Barcelona School of Economics

WORKING PAPERS

Opinion Aggregation: Borda and Condorcet Revisited

BSE Working Paper 1262
June 2021 (Revised: September 2022)

Salvador Barberà, Walter Bossert

Opinion aggregation: Borda and Condorcet revisited*

Salvador Barberà
MOVE, Universitat Autònoma de Barcelona
and Barcelona School of Economics
Facultat d'Economia i Empresa
08193 Bellaterra (Barcelona)
Spain
salvador.barbera@uab.cat

Walter Bossert
Centre Interuniversitaire de Recherche en
Economie Quantitative (CIREQ)
University of Montreal
P.O. Box 6128, Station Downtown
Montreal QC H3C 3J7
Canada
walter.bossert.wb@gmail.com

This version: September 12, 2022

Abstract. In a world admitting a fixed finite set of alternatives, an opinion is an ordered pair of alternatives. Such a pair expresses the idea that one alternative is superior to another in some sense, and an opinion aggregator assigns an aggregate relation on the set of alternatives to every possible state of opinion. Our primary motivation is to extend the standard model of social choice theory to a more general one in which no specific reference to agents generating or holding opinions is needed. Although our analysis has some bearing on those cases where opinions reflect the goodness relations of agents in a society, it is not limited to them. In addition to that interpretation, opinions can also be used to represent other forms of comparative assessments emerging from different sources. The main results of the paper provide characterizations of suitably formulated versions of the Borda rule and Condorcet's majority rule, essential aggregation methods that remain well-defined in our larger context. *Journal of Economic Literature* Classification Nos.: D71, D72.

* Salvador Barberà acknowledges financial support from the Spanish Agencia Estatal de Investigación (AEI), through the Severo Ochoa Programme for Centres of Excellence in R&D (Barcelona School of Economics CEX2019-000915-S), from the Spanish Ministry of Economy and Competitiveness and FEDER, through grants ECO2017-83534-P and PID2020-116771GB-I00, and from the Generalitat de Catalunya, through grant 2017SGR0711. We thank the editor Pierpaolo Battigalli, an associate editor, three referees, Susumu Cato, Bhaskar Dutta, Anke Gerber, Sean Horan, Matthew Jackson, Bernardo Moreno, Hervé Moulin, and Yves Sprumont for comments and suggestions.

1 Introduction

In a world admitting a fixed finite set of alternatives, an opinion is an ordered pair of alternatives. Such a pair expresses the idea that one alternative is superior to another in some sense, and an opinion aggregator assigns an aggregate relation on the set of alternatives to every possible state of opinions. Our primary motivation is to extend the standard model of social choice theory to a more general one in which no specific reference to agents generating or holding opinions is needed. Many aggregators have been defined where opinions reflect the goodness relations of agents. These are often assumed to satisfy additional properties, such as completeness or transitivity, although such assumptions can be relaxed to some extent with no harm. But not always. Consider, for example, the classical Borda (1781) rule. This method assigns scores to the alternatives that are then used to establish an aggregate ranking, and the way to reach these scores admits several equivalent descriptions when opinions are derived from complete, transitive, and antisymmetric individual relations. The score of an alternative can be derived by adding the weights assigned to the positions it occupies in the goodness relations of the agents, or the number of times that this alternative defeats another, or the size of the differences between the number of its pairwise wins and losses. We shall argue that only the third formulation is adequate to use in our general framework, and we will provide a characterization of it. In that case, it is fortunate that the spirit of a classical rule can be fully maintained in the larger framework: Borda's method is well defined even if the origin of opinions is not attached to specific individuals; moreover, it still satisfies properties that make it attractive and that we shall spell out when characterizing it. Other rules are not so robust and cannot be properly extended unless individuals are well identified; this is the case, for example, for the remaining scoring rules (Young, 1975) that assign different weights to alternatives according to their positions in the individual rankings.

