

An Axiomatic Characterization of the Borda Mean Rule

Florian Brandl · Dominik Peters

July 14, 2018

Abstract A social dichotomy function maps a collection of weak orders to a set of dichotomous weak orders. Every dichotomous weak order partitions the set of alternatives into approved alternatives and disapproved alternatives. The Borda mean rule returns all dichotomous weak orders that approve all alternatives with above-average Borda score and disapprove 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. Our result holds for all domains of weak orders that are sufficiently rich, including the domain of all linear orders and the domain of all weak orders.

1 Introduction

The typical objective of social choice is to choose the best alternatives based on voters' preferences. Preferences are given as a weak order for each voter. Using this information, *social choice functions* choose a set of winning 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. [Duddy et al. \(2014\)](#) note that social choice functions, viewed as producing a partition into winning and non-winning alternatives, are not the right tool for this task. For example, consider a class of students that is to be divided into beginners and advanced learners based on how they are ranked by teachers. The goal is to form two groups of students such that

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

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

	1	2	3
a	a	a	a
b	b	b	b
c	d	e	c
d	e	c	d
e	c	d	e

Fig. 1: A preference profile for three voters and five alternatives. Each column depicts the preferences of one voter whose preference ranking is listed from top to bottom. The symmetric Borda scores of a and b are 12 and 6, respectively; each of c , d , and e has a symmetric Borda score of -6 . Hence, the Borda mean rule returns the dichotomous weak order that approves a and b and disapproves c , d , and e . Note that even though b is Pareto dominated by a , b is approved.

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 would uniquely choose the unanimously top-ranked student. Hence, the group of advanced learners would consist of only this one student; all other students would be 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, social dichotomy functions 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 switch place.

We consider a social dichotomy function that is based on *Borda scores*. A *preference profile* specifies for each voter a weak order on an alternative set A . Given a preference profile, the (symmetric) Borda score of an alternative a is obtained as follows: a gains a point for each voter i and each alternative b such that $a \succ_i b$; a loses a point if $b \succ_i a$; no points are assigned if $a \sim_i b$. Hence, the average Borda score is 0. *Borda's rule* is the social choice function that returns the alternatives with maximum Borda score. 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 defined by Duddy et al. (2014) and further discussed by Duddy et al. (2016, Section 5) and Zwicker (2016). Notice that the Borda mean rule satisfies reversal symmetry. An example illustrating the Borda mean rule is given in Figure 1.

Reversal symmetry (similarly defined) is also a natural property for *social preference functions*, which return a set of *linear orders* of the alternatives. 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. Young (1995) predicted that “the time will come when it is considered a standard tool for political and group decision making”. Given a preference profile, the rule assigns to each possible weak order \succsim a *Kemeny score*: the order gains a point for each voter i and each pair of alternatives $a, b \in A$ such that $a \succ b$ and $a \succ_i b$; the order loses a point if $b \succ_i a$; no points are assigned if $a \sim_i b$. Thus, the Kemeny score of \succsim indicates how much pairwise “agreement” there is between the output order \succsim and the input preferences. Kemeny’s rule returns the set of all *linear orders* with maximum Kemeny score.

Zwicker (2016) introduced the idea of using Kemeny scores to define aggregation rules for other output types. For example, suppose we maximize the Kemeny score over the family of relations $\{x\} \succ A \setminus \{x\}$ that have a unique most-preferred element and that are indifferent between all other alternatives. This yields precisely the relations whose most-preferred element is a winner of *Borda’s rule*. In his paper, Zwicker (2016) proposed the *k-Kemeny rule* which returns the *k*-chotomous weak order of highest Kemeny score; a weak order \succsim is called *k-chotomous* if its induced indifference relation \sim partitions A into at most k indifference classes. Thus, *k*-chotomous orders are ordered *k*-partitions. 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 proposals for voting rules; which of them should we use, and in which contexts? 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, such as the rule’s interpretation as a maximum likelihood estimate (Young, 1988)). 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). Thus, 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 to Young’s (1974a) characterization of Borda’s rule. In particular, we also use linear algebra and exploit the orthogonal decomposition

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

of weighted tournaments popularized in social choice theory by Zwicker (1991).² In contrast to Young and Levenglick (1978), 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.) The quasi-Condorcet property is stronger than the *cancellation axiom* and thus, in conjunction with reinforcement, implies that the output can only depend on the weighted majority relation (see Lemma 2). Neither of the four axioms in our collection can be dropped without the characterization result breaking down (see Section 8).

Our result applies to various possible input types. In particular, it applies when votes are given by linear orders, or when they are given by arbitrary weak orders. It also applies to j -chotomous weak orders whenever $j \geq 3$; thus, our result characterizes Zwicker’s (2016) $(j, 2)$ -Kemeny rule for each $j \geq 3$, which is the Borda mean rule as applied to profiles of j -chotomous weak orders. More generally, our proof works whenever the domain of allowed preference orders forms a *McGarvey domain*, that is, whenever every possible weighted majority tournament (with only even weights or only odd weights) can be induced by a profile using such orders. The domains of linear orders, of weak orders, and of j -chotomous orders ($j \geq 3$) are McGarvey domains. In contrast, the domain of dichotomous orders is not a McGarvey domain, and our proof does not apply to this domain.

2 Related Work

Duddy et al. (2016) study a setting in which every voter holds a binary evaluation of the alternatives or, equivalently, a dichotomous weak order. A *binary aggregation function* maps the voters’ binary evaluations to an ordered tripartition of approved, tied, and disapproved alternatives. Thus, the output of a binary aggregation function assigns to each alternative one of the values $+1$, 0 , or -1 . Duddy et al. (2016) propose the *mean rule*, which assigns $+1$ to all alternatives with above-average approval score, assigns 0 to alternatives whose approval score is exactly average, and assigns -1 to all alternatives with below-average approval score. They explain that the mean rule can be used in *judgement aggregation* for certain agendas, and connect the mean rule with the scoring rules for judgement aggregation introduced by Dietrich (2014). Further, Duddy

² Similar decompositions have been explored earlier by Harary (1959, 1969) and Croom (1978).

et al. (2016) prove that the mean rule is the only binary aggregation function satisfying axioms called 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). While the mean rule is very closely related to the Borda mean rule (when evaluated on profiles of dichotomous weak orders), the formal setting of Duddy et al. (2016) differs significantly from ours. In particular, our characterization result is logically independent from theirs, since our reinforcement axiom neither implies nor is implied by their consistency axiom. Further, our proof does not work for the case where voters are only allowed to submit dichotomous weak orders, and their proof does not apply to the kinds of input types that we study.

