underline{x}) + \delta$, and $\gamma'(x) = \gamma(x)$ for all $x \in A \setminus \{\bar{x}, \underline{x}\}$. Note that either $\bar{x}$ or $\underline{x}$ now has Borda score 0 in T', so $f(T') = BM(T')$ by induction. Lemma 5 implies that $f(S_{\underline{x}}^{\bar{x}}) = BM(S_{\underline{x}}^{\bar{x}}) = \{\succcurlyeq \in \mathscr{D} \colon \bar{x} \succ \underline{x}\}$. Also, by the definition of the Borda mean rule, $BM(T') \cap BM(S_{\underline{x}}^{\bar{x}}) \neq \varnothing$. From this and $T = T' + \delta S_{\underline{x}}^{\bar{x}}$ it follows from reinforcement that

$$f(T) = f(T') \cap f(S_{\underline{x}}^{\bar{x}}) = BM(T') \cap BM(S_{\underline{x}}^{\bar{x}}) = BM(T). \qquad \square$$

The outcome of f does not depend on the cyclic part of a weighted tournament as proven in Corollary 1. Lemma 6 shows that f is equal to the Borda mean rule for purely cocyclic tournaments. Together, these two statements imply that f is equal to the Borda mean rule for all weighted tournaments. This completes the proof of Theorem 1.

8 Independence of the Axioms

We show that all four axioms are indeed required for the characterization by giving an SDF that satisfies all but one of the axioms for each of the four axioms.

- *Neutrality:* Fix two alternatives $a, b \in A$ and define a skewed variant of the Borda mean rule by first doubling the weight of the edge between a and b and then calculating the outcome of the Borda mean rule.
- *Reinforcement:* Apply the sign-function to all majority margins (i.e., replace positive numbers by $+1$ and replace negative numbers by -1) before calculating the outcome of the Borda mean rule. This yields the *Copeland mean rule* that approves all alternatives with above-average Copeland score and disapproves those with below-average Copeland score.
- *Faithfulness:* Reverse the sign of all majority margins before calculating the outcome of the Borda mean rule. This yields the reverse Borda mean rule.
- *Quasi-Condorcet property:* Whenever all alternatives have Borda score zero (the weighted tournament is purely cyclic) then return all dichotomies. Otherwise, return the Borda winners, in the sense of returning $\{D_{\{x\}} \colon x \text{ is a Borda winner}\}$. By case analysis, one can check that this rule satisfies reinforcement. Notice that it does not satisfy reversal symmetry.

The last example implies that, in our main result, we cannot weaken the quasi-Condorcet property to cancellation.

9 Conclusions and Future Work

We have presented a characterization of the Borda mean rule as a social dichotomy function, showing that it fills the same space as does Kemeny's rule among social preference functions. It would be interesting to see other SDFs proposed and discussed in the literature; for now, the Borda mean rule seems like a very attractive example of an SDF.

Several questions remain for future work. Is there an alternative proof of our characterization that does not need linear algebra, such as in the proof of Hansson and Sahlquist (1976) for Borda's rule and of Debord (1992) for the k-Borda rule? (See Remark 1.) We can also ask whether the Borda mean rule can be characterized using different axioms. It seems particularly desirable to replace the quasi-Condorcet property with a more intuitive axiom. For example, does our result still hold if we were to replace the quasi-Condorcet property with cancellation and reversal symmetry? Or if we replace it with cancellation together with the requirement that Condorcet winners are always approved and Condorcet losers are always disapproved? These results are not ruled out by our examples in Section 8; to establish them, one would only need to reprove the conclusion of Lemma 5.

The Borda mean rule is particularly natural if voters' preferences are themselves dichotomous; in this setting, the Borda mean rule is often called the *mean rule* (Duddy et al., 2016). Our proof does not characterize the Borda mean rule if it is defined only over dichotomous preference profiles, because the quasi-Condorcet property is equivalent to cancellation on this domain. It would be interesting to have an axiomatic characterization of the mean rule using reinforcement. An axiomatic characterization using a different consistency notion is already known (Duddy et al., 2016).

We have noted that the Borda mean rule can also be seen as the 2-Kemeny rule. It seems plausible that our axioms in fact also characterize the k-Kemeny rule for each $k \geq 3$. However, it seems that different techniques (closer to the ones employed by Young and Levenglick (1978)) are necessary to show this.

Finally, is there a similar characterization of *scoring mean rules* based on other scoring rules, in the style of Young (1975)?

Acknowledgements This paper was partly written while the authors visited Carnegie Mellon University. We thank our host Ariel Procaccia, and COST Action IC1205 on Computational Social Choice for support. We thank Bill Zwicker and Piotr Skowron for helpful discussions.

References

J. P. Barthélemy and M. F. Janowitz. A formal theory of consensus. *SIAM Journal on Discrete Mathematics*, 4(3):305–322, 1991.

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

P. Y. Chebotarev and E. Shamis. Characterizations of scoring methods for preference aggregation. *Annals of Operations Research*, 80:299–332, 1998.

B. Debord. An axiomatic characterization of Borda's k-choice function. *Social Choice and Welfare*, 9(4):337–343, 1992.

C. Duddy, N. Houy, J. Lang, A. Piggins, and W. S. Zwicker. Social dichotomy functions. Working paper, 2014.

C. Duddy, A. Piggins, and W. S. Zwicker. Aggregation of binary evaluations: a Borda-like approach. *Social Choice and Welfare*, 46(2):301–333, 2016.

P. Faliszewski, P. Skowron, A. Slinko, and N. Talmon. Multiwinner voting: A new challenge for social choice theory. In U. Endriss, editor, *Trends in Computational Social Choice*, chapter 2. 2017. Forthcoming.

P. C. Fishburn. Condorcet social choice functions. *SIAM Journal on Applied Mathematics*, 33(3):469–489, 1977.

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

B. Hansson and H. Sahlquist. A proof technique for social choice with variable electorate. *Journal of Economic Theory*, 13(2):193–200, 1976.

J. G. Kemeny. Mathematics without numbers. *Daedalus*, 88:577–591, 1959.

M. Lackner and P. Skowron. Consistent approval-based multi-winner rules. *arXiv preprint arXiv:1704.02453*, 2017.

F. R. McMorris, H. M. Mulder, and R. C. Powers. The median function on median graphs and semilattices. *Discrete Applied Mathematics*, 101(1):221–230, 2000.

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

E. M. S. Niou. A note on Nanson's rule. *Public Choice*, 54:191–193, 1987.

P. Skowron, P. Faliszewski, and A. Slinko. Axiomatic characterization of committee scoring rules. In *Proceedings of the 6th International Workshop on Computational Social Choice (COMSOC)*, 2016.

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

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

H. P. Young. A note on preference aggregation. *Econometrica*, 42(6):1129–1131, 1974b.

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

H. P. Young. Optimal voting rules. *Journal of Economic Perspectives*, 9(1):51–64, 1995.

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

W. S. Zwicker. The voter's paradox, spin, and the Borda count. *Mathematical Social Sciences*, 22(3):187–227, 1991.

W. S. Zwicker. Cycles and intractability in social choice theory. In *Proceedings of the 6th International Workshop on Computational Social Choice (COMSOC)*, 2016.