Clearly, Borda's is not the only rule that can be defined without reference to specific individuals. In particular, the simple majority rule advocated by Condorcet (1785) can also be well defined in our context, along with others. And these two competing forms of aggregating individual goodness relations also share another important characteristic that we highlight. This is the observation that they both are based on the differences between the number of wins and the number of losses that each alternative experiences when confronted with others in pairwise contests. The notion that this difference is what really matters may, on the surface, appear to be be debatable, but it is natural and unifying—and it precludes the difficulties alluded to in the previous paragraph that emerge when only wins are taken into consideration in defining the Borda rule. We prove that all difference-based rules share a cancellation property that is, conceptually, familiar from earlier contributions such as that of Young (1974). Our main results provide characterizations of the methods advocated by Borda and by Condorcet, and they necessarily differ from the existing ones because one can no longer use axioms that rely on the identification of those individuals who hold the opinions to be aggregated. While the axiomatic treatment of majority is quite close to that provided by May (1952) and reestablished by Sen (1970, Theorem 5*1), the characterization of Borda's rule is based on a significant departure from previous approaches that are valid on smaller domains. It is also worth remarking that, although usually expressed in the context where individual relations are complete and transitive, known characterizations, like those by Young (1974) or Nitzan and Rubinstein (1981), could be naturally extended to avoid the use of these regularity conditions. Indeed, such properties are not needed when applying our rules to the larger framework, and several interesting conclusions that we think have not been previously emphasized are derived from this fact. Specifically, we show that several important rules that only require the expression of partial information rather than of full profiles on the part of the agents can be rephrased in our language that involves opinions. Moreover, we prove that the plurality rule is equivalent to the use of Borda's rule or the Condorcet method of majority decision when translated into the opinion-aggregation setting—and the same observation applies to approval voting and some of its variants.

In summary, we propose a new and larger framework to investigate aggregation rules, we offer a characterization of the important class of difference-based rules, we observe that the two most classical aggregation methods, which clearly belong to this large class, can be naturally extended to and characterized in the framework, and we show that both are equivalent to the plurality rule and to approval voting in the case where the individuals who express opinions are identified but submit ballots that contain less information than full profiles of individual goodness relations.

Although our analysis has some bearing on those cases where opinions reflect the goodness relations of agents in a society, it is not limited to them. In addition to this interpretation, opinions can also be used to represent other forms of comparative assessments emerging from different sources. For example, the opinion (x, y) may indicate that a book x is better than another book y, an exam x handed in by one student deserves a better grade than another exam y submitted by another student, a sports team x defeats another team y, a statement x is more likely to be true than another statement y. Another interesting application emerges when it comes to observed choices. If, in a repeated series of observations, an object is chosen over another, it may very well be the case that several of these acts of choice are performed by one and the same agent—and this is more difficult to model if only a single goodness relation is to be assigned to each agent; note that it is perfectly possible that the same agent chooses x over y in one instance and y over x in another. Furthermore, there seems to be evidence that some cognitive processes operate by aggregating different (and potentially conflicting) impulses; see, for instance, Jackson and Yariv (2015, p. 151) for a discussion. Therefore, the opinion-based framework may serve as an adequate model of decision-making that can be observed in the human brain. Clearly, this list is by no means intended to be complete but serves as an indication of the broad applicability of the notion of opinions.

A state of opinion is a function that identifies the number of times each opinion appears in the state under consideration. The special case of the empty state of opinion—a function that assigns the number zero to each possible opinion—is included as a possibility. We do not assume anything about who expresses the opinions but note that we can accommodate cases in which a single agent may express any finite number of opinions or none at all. For instance, the standard case familiar from social choice theory where each agent can express one and only one goodness relation is covered, and so is the possibility of each agent merely passing judgment on a single pair of alternatives. But, as we already hinted at, the notion of an opinion is more general because several opinions regarding the same pair of alternatives can be expressed by the same agent. A more detailed discussion of the significance of states of opinion and their relationships to more conventional modeling devices is provided once the requisite definitions have been introduced formally.

An opinion aggregator is a function that assigns a relation on the set of alternatives to any possible state of opinion, with the convention that the universal equal-goodness relation is associated with the empty state of opinion. The latter requirement is very natural: if no opinion is expressed whatsoever, there are no grounds for favoring any alternative over any other. The distinctive feature of our approach is that, unlike in the standard theories of social choice or judgment aggregation, we do not impose initial assumptions about linkages between the opinions that constitute a state. Nor do we require them, or the relations into which they are aggregated, to possess a priori any form of coherence. In particular, this implies that properties such as completeness or acyclicity are not required. See also Fishburn (1984) and Chebotarev (1994), for instance, who dispense with the completeness assumption but retain the basic framework phrased in terms of individual agents.