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., 2017a). 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 alternatives. Indeed, multiwinner rules typically do not satisfy reversal symmetry. 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, but simpler to prove due to the cardinality constraint on the set of winners. The k -Borda rule can be equivalently defined as the rule that returns the Kemeny score-optimal dichotomous orders with exactly k approved alternatives.

Recently, there has also been some discussion of multiwinner voting rules with a *variable* number of winners. The Borda mean rule is an example of such a rule. Kilgour (2016) reviews several such rules for the case of approval (dichotomous) preferences, and Faliszewski et al. (2017b) study their computational complexity.

A recent paper by Lang et al. (2016) proposes several schemes of rules that can be used to aggregate preferences into an arbitrary structure. For example, they propose a Kemeny scheme that can be used to find an aggregate ranking, or a committee, or a single winner, or an ordered committee, etc. Applying their Kemeny scheme to the output type of a dichotomy (i.e., an ordered partition into two pieces) yields the Borda mean rule. They also propose two other schemes that specialized to dichotomies, yielding different rules. The first is based on minimizing a Hamming distance and yields the *Copeland mean rule*, which approves an alternative whenever its Copeland score is above-average. The second generalizes the *Ranked Pairs rule* due to Tideman (1987), and yields a Ranked Pairs rule for dichotomies.

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 (cancellation). [Smith \(1973\)](#) independently found a characterization of scoring rules as social welfare functions; [Young \(1974b\)](#) gave an alternative proof of that result.

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 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](#)). [Nehring and Pivato \(2018\)](#) characterized the median procedure in judgement aggregation using reinforcement and *supermajority efficiency*, which is reminiscent of the quasi-Condorcet property. Their result also yields an alternative characterization of Kemeny’s rule.

3 Definitions

We use $\mathbb{N} = \{1, 2, \dots\}$ to denote an infinite set of potential voters. Let A be a finite set of alternatives, where $|A| = m$. The preferences of an agent $i \in \mathbb{N}$ are given by a binary relation $\succsim_i \subseteq A \times A$ which is complete and transitive; such a relation is called a *weak order*. We will write $a \succ_i b$ if $a \succsim_i b$ but $b \not\succsim_i a$, and $a \sim_i b$ if both $a \succsim_i b$ and $b \succsim_i a$. The *reverse* $-\succsim$ of a weak order \succsim is defined by $(a, b) \in -\succsim$ if and only if $(b, a) \in \succsim$. If σ is a permutation of A , we can naturally define the relation $\sigma(\succsim) = \{(\sigma(a), \sigma(b)) : (a, b) \in \succsim\}$, and extend this definition to sets and profiles of weak orders.

A weak order \succsim is called a *linear order* if it is antisymmetric, so that $a \sim b$ only if $a = b$. A weak order \succsim 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 $\succsim = A \times A$ is complete indifference. 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 $\mathcal{R}(A)$ for the set of all weak orders over A , $\mathcal{L}(A)$ for the set of linear orders over A , and $\mathcal{R}_2(A)$ for the set of dichotomous weak orders over A . When the set A is clear from the context, we write \mathcal{R} , \mathcal{L} , and \mathcal{R}_2 , respectively.

An *electorate* N is a finite and non-empty subset of \mathbb{N} . The set of all electorates is denoted by $\mathcal{F}(\mathbb{N})$. A (*preference*) *profile* $P \in \mathcal{R}^N$ on electorate N is a function assigning a weak order to each voter in N . The preferences of voter i in profile P are then denoted by \succsim_i .

Let $\mathcal{D} \subseteq \mathcal{R}$ be a domain of weak orders that the voters are allowed to submit. Typical choices for \mathcal{D} will be \mathcal{R} or \mathcal{L} . A *social dichotomy function* f is a map from the set of all profiles in \mathcal{D}^N for some $N \in \mathcal{F}(\mathbb{N})$ to non-empty subsets of \mathcal{R}_2 , so that $f(P) \subseteq \mathcal{R}_2$ for all profiles P .³ We denote by $-f$ the social dichotomy function that

³ One could require that the output of a social dichotomy function never contains the weak order denoting complete indifference. All of our arguments remain valid for this alternative definition.

returns the reverse of the weak orders returned by f ; thus, $(-f)(P) = -(f(P))$, for all profiles P .

Given a profile $P \in \mathcal{R}^N$ on the electorate N and two alternatives $a, b \in A$, let us write

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

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 $m_{ab} = -m_{ba}$).⁴ By $T(P) \in \mathbb{Z}^{m \times m}$ we denote the matrix of majority margins induced by P . This operation extends to sets of profiles by taking the union of the images for single profiles. We can interpret T as a *weighted tournament* 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} . For a permutation σ of A , we denote by $\sigma(T)$ the weighted tournament whose vertices are relabelled according to σ , so that the arc from $\sigma(a)$ to $\sigma(b)$ has weight m_{ab} for all $a, b \in A$.

The domains \mathcal{D} we consider are required to be sufficiently rich, a notion defined in terms of the weighted tournaments they induce. A domain \mathcal{D} is a *McGarvey domain* if every weighted tournament that can be induced by a preference profile of linear orders can be induced by a profile on the domain. Formally,

$$T(\{P \in \mathcal{L}^N : N \in \mathcal{F}(\mathbb{N})\}) \subseteq T(\{P \in \mathcal{D}^N : N \in \mathcal{F}(\mathbb{N})\}). \quad (\text{McGarvey domain})$$

It has been shown by [Debord \(1987\)](#) that the set of linear orders can induce exactly those integral weighted tournaments whose off-diagonal entries all have the same parity. Hence, a domain is a McGarvey domain if and only if it can induce all those weighted tournaments.⁵ Examples of McGarvey domains are \mathcal{L} , \mathcal{R} , and the set of all j -chotomous weak orders on A for $j \geq 3$. The set of dichotomous weak orders \mathcal{B}_2 and the set of weak orders that are single-peaked with respect to some fixed linear order are not McGarvey domains, since the majority relation is transitive for every profile from either of these domains.

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

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

the net weighted out-degree of a in the weighted tournament induced by P . This definition of Borda scores makes sense for profiles of arbitrary weak orders. For the case of linear orders, 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 $(m - 1, m - 2, \dots, 1, 0)$; 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

⁴ A matrix $M \in \mathbb{R}^{m \times m}$ is skew-symmetric if $M = -M^T$.

⁵ In fact, it is sufficient for our proof that the domain \mathcal{D} is rich enough to induce all tournaments with even entries.

of Borda scores.⁶ 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.

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. The Borda mean rule BM is the social dichotomy function with

$$BM(P) = \left\{ \succsim \in \mathcal{R}_2 : a \succ b \text{ for all } a, b \in A \text{ with } \beta(a) > 0 \text{ and } \beta(b) < 0 \right\}$$

for every preference profile $P \in \mathcal{D}^N$ on any electorate N . Thus, the Borda mean rule returns all dichotomous weak orders 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. This implies that the Borda mean rule returns exactly 2^k dichotomous weak orders, where k is the number of alternatives with exactly average Borda score. These alternatives appear in all possible ways in the upper and in the lower indifference class.

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 weak orders of maximum Kemeny score (as defined in Section 1):

$$BM(P) = \arg \max_{\succsim \in \mathcal{R}_2} \sum_{x \succ y} m_{xy}.$$

Hence, any $\succsim \in BM(P)$ maximizes the pairwise agreement with the voters' preferences among all dichotomous weak orders.

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.⁷ 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 $\beta(b) < 0$ if b is the Condorcet loser. We can argue alternatively that the Kemeny score of a dichotomy \succsim 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 2. In the tournament T , all alternatives have Borda score 0 (later we will say that T is *purely cyclic*), and so we have $BM(T) = \mathcal{R}_2$. The

⁶ Ours is not the only sensible extension of Borda scores to weak orders. An alternative choice for the score for $a \in A$ would be $\sum_{i \in N} (m - \text{rank}_i(a))$, where $\text{rank}_i(a) = k$ iff a is in the k th highest indifference class of \succsim_i . This alternative definition notably cannot be written only in terms of the majority margins m_{ab} . The two definitions are discussed by Chebotarev and Shamis (1998, Sec. 3.4) and Gärdenfors (1973, Sec. 7).

⁷ Adapting Fishburn's (1977) classification of the informational requirements of social choice functions to social dichotomy functions (whose input could contain weak orders), the Borda mean rule is a C2 rule.




Fig. 2: 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.

tournament T' is *purely cocyclic* with Borda scores 3, 0, -3 for x, y, z , respectively. Hence $BM(T') = \{\{x, y\} \succ \{z\}, \{x\} \succ \{y, z\}\}$. The tournament $T'' = T + T'$ has both a cyclic part and a cocyclic part. The Borda scores are 3, 0, -3 for x, y, z , which implies that $BM(T'') = \{\{x, y\} \succ \{z\}, \{x\} \succ \{y, z\}\}$.

5 Axioms

Duddy et al. (2014) argue that social dichotomy functions should satisfy *reversal symmetry*: if all voters reverse their preferences, then the approved set becomes the disapproved set and *vice versa*. Formally, a social dichotomy function satisfies reversal symmetry if

$$f(-P) = -f(P) \quad \text{for all } P \in \mathcal{D}^N \text{ and } N \in \mathcal{F}(\mathbb{N}). \quad (\text{Reversal symmetry})$$

While the Borda mean rule satisfies reversal symmetry, we do not impose this axiom for our characterization (it must therefore be implied by our other axioms). Instead, we use the same four axioms that also feature in Young and Levenglick’s (1978) characterization of Kemeny’s rule. First, we require social dichotomy functions to satisfy *neutrality*: renaming the alternatives in a preference profile leads to the same renaming in the output relations. Neutrality thus prescribes that a social dichotomy function 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, a social dichotomy function f satisfies neutrality if

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

When dealing with variable electorates, it seems reasonable to require that the rankings that are returned for two disjoint electorates should be precisely those that are returned when the electorates are merged. If the output for two electorates does not intersect, the condition says nothing. This is known as *reinforcement*. A social dichotomy function f satisfies reinforcement if

$$f(P) \cap f(P') \neq \emptyset \text{ implies } f(P) \cap f(P') = f(P \cup P') \quad \text{for all } P \in \mathcal{D}^N \text{ and } P' \in \mathcal{D}^{N'} \\ \text{with } N \cap N' = \emptyset. \quad (\text{Reinforcement})$$

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




Fig. 3: A tournament T with two dummy alternatives z and w . The arc from x to y has weight 1.

(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 also called consistency) for social welfare functions.

Next, 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 $m_{xy} = 0$ for all $y \in A$. The *quasi-Condorcet property* asserts that dummy alternatives can be placed arbitrarily in the output relation. Intuitively, this axiom claims that we have no relevant information about the appropriate placement of dummy alternatives in the output. For a set $B \subseteq A$, we say that two dichotomous weak orders \succsim and \succsim' agree on B if for all $x, y \in B$, we have $x \succsim y$ if and only if $x \succsim' y$. Then, a social dichotomy function f satisfies the quasi-Condorcet property if

$$\begin{aligned} \succsim \in f(P) \text{ if and only if } \succsim' \in f(P) \quad \text{whenever } \succsim \text{ and } \succsim' \text{ agree on } A \setminus D, \\ \text{where } D = \{x \in A : m_{xy} = 0 \text{ for all } y \in A\}, \text{ for all } P \in \mathcal{D}^N \text{ and } N \in \mathcal{F}(\mathbb{N}). \end{aligned}$$

(Quasi-Condorcet property)

Since it is the least intuitive of our axioms, we give an example that illustrates the consequences of the quasi-Condorcet property. Consider a preference profile P that induces the weighted tournament depicted in Figure 3. Note that z and w are dummies in T , i.e., $A \setminus B = \{z, w\}$. If $\succsim \in f(P)$ with $\{x, z, w\} \succsim \{y\}$, then the quasi-Condorcet property implies that the following three dichotomous weak orders are also contained in $f(P)$.

$$\{x\} \succsim' \{y, z, w\} \quad \{x, z\} \succsim' \{y, w\} \quad \{x, w\} \succsim' \{y, z\}$$

In Corollary 1, we show that in conjunction with the other axioms, the quasi-Condorcet property implies that not only dummy alternatives can be placed arbitrarily in the output relation, but in fact all alternatives with Borda score 0.