In the spirit of Borda and Condorcet, we characterize the Borda opinion aggregator and the majority opinion aggregator as illustrations of how our framework of opinion aggregation can be put into practice. The ideas underlying Borda's method diverge considerably from the majoritarian ideas expressed by Condorcet. Both had supporters as well as detractors among their contemporaries. For instance, Morales (1797) was a committed advocate of Borda, whereas Daunou (1803) favored Condorcet's ideas and was rather critical of Borda's method. Daunou also formulated his own proposal of a voting method, which is analyzed and characterized by Barberà, Bossert, and Suzumura (2021). Herrero and Villar (2021) examine a rule that represents a combination of the proposals by Borda and Condorcet.

As elaborated on earlier, one motivation for the use of the Borda method and the majority rule is that they are widely applicable and can be defined in our general context without the need for any further assumptions regarding the structure of the states of opinion. A second important reason for our focus is that, in natural contexts where different well-known rules can be defined, the two aggregators coincide with these rules. Consider, for example, the case of plurality voting, and the kind of special states of opinion that the balloting under this rule can generate. By declaring their top alternative, agents express their opinions between that alternative and any other. We show that if attention is restricted to such states of opinion, the aggregate ranking according to the plurality rule coincides with the aggregate ranking according to the Borda opinion aggregator and with that according to the majority opinion aggregator. As another example, consider approval voting where each agent submits a set of alternatives—namely, the set of approved-of alternatives. Approval voting ranks the alternatives on the basis of their approval scores—that is, the number of times an alternative is approved of by an agent. A natural interpretation is that such a ballot is represented by a state of opinions that consists of all opinions (x, y), where x is approved of and y is not. Again, the ranking of the alternatives according to the approval-voting rule coincides with the ranking according to the Borda opinion aggregator and with that according to the majority opinion aggregator.

A third reason that motivates us is of a historical character. Both Borda's (1781) own description of his proposed voting method and its defense by Morales are very inspiring.

In fact, the notion of an opinion appears in Morales's (1797) Memoir on the Calculus of Opinion—a contribution that was well-respected and appreciated not only by Borda himself but several other contemporaries. The term opinion is used by Morales in the same way as in our model but, in addition, Morales strongly defends this notion as a natural unit of measurement, and he advocates the view that all opinions should count equally, no matter who holds them. Borda's presentation also refers to what he calls merit as a magnitude admitting a definite and fixed value, unique up to affine scaling. In a similar manner, he insists on the fact that the difference of merit between two successive alternatives in a voter's relation remains the same independently of the position they occupy. Of course, we do not maintain this significance of differences as a starting point, but our axiomatic characterization can be interpreted as a foundation for its ex-post defense.

We hope that the observations presented here will further underline the importance and significance of the methods advocated by Borda and by Condorcet. As alluded to earlier, these two methods can be defined within our framework, whereas many other rules only make sense in the more limited context of classical social choice theory based on profiles of individual orderings. That the Borda method does not require stringent assumptions on individual relations is already pointed out by Young (1974, p. 51) who states that his characterization of the classical Borda social choice function can be rephrased in a way that applies to general profiles of individual relations that do not have to be orderings. Likewise, Nitzan and Rubinstein (1981) characterize the Borda rule without assuming individual goodness relations to be transitive. However, our model of collective choice in terms of opinion aggregation has, to the best of our knowledge, not been examined in the earlier literature. That Condorcet's method of majority decision does not rely on any properties of the individual inputs such as completeness or acyclicity is well-known and follows immediately from inspecting the requisite criterion to rank the alternatives. Each of these well-established methods of collective choice allows for a characterization in our model of opinion aggregation. Both axiomatizations differ from analogous results in the traditional setting in that no reference to individual agents is made here. There is, however, a conceptual difference between our two characterizations. The majoritarian criterion only uses information on x and y to rank any two alternatives x and y, to the exclusion of all other alternatives. This feature is associated with a rather forceful neutrality property than cannot but be present in some guise or other. As a consequence, our characterization of the majority opinion aggregator shares the central features of May's (1952) result—other than being able to avoid all references to specific individuals. In contrast, our axiomatization of the Borda opinion aggregator is, as we believe, considerably more novel and utilizes the rich framework of opinion aggregation to its full advantage. See, for example, Young (1974), Hansson and Sahlquist (1976), Nitzan and Rubinstein (1981), and Mihara (2017) for characterizations of the Borda rule in the traditional setting.