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

$$f(P) = \mathcal{R}_2 \quad \text{for all } P \in \mathcal{D}^N \text{ and } N \in \mathcal{F}(\mathbb{N}) \text{ with } m_{xy} = 0 \text{ for all } x, y \in A.$$

(Cancellation)

To see that the quasi-Condorcet property implies cancellation, observe that whenever all majority margins are zero in some profile P , every alternative is a dummy. Thus, every two dichotomous weak orders agree on the (empty) set of alternatives that are not dummies. Hence, the quasi-Condorcet property requires that $f(P) = \mathcal{R}_2$.

Our axioms so far do not rule out the trivial social dichotomy function which always returns all dichotomies. We need an axiom that prescribes some degree of correlation of the voters' preferences with the aggregated dichotomies. An arguably minimal axiom of this nature is *faithfulness*, which requires that whenever the electorate consists of just one voter, the aggregated preference should not contradict that voter's preferences. Formally, a social dichotomy function f satisfies faithfulness if

$$\succ_i \subseteq \succ \quad \text{for all } \succ \in f(P), P \in \mathcal{D}^{\{i\}}, \text{ and } i \in \mathbb{N}. \quad (\text{Faithfulness})$$

Thus, if P is a profile with only one voter who strictly prefers a to b , then $f(P)$ cannot contain any dichotomy which approves b but disapproves a . As an example, if P is a profile containing only the vote of voter i , and \succ_i is the linear order $a \succ_i b \succ_i c \succ_i d$, then faithfulness requires that $f(P)$ contains only dichotomous orders that approve an initial segment (an upper contour set) of \succ_i . Thus $\{a, b\} \succ \{c, d\}$ may be a member of $f(P)$, but $\{a, d\} \succ \{b, c\}$ cannot be a member of $f(P)$.

Faithfulness is a weak axiom. It is even weaker than the following (weak) Pareto axiom: whenever a profile $P \in \mathcal{D}^N$, $N \in \mathcal{F}(\mathbb{N})$, is such that $a \succ_i b$ for all $i \in N$, then $a \succ b$ for all $\succ \in f(P)$. In the proof of our characterization, the faithfulness axiom is only invoked once, to rule out the trivial social dichotomy function and the reverse Borda mean rule $-BM$.

6 The Linear Algebra of Weighted Tournaments

Let us define a few special weighted tournaments that will be useful later, see Figure 4 for drawings. Given three distinct alternatives $x, y, z \in A$, we write C_{xyz} for the weighted tournament with $m_{xy} = m_{yz} = m_{zx} = 1$, $m_{yx} = m_{zy} = m_{xz} = -1$, 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 $S_y^x = D_{\{x\}} + D_{A \setminus \{y\}}$ which consists of a single “top” alternative x , a single “bottom” alternative y , and all other alternatives in between; x defeats y by a majority margin of 2.

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 the vector space of rational-valued skew-symmetric $m \times m$ matrices (and, equivalently, of weighted tournaments). Note that the dimension $\dim V$ of V is $\binom{m}{2}$. Identifying a




Fig. 4: Some types of tournaments. Unlabelled arcs have weight 1.

skew-symmetric matrix with a vector in $\mathbb{Q}^{m(m-1)/2}$, this vector space can be endowed with the usual inner product.⁸ The *cycle space* $V_{\text{cycle}} = \langle C_{xyz} : x, y, z \in A \text{ distinct} \rangle$ is defined as the span (the set of all linear combinations) of all 3-cycles C_{xyz} , or equivalently, the span of all simple cycles $C_{v_1 v_2 \dots v_k}$, since $C_{v_1 v_2 \dots v_k} = \sum_{i=2}^{k-1} C_{v_1 v_i v_{i+1}}$. The *cocycle space* $V_{\text{cocycle}} = \langle D_{\{x\}} : x \in A \rangle$ is defined as the span of all tournaments $D_{\{x\}}$, or equivalently, the span of all D_X with $X \subseteq A$, since $D_X = \sum_{x \in X} D_{\{x\}}$. As we now show, V can be decomposed as the sum $V = V_{\text{cycle}} + V_{\text{cocycle}}$, and these two subspaces are orthogonal, i.e., for all $T_{\text{cycle}} \in V_{\text{cycle}}$ and $T_{\text{cocycle}} \in V_{\text{cocycle}}$, we have $T_{\text{cycle}} \cdot T_{\text{cocycle}} = 0$. In the following proof, we will also establish that $\dim V_{\text{cycle}} = \binom{m}{2} - (m-1)$ and $\dim V_{\text{cocycle}} = m-1$.

Proposition 1 *The subspaces V_{cycle} and V_{cocycle} are orthogonal and $V = V_{\text{cycle}} + V_{\text{cocycle}}$.*

Proof Since the inner product is bilinear, it suffices to check orthogonality on spanning sets of V_{cycle} and V_{cocycle} . So consider some C_{xyz} and some $D_{\{v\}}$ with $x, y, z, v \in A$. If $v \notin \{x, y, z\}$, then $C_{xyz} \cdot D_{\{v\}} = 0$. Otherwise, without loss of generality, $v = x$ and so $C_{xyz} \cdot D_{\{x\}} = 1 - 1 = 0$.

To see that $V = V_{\text{cycle}} + V_{\text{cocycle}}$, we use a dimension analysis. Write $A = \{1 \dots m\}$. We claim that $D_{\{1\}}, \dots, D_{\{m-1\}}$ are linearly independent, so that $\dim V_{\text{cocycle}} \geq m-1$. Let $\lambda_1, \dots, \lambda_{m-1} \in \mathbb{Q}$ be such that $\sum_{i=1}^{m-1} \lambda_i D_{\{i\}} = 0$. This implies that $\lambda_i = \lambda_j$ for all $i, j \in \{1, \dots, m-1\}$ to ensure that $m_{ij} = 0$. Moreover, $0 = m_{1m} = \lambda_1$. Hence $\lambda_i = 0$ for all $i \in \{1, \dots, m-1\}$.

We prove that $\dim V_{\text{cycle}} \geq \binom{m}{2} - (m-1) = \binom{m-1}{2}$ by showing that $\{C_{ijm} : i, j \in [m-1], i < j\}$ is a set of linearly independent vectors. For $i, j \in [m-1], i < j$, let $\lambda_{ij} \in \mathbb{Q}$ be such that $\sum_{i=1}^{m-1} \sum_{j=i+1}^{m-1} \lambda_{ij} C_{ijm} = 0$. Since $0 = m_{ij} = \lambda_{ij}$, it follows that $\lambda_{ij} = 0$ for all $i, j \in [m-1], i < j$. \square

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 cyclic* if $T = T_{\text{cycle}}$ so that $T_{\text{cocycle}} = 0$, and we say that T is *purely cocyclic* if $T = T_{\text{cocycle}}$ so that $T_{\text{cycle}} = 0$. Of the examples in Figure 4, C_{xyz} is purely cyclic, and D_X and S_y^x are purely cocyclic. In Figure 2, the tournament $T'' = T + T'$ can be decomposed into its cyclic component T and its cocyclic component T' .

⁸ The inner product of two vectors $u, v \in \mathbb{Q}^n$ is defined to be $u \cdot v = \sum_{i=1}^n u_i v_i$.

If T is a purely cyclic tournament, then every alternative has Borda score 0. Hence the Borda scores induced by a given tournament T are the same as the Borda scores induced by its cocyclic component T_{cocycle} . In fact, knowledge of the Borda scores is enough to construct T_{cocycle} , as is shown by the following 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 $\gamma: A \rightarrow \mathbb{R}$ such that $m_{xy} = \gamma(x) - \gamma(y)$ for all $x, y \in A$. In fact, if T is purely cocyclic, then it is difference generated by $\gamma(x) := \beta(x)/m$, i.e., by Borda scores, suitably rescaled.*

The function γ is unique up to adding a constant. We can normalize γ by requiring that $\sum_{x \in A} \gamma(x) = 0$, in which case we then have $\gamma(x) = \beta(x)/m$ for all $x \in A$.

For example, the tournament S_y^x is difference generated with $\gamma(x) = 1$, $\gamma(y) = -1$, and $\gamma(z) = 0$ for $z \in A \setminus \{x, y\}$. The tournament $D_{\{x\}}$ is difference generated with $\gamma(x) = (m-1)/m$ and $\gamma(z) = -1/m$ for all $z \in A \setminus \{x\}$.

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 $\alpha_1 T_1 + \alpha_2 T_2$ is difference generated by $\alpha_1 \gamma_1 + \alpha_2 \gamma_2$. 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 V_{cycle} , since V_{cycle} is spanned by 3-cycles. Thus, by Proposition 1, T is purely cocyclic, i.e., $T \in V_{\text{cocycle}}$.

For the second statement, recall that $D_{\{x\}}$ is difference generated by $\gamma(x) = \beta(x)/m$, since $\beta(x) = m-1$, and $\beta(z) = -1$ for all $z \neq x$, so that $\gamma(x) - \gamma(z) = \beta(x)/m - \beta(z)/m = 1$. Similar to above, it then follows that linear combinations of these tournaments are generated by $\gamma(x) = \beta(x)/m$. \square

With this result, it is easy to find the decomposition of a given tournament. First construct the cocyclic component T_{cocycle} using the Borda scores, and then obtain $T_{\text{cycle}} = T - T_{\text{cocycle}}$.

7 Characterization

We are now ready to state and prove our main result. For the remainder of this section, let $\mathcal{D} \subseteq \mathcal{R}$ be some fixed McGarvey domain.

Theorem 1 *A social dichotomy function 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. Let f be a social dichotomy function satisfying neutrality, reinforcement, faithfulness, and the quasi-Condorcet property. We show that f is the Borda mean rule. The proof is split up into five lemmas; in the statement of each lemma, we mention which axioms are used in their proofs. Whenever cancellation suffices as a weakening of the quasi-Condorcet property, we note this as well.

Our first two lemmas are similar to lemmas in Young’s (1974a) characterization of Borda’s rule. Their conclusions do not depend on the type of output of f , and in particular also hold for social choice functions and social preference functions (for the appropriate definition of cancellation). Young (1974a) operates in the context of profiles of linear orders, but we can adapt the arguments to work for any McGarvey domain.

Lemma 2 (Young, 1974a, Young and Levenglick, 1978) *If a social dichotomy function 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 two profiles that induce the same majority margins, i.e., $T(P_1) = T(P_2)$. Assume first that P_1 and P_2 are defined on disjoint electorates. Let Q and Q' be two profiles such that the electorates of all four profiles are pairwise disjoint and $T(Q) = T(P_1)$ and $T(Q') = -T(P_1 \cup Q)$. The profile Q' exists, since all majority margins in $P_1 \cup Q$ are even (since $T(P_1 \cup Q) = T(P_1) + T(Q) = 2 \cdot T(P_1)$) and \mathcal{D} is a McGarvey domain by assumption. Observe that $T(P_1 \cup Q \cup Q') = T(P_2 \cup Q \cup Q') = 0$. Thus, it follows from cancellation that $f(P_1 \cup Q \cup Q') = f(P_2 \cup Q \cup Q') = \mathcal{R}_2$. By reinforcement, we get

$$\begin{aligned} f(P_1) &= f(P_1) \cap \mathcal{R}_2 = f(P_1) \cap f(P_2 \cup Q \cup Q') \\ &= f(P_1 \cup P_2 \cup Q \cup Q') \\ &= f(P_1 \cup Q \cup Q') \cap f(P_2) = \mathcal{R}_2 \cap f(P_2) = f(P_2). \end{aligned}$$

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 $f(P_1) = f(P_3) = f(P_2)$. \square

Since the quasi-Condorcet property implies cancellation, this shows that f only depends on the majority margins induced by a preference profile. In particular, this implies that f is *anonymous*, i.e., the outcome is invariant under renaming the voters.