We conclude this introduction with an informal description of the properties that we employ in this paper. Our first formal result identifies all opinion aggregators that are based on net wins—that is, on the difference between the number of wins and the number of losses experienced by an alternative in a state of opinion. This is achieved by imposing a pairwise cancellation axiom that is familiar from the earlier literature; for instance, Young (1974) employs such a property in his characterization of the Borda social choice function.

Intuitively, it requires that opposing opinions cancel each other out when determining the aggregate relation.

Our first step in identifying a system of axioms that characterizes the Borda opinion aggregator consists of strengthening the pairwise opinion cancellation alluded to in the previous paragraph. In analogy to pairwise opinion cancellation, the property deals with how a favorable opinion for an alternative x over a second alternative y is to be traded off against a favorable opinion for y over a third alternative z. If x and z are one and the same alternative, the pairwise variant of the property results. If these two alternatives differ, however, the situation is more subtle: although a win is compensated by a loss as far as alternative y is concerned, it must be taken into consideration that removing a favorable opinion for x over y leaves x with one win less, and an analogous observation applies to the removal of a loss of z against y. Thus, these changes must be accounted for in the version of the axiom that also applies to distinct alternatives x and z. We stress that one of the possible motivations of the axiom goes back to Morales's (1797) view that all opinions should count equally.

The second property we employ in axiomatizing the Borda opinion aggregator is an opinion-monotonicity requirement. Loosely speaking, the axiom requires that if the situation of some undominated alternatives improves by adding a favorable opinion for each of them, then they improve in the aggregate ranking while remaining equally good among each other. A dual requirement is imposed on the response of the opinion aggregator to the deterioration of alternatives that are not better than any of the others. This is an intuitively appealing property that ensures a positive response to a change in a state of opinion that unambiguously favors some of the alternatives. We note that the two axioms just described exploit the structure of the entities being studied here—the states of opinion.

As mentioned earlier, we assume throughout that an opinion aggregator assigns the universal equal-goodness relation to the empty state of opinion. The universal equal-goodness relation is complete and transitive but, a priori, none of these coherence properties need to be satisfied by other aggregate relations that emerge by applying the rule. Because the Borda opinion aggregator generates only complete and transitive aggregate relations, these requirements must either be imposed or implied by other axioms. It turns out that completeness is a consequence of other conditions but transitivity is not. Therefore, the final property we impose is that all relations generated by an opinion aggregator be transitive.

Turning to our characterization of the majority opinion aggregator, we note first that this method always generates complete aggregate relations. Because this property is not implied by our remaining axioms, it needs to be required explicitly for all states of opinion other than the empty state. The other two axioms employed parallel those of May (1952). The first is opinion neutrality, the opinion-aggregation variant of May's (1952) well-known strengthening of Arrow's (1951; 1963) independence of irrelevant alternatives, and the second is a suitably formulated version of opinion responsiveness.

2 Opinions and opinion aggregators

Consider a finite and non-empty set X of alternatives. For distinct $x, y \in X$, an opinion on x and y is an ordered pair $(x, y) \in D = X^2 \setminus \{(z, z) \mid z \in X\}$. The first element in the pair is the winner and the second is the loser in that opinion. A natural interpretation is that an opinion (x, y) represents the view that x is better than y according to an agent. However, we do not assign an individual label to an opinion, which means that if two opinions (x, y) and (z, w) are observed, they may reflect the views of a single agent or of two distinct agents. It is irrelevant which of these two options applies—all opinions are treated impartially, no matter who holds them.