We next observe that reinforcement and anonymity imply that f is *homogeneous*. Suppose $P \in \mathcal{D}^N$ and $P' \in \mathcal{D}^{N'}$ are profiles on disjoint electorates, and such that there is a bijection $\phi : N \rightarrow N'$ with $P'(\phi(i)) = P(i)$ for all $i \in N$. Thus, P and P' are “copies”. By anonymity, $f(P) = f(P')$, and thus by reinforcement $f(P \cup P') = f(P) \cap f(P') = f(P)$. By induction, if nP is a profile consisting of n copies of P , then $f(nP) = f(P)$.

In light of Lemma 2, we can view f as a function that takes as input a weighted tournament that is induced by some preference profile $P \in \mathcal{D}^N$ with $N \in \mathcal{F}(\mathbb{N})$. The next lemma will show that f admits a unique extension to all rational weighted

tournaments. We say that a function $\hat{f} : V \rightarrow \mathcal{R}_2$ extends f if $\hat{f}(T(P)) = f(P)$ for all profiles $P \in \mathcal{D}^N$ with $N \in \mathcal{F}(\mathbb{N})$. Our axioms naturally extend to functions defined on tournaments. In particular, \hat{f} satisfies neutrality if $\hat{f}(\sigma(T)) = \sigma(f(T))$ for all $T \in V$ and $\sigma \in \Pi(A)$; it satisfies reinforcement if $f(T + T') = f(T) \cap f(T')$ for all $T, T' \in V$ with $f(T) \cap f(T') \neq \emptyset$; and it satisfies the quasi-Condorcet property if for all $T \in V$, we have $f(T) = f(T) \setminus B$ where $A \setminus B = \{x \in A : m_{xy} = 0 \text{ for all } y \in A\}$.

Lemma 3 (Young, 1974a) *Suppose f satisfies reinforcement and cancellation. Then there exists a unique $\hat{f} : V \rightarrow \mathcal{R}_2$ which extends f and satisfies reinforcement. If f further satisfies neutrality or the quasi-Condorcet property, then so does \hat{f} .*

Proof Let $T \in V$ be a rational tournament. Take $n \in \mathbb{N}$ to be an integer such that nT is a tournament whose weights are even integers, i.e., such that $\frac{1}{2}nT \in \mathbb{Z}^{m \times m}$. Because \mathcal{D} is a McGarvey domain, there is a preferences profile $P \in \mathcal{D}^N$ for some $N \in \mathcal{F}(\mathbb{N})$ which induces nT . We define $\hat{f}(T) := f(P)$. Then $\hat{f} : V \rightarrow \mathcal{R}_2$ is well-defined by Lemma 2 and because f is homogeneous.

Next, we show that \hat{f} satisfies reinforcement. From the definition of \hat{f} , it follows immediately that $\hat{f}(nT) = \hat{f}(T)$ for all $T \in V$ and $n \in \mathbb{N}$. Let $T, T' \in V$ be such that $f(T) \cap f(T') \neq \emptyset$. Take $n \in \mathbb{N}$ with $\frac{1}{2}nT, \frac{1}{2}nT' \in \mathbb{Z}^{m \times m}$, and let P, P' be profiles on disjoint electorates which induce nT and nT' , respectively. Note that $P \cup P'$ induces $nT + nT'$. Then

$$\begin{aligned} \hat{f}(T + T') &= \hat{f}(n(T + T')) = \hat{f}(nT + nT') \\ &= f(P \cup P') = f(P) \cap f(P') = \hat{f}(T) \cap \hat{f}(T'). \end{aligned}$$

since $f(P) \cap f(P') \neq \emptyset$ and f satisfies reinforcement.

The extension \hat{f} is unique: let $\tilde{f} : V \rightarrow \mathcal{R}_2$ be any function which satisfies reinforcement and extends f so that $\tilde{f}(T(P)) = f(P)$ for all profiles P . Then $\tilde{f} = \hat{f}$: let $T \in V$, and take $n \in \mathbb{N}$ with $\frac{1}{2}nT \in \mathbb{Z}^{m \times m}$ and a profile P inducing nT . Then

$$\tilde{f}(T) = \tilde{f}(nT) = f(P) = \hat{f}(T),$$

where the first equality follows because $\tilde{f}(T + T) = \tilde{f}(T)$ by reinforcement.

The fact that \hat{f} satisfies neutrality, faithfulness, and the quasi-Condorcet property for all $T \in V$ follows easily from the definition of \hat{f} and our assumption that f satisfies these axioms. \square

Using Lemma 3, we will from now on take f to be defined on all of V , the set of rational-valued skew-symmetric matrices (equivalently, weighted tournaments).

We next show that f only depends on the Borda scores of the alternatives. To achieve this result, we will show that f is trivial on purely cyclic tournaments, in the sense of returning all dichotomies; this implies 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), f can then only depend on the Borda scores. As a first step, we show that f is trivial for the building blocks C_{xyz} of the cycle space, by an argument using neutrality.

Lemma 4 *If f satisfies neutrality, reinforcement, and cancellation, then $f(C_{xyz}) = \mathcal{R}_2$ for all distinct $x, y, z \in A$.*

Proof Let $x, y, z \in A$ be three distinct alternatives, write $C = C_{xyz}$, and take any $\succ \in f(C)$. Since \succ is dichotomous, it cannot be that x, y, z are all in distinct indifference classes, and so \succ is indifferent between at least two of them. Hence, we may assume without loss of generality that $x \sim y$. Let $\sigma = (xy)$ be the permutation that transposes x and y and leaves other alternatives fixed. We have $\succ = \sigma(\succ) \in \sigma(f(P)) = f(\sigma(C))$, where the last equality follows from neutrality. Next, note that $\sigma(C) = C_{zyx}$ which is the reverse of C . Hence, $C + \sigma(C) = 0$. Since $\succ \in f(C) \cap f(\sigma(C))$, by reinforcement,

$$f(C) \cap f(\hat{\sigma}(C)) = f(C + \hat{\sigma}(C)) = f(0) = \mathcal{R}_2,$$

where the last step follows from cancellation. Hence $f(C) = \mathcal{R}_2$. \square

Next, we lift the result for 3-cycles C_{xyz} to apply to all tournaments in V_{cycle} .

Corollary 1 *If a social dichotomy function f satisfies neutrality, reinforcement, and cancellation, 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

$$T_{\text{cycle}} = \sum_{\substack{x, y, z \in A \\ \text{distinct}}} \lambda_{xyz} C_{xyz},$$

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, we have $f(C_{xyz}) = \mathcal{R}_2$ for all distinct $x, y, z \in A$. Thus, by reinforcement,

$$f(T_{\text{cycle}}) = \bigcap_{\lambda_{xyz} > 0} f(C_{xyz}) = \mathcal{R}_2.$$

Thus, again by reinforcement,

$$f(T) = f(T_{\text{cycle}}) \cap f(T_{\text{cocycle}}) = f(T_{\text{cocycle}}). \quad \square$$

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 alternative) can be moved around freely, now we see that this is the case for all alternatives with Borda score 0. Note that dummy alternatives have Borda score 0. Also, for every weighted tournament T , there is a weighted tournament T' such that each alternative has the same Borda score in T and T' , and every alternative with Borda score 0 is a dummy alternative in T' . The tournament T' can be constructed by subtracting a purely cyclic tournament from T that contains exactly those edges in T that are adjacent to some alternative with Borda score 0. Since T and T' have the same Borda scores and f only depends on Borda scores, f must treat dummy alternatives and alternatives with Borda score 0 identically.




Fig. 5: The tournaments used in the proof of Lemma 5. Each alternative has the same Borda score in either tournament: $\beta(x) = m$, $\beta(y) = -m$, $\beta(z_i) = 0$ for all $i \in \{1, \dots, m-2\}$.

Next we observe that f is equivalent to the Borda mean rule for the purely cocyclic tournaments S_y^x shown in Figure 4c. 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. The proof of Lemma 5 is the only place where we use the full force of the quasi-Condorcet property rather than only cancellation.

Lemma 5 *If a social dichotomy function f satisfies neutrality, reinforcement, and the quasi-Condorcet property, then $f(S_y^x) \in \{BM(S_y^x), -BM(S_y^x), \mathcal{R}_2\}$ for all $x, y \in A$.*

Proof Recall that $|A| = m$. Let G_y^x be the weighted tournament given by

$$m_{xy} = m, \quad m_{yx} = -m, \quad \text{and } m_{zv} = 0 \text{ for } z, v \in A \setminus \{x, y\}.$$

For each alternative, its Borda score in S_y^x is equal to its Borda score in G_y^x (see Figure 5). By Corollary 1, $f(G_y^x) = f(S_y^x)$, and so it suffices to show that $f(G_y^x) \in \{BM(G_y^x), -BM(G_y^x), \mathcal{R}_2\}$. Note that $BM(G_y^x) = BM(S_y^x) = \{\succ \in \mathcal{R}_2 : x \succ y\}$.

Suppose first that there is $\succ \in f(G_y^x)$ such that $x \sim y$. Neutrality applied to the permutation $\sigma = (xy)$ implies that $\succ \in f(\sigma(G_y^x)) = f(G_x^y)$. Hence, $f(G_y^x) \cap f(G_x^y) \neq \emptyset$, so reinforcement and cancellation imply that

$$f(G_y^x) \cap f(G_x^y) = f(G_y^x + G_x^y) = f(0) = \mathcal{R}_2.$$

Hence $f(G_y^x) = \mathcal{R}_2$, and we are done.

In G_y^x , all alternatives except x and y are dummies. Thus, if there is $\succ \in f(G_y^x)$ such that $x \succ y$, then by the quasi-Condorcet property, $BM(G_y^x) = \{\succ \in \mathcal{R}_2 : x \succ y\} \subseteq f(G_y^x)$. Similarly, if there is $\succ \in f(G_y^x)$ such that $y \succ x$, then $-BM(G_y^x) \subseteq f(G_y^x)$.

Suppose next that there are $\succ, \succ' \in f(G_y^x)$ such that $x \succ y$ and $y \succ' x$. Then $BM(G_y^x) \cup -BM(G_y^x) \subseteq f(G_y^x)$. Neutrality applied to the permutation $\sigma = (xy)$ implies

$$-BM(G_y^x) \cup BM(G_y^x) = \sigma(BM(G_y^x)) \cup \sigma(-BM(G_y^x)) \subseteq f(\sigma(G_y^x)) = f(G_x^y).$$

Thus, $f(G_y^x) \cap f(G_x^y) \neq \emptyset$, so reinforcement and cancellation imply that

$$f(G_y^x) \cap f(G_x^y) = f(G_y^x + G_x^y) = f(0) = \mathcal{R}_2.$$

Hence $f(G_y^x) = \mathcal{R}_2$, and we are done.

In the remaining case, either $x \succ y$ for all $\succ \in f(G_y^x)$ and hence $f(G_y^x) = BM(G_y^x)$; or $y \succ x$ for all $\succ \in f(G_y^x)$ and hence $f(G_y^x) = -BM(G_y^x)$. \square

Denote by $TRIV$ the social dichotomy function that always returns all dichotomous weak orders, i.e., $TRIV(T) = \mathcal{R}_2$ for all $T \in V$. In combination with our axioms, either one of the three cases characterized in Lemma 5 pins down f on all tournaments.

Lemma 6 *If f satisfies neutrality, reinforcement, and the quasi-Condorcet property, then $f \in \{BM, -BM, TRIV\}$.*

Proof Take any distinct $x, y \in A$ and choose $g \in \{BM, -BM, TRIV\}$ so that $f(S_y^x) = g(S_y^x)$. By Lemma 5, the choice of g is well-defined, unique, and, by neutrality of f , independent of x, y .

We will show that $f = g$. By Corollary 1, it suffices to show that f is equal to g for purely cocyclic tournaments. Let T be a purely cocyclic tournament. We prove the statement by induction on the number of alternatives with non-zero Borda score in T .