The absence of individuals assigned to observed opinions may be seen as a disadvantage of our approach. A possible argument in this vein is that this feature prevents us from disqualifying or discounting the opinions of agents whose goodness relations are inconsistent in the sense that their opinions display violations of transitivity, for instance. Our response to this concern is that we are somewhat reluctant to advocate an asymmetric treatment of the agents, depending on a notion of perceived inconsistency or irrationality. For example, violations of transitivity are well-documented in the psychology literature because they may arise from thresholds of perception; see, for instance, Armstrong (1939), Luce (1956), Luce and Raiffa (1957), and Fishburn (1970) for detailed discussions. Thus, phenomena of this nature do not necessarily indicate a lack of consistency or coherence in the views of an agent. In addition, the observation that seemingly contradictory opinions may originate from separate brain impulses provides a further argument against the dismissal of opinions on the grounds of what may, at first sight, appear to constitute inconsistent behavior.

A state of opinion is a function $o \colon D \to \mathbb{N}_0$, where $o(x, y)$ is the number of times the opinion $(x, y) \in D$ appears in the state of opinion o. The cardinality $|o|$ of a state of opinion $o \in \mathcal{O}$ is the total number of opinions that appear in o, that is, $|o| = \sum_{(x,y)\in D} o(x, y)$. The empty state of opinion $o_\emptyset \in \mathcal{O}$ is the function defined by $o_\emptyset(x, y) = 0$ for all $(x, y) \in D$ and its cardinality is equal to zero.

States of opinion are very general in that they can capture scenarios that are difficult to model in the more traditional individual-profile framework. For instance, suppose we can repeatedly observe individual choices. This induces a state of opinion o in which each choice of an alternative x in the presence of another alternative y is represented by an opinion (x, y) in o. It is, of course, possible that the same pair (x, y) emerges multiple times—and that the pair (y, x) is also observed in other acts of choice. The formulation in terms of opinions has an advantage over the profile model because of its increased flexibility. It is perfectly possible to accommodate seemingly contradictory choices such that the same agent chooses x over y in one observation and y over x in another. While it is possible to deal with such a situation in a profile setting by declaring x and y equally good according to the goodness relation of the individual in question (thus making both alternatives possible choices), this move requires the imposition of an additional assumption that is not generated by the observed behavior itself. In contrast, a state of opinion can easily accommodate the simultaneous presence of (multiple copies of) both (x, y) and (y, x) because no explicit reference to individuals is required. In line with Morales (1797), all opinions are treated equally, no matter who holds them—and this is the fundamental equal-treatment assumption that underlies this approach. We note that this form of impartiality is reflected in the definition of the aggregation procedure based on the Borda rule and in one of our fundamental axioms that we focus on in this paper.

As already alluded to in the introduction, there is some evidence that cognitive processes aggregate different (and potentially conflicting) impulses in the brain; see, for instance, Jackson and Yariv (2015) for details. In an opinion-based framework, this phenomenon is easily accommodated because nothing prevents us from interpreting the joint presence of the pairs (x, y) and (y, x) as being two conflicting impulses of this nature experienced by one and the same individual. This is not possible in the context of individual goodness relations: by definition of a betterness relation, an alternative x is better than an alternative y if x is at least as good as y but not the reverse. As an immediate consequence, it is not possible for x to be better than y and, at the same time, for y to be better than x from the viewpoint of the same individual.

We now turn to the aggregation of opinions into an aggregate relation. Let $\mathcal{B}$ be the set of all binary relations on X. An opinion aggregator is a function $f\colon \mathcal{O} \to \mathcal{B}$ such that $f(o_\emptyset) = X^2$. Thus, $f(o)$ is the aggregate relation assigned to the state of opinion $o \in \mathcal{O}$, and the universal equal-goodness relation X^2 on X is associated with the empty state of opinion $o_\emptyset$. We use $P(f(o))$ and $I(f(o))$ to denote the asymmetric and symmetric parts of the relation $f(o)$.

Two opinion aggregators are of particular interest in this paper. The Borda opinion aggregator f^B ranks alternatives by comparing the differences between the numbers of wins and losses in a state of opinion. For a state of opinion $o \in \mathcal{O}$ and an alternative $x \in X$, the Borda opinion score $b(x; o)$ is defined as

$$b(x; o) = \sum_{z \in X \setminus \{x\}} [o(x, z) - o(z, x)].$$

The Borda opinion aggregator is obtained by letting

$$(x, y) \in f^B(o) \;\Leftrightarrow\; b(x; o) \geq b(y; o)$$

for all $o \in \mathcal{O}$ and for all $x, y \in X$.