If there are no such alternatives (every alternative has Borda score 0) then $T = 0$. By cancellation (implied by the quasi-Condorcet property), we have $f(T) = \mathcal{R}_2 = g(T)$.

Now assume that not all alternatives have Borda score 0 in T , and that $f(T') = g(T')$ for all $T' \in V$ in which fewer alternatives have non-zero Borda score than in T . By Lemma 1, T is difference generated by the function $\gamma: A \rightarrow \mathbb{R}$, $\gamma(a) = \beta(a)/m$, i.e., by Borda scores, suitably rescaled. Since not all Borda scores are 0, γ is not constant. Let $\bar{x} \in \arg \max_{x \in A} \gamma(x)$ and $\underline{x} \in \arg \min_{x \in A} \gamma(x)$. Since $\sum_{x \in A} \gamma(x) = 0$, we have $\gamma(\bar{x}) > 0$ 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 $\gamma': A \rightarrow \mathbb{R}$ with

$$\gamma'(\bar{x}) = \gamma(\bar{x}) - \delta, \gamma'(\underline{x}) = \gamma(\underline{x}) + \delta, \text{ and } \gamma'(x) = \gamma(x) \text{ for all } x \in A \setminus \{\bar{x}, \underline{x}\}.$$

By the choice of γ , we have that $\gamma'(a) = \beta'(a)/m$ for all $a \in A$, where $\beta'(a)$ is the Borda score of a in T' . Note that either \bar{x} or \underline{x} now has Borda score 0 in T' . Thus, there are fewer alternatives with non-zero Borda score in T' than in T , and so $f(T') = g(T')$ by assumption. Also, $f(S_{\bar{x}}^{\bar{x}}) = g(S_{\bar{x}}^{\bar{x}})$, by definition of g . For each of the three possible choices of g , we have $g(T') \cap g(S_{\bar{x}}^{\bar{x}}) \neq \emptyset$. From this and $T = T' + \delta S_{\bar{x}}^{\bar{x}}$, it follows by reinforcement of f and g that

$$f(T) = f(T') \cap f(S_{\bar{x}}^{\bar{x}}) = g(T') \cap g(S_{\bar{x}}^{\bar{x}}) = g(T). \quad \square$$

Neither $-BM$ nor $TRIV$ satisfies faithfulness. To see this, consider $P \in \mathcal{D}^{(i)}$, $i \in \mathbb{N}$, such that \succ_i is not complete indifference, which exists since \mathcal{D} is a McGarvey domain. Let $x, y \in A$ be such that x is most-preferred and y is least-preferred for i ; that is, $x \succ_i z \succ_i y$ for all $z \in A$. Then $x \succ_i y$, and it is easy to check that $\beta(x) > 0 > \beta(y)$. Faithfulness requires that if $\succ \in f(P)$, then $\succ_i \subseteq \succ$; in particular, $x \succ y$. However, by their definitions, both $TRIV(P)$ and $-BM(P)$ contain dichotomies with $y \succ x$.

In combination with Lemma 6, 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 a social dichotomy function 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 all dichotomous weak orders returned by the Borda mean rule ($-BM$) or always return all dichotomies (*TRIV*). Lemma 6 shows that these are in fact the only other social dichotomy functions that satisfy the remaining axioms.
- *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 $\{D_{\{x\}} : x \text{ is a Borda winner}\}$. 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 social dichotomy functions discussed in the literature. For example, one might consider mean rules based on other positional scoring rules. Lang et al. (2016) propose a version of the Ranked Pairs rule that returns dichotomies, and Kilgour (2016) proposes some multiwinner voting rules with committees of variable size, which can be interpreted as social dichotomy functions. For now, the Borda mean rule seems like a very attractive example of an social dichotomy function.

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? 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 the conjunction of 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. In a different formal setting, an axiomatic characterization using another 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 Klaus Nehring, Piotr Skowron, and Bill Zwicker 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.
- F. H. Croom. *Basic Concepts of Algebraic Topology*. Springer-Verlag, 1978.
- B. Debord. Caractérisation des matrices des préférences nettes et méthodes d'agrégation associées. *Mathématiques et sciences humaines*, 97:5–17, 1987.
- B. Debord. An axiomatic characterization of Borda's k -choice function. *Social Choice and Welfare*, 9(4):337–343, 1992.
- Franz Dietrich. Scoring rules for judgment aggregation. *Social Choice and Welfare*, 42(4):873–911, 2014.
- 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. 2017a.
- P. Faliszewski, A. Slinko, and N. Talmon. The complexity of multiwinner voting rules with variable number of winners. Technical report, <https://arxiv.org/abs/1711.06641>, 2017b.
- P. C. Fishburn. Condorcet social choice functions. *SIAM Journal on Applied Mathematics*, 33(3):469–489, 1977.
- P. Gärdenfors. Positionalist voting functions. *Theory and Decision*, 4(1):1–24, 1973.
- B. Hansson and H. Sahlquist. A proof technique for social choice with variable electorate. *Journal of Economic Theory*, 13(2):193–200, 1976.
- F. Harary. Graph theory and electric networks. *IRE Transactions on Circuit Theory*, 6(5):95–109, 1959.
- F. Harary. *Graph Theory*. Addison-Wesley, 1969.
- J. G. Kemeny. Mathematics without numbers. *Daedalus*, 88:577–591, 1959.
- D. M. Kilgour. Approval elections with a variable number of winners. *Theory and Decision*, 81(2):199–211, 2016.
- M. Lackner and P. Skowron. Consistent approval-based multi-winner rules. Technical report, <https://arxiv.org/abs/1704.02453>, 2017.
- J. Lang, J. Monnot, A. Slinko, and W. S. Zwicker. Beyond electing and ranking: Collective dominating chains, dominating subsets and dichotomies. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 24–32. IFAAMAS, 2016.
- 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.
- K. Nehring and M. Pivato. The median rule in judgement aggregation. 2018. Working paper.
- 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.
- T. N. Tideman. Independence of clones as a criterion for voting rules. *Social Choice and Welfare*, 4(3):185–206, 1987.
- 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. Condorcet’s theory of voting. *The American Political Science Review*, 82(4):1231–1244, 1988.

- 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. Technical report, <https://arxiv.org/abs/1608.03999>, 2016.