In the context of goodness aggregation, the Borda rule is a special case of a scoring rule (Young, 1975). As already observed by Young (1974), the Borda rule sets itself apart from other scoring rules because it is well-defined even if individual goodness relations are neither complete nor transitive. Closely related to the scoring rules are the positional voting functions analyzed in Gärdenfors (1973); see also Bossert and Suzumura (2020).

Our definition of the Borda opinion aggregator is analogous to the definition of the traditional Borda (1781) rule for more structured voting rules that depend on individual antisymmetric orderings as their informational basis. In this setting, the Borda rule can equivalently be described by simply comparing the number of wins rather than the difference between the number of wins and the number of losses. This is the case because completeness and antisymmetry together imply immediately that the number of losses is uniquely determined once the number of wins is known. However, once completeness is no longer assumed (as is the case in our framework), these two variants can differ dramatically. To illustrate, consider a situation with a set of three alternatives given by $X = \{x, y, z\}$. Suppose a state of opinion $o \in \mathcal{O}$ is given by

$$o(x, z) = o(y, z) = 1 \;\text{ and }\; o(z, y) = 2$$

and the remaining values of o are equal to zero. According to our definition, the Borda opinion score for each alternative is given by the requisite difference between the number of wins and the number of losses so that, in particular, $b(x;o) = 1 - 0 = 1$ and $b(y;o) = 1 - 2 = -1$. Thus, according to the Borda opinion aggregator, $(x,y) \in P(f^B(o))$—that is, x is better than y. In contrast, if we were to use the number of wins rather than the difference between the numbers of wins and losses, x and y would be ranked as equally good for the state of opinion o because both record a number of wins equal to one. This latter relative ranking seems difficult to reconcile with the spirit of the Borda method because y is beaten two times, whereas x does not suffer any such losses. The use of the Borda opinion aggregator as defined above avoids conclusions of this nature.

The majority opinion aggregator f^M is defined by letting

$$(x,y) \in f^M(o) \;\Leftrightarrow\; o(x,y) \geq o(y,x)$$

for all $o \in \mathcal{O}$ and for all $x, y \in X$.

As illustrated above, it is of crucial importance to employ the difference between the number of wins and the number of losses to determine the Borda score of an alternative; restricting attention to the number of wins leads to conclusions that go squarely against the spirit of Borda's method. We note that the majority opinion aggregator is also based on the difference between the number of wins and the number of losses; this follows immediately from rewriting its definition as

$$(x,y) \in f^M(o) \;\Leftrightarrow\; o(x,y) - o(y,x) \geq 0$$

for all $o \in \mathcal{O}$ and for all $x, y \in X$. We consider this exclusive reliance on the differences between wins and losses to be the cornerstone of our approach to opinion aggregation and provide a foundation of it by means of an intuitively appealing equivalence result the scope of which extends well beyond the methods of Borda and of Condorcet. In particular, the following theorem illustrates how an adaptation of a well-established cancellation condition (see, for instance, Young, 1974) can be used to characterize a general class of opinion aggregators that are based on the differences between wins and losses—or, in other words, on net wins.

An opinion aggregator f is based on net wins if, for all $o, o' \in \mathcal{O}$,

$$o'(v,w) - o'(w,v) = o(v,w) - o(w,v) \;\text{ for all } (v,w) \in D$$

implies

$$f(o') = f(o).$$

The class of opinion aggregators that are based on net wins are characterized by means of a pairwise cancellation axiom. The condition requires that if there is an alternative y who loses to another alternative x according to one opinion in a state of opinion o, and wins against x according to another opinion in o, the two opinions (x,y) and (y,x) cancel each other out when determining the aggregate relation for the state of opinion under consideration. That is, if the numbers of instances of (x,y) and of (y,x) are reduced by one and no other changes occur, the aggregate relations before and after the elimination of these two opinions are identical. This is plausible because only x and y are affected by the change and the win of x over y is canceled out by the loss of x against y. An analogous property phrased in the context of profile aggregation is well-established in the literature; it is, for instance, employed by Young (1974) in his characterization of the Borda rule.

Pairwise opinion cancellation. For all $o, o' \in \mathcal{O}$ and for all $(x, y) \in D$, if $o'(x, y) = o(x, y) - 1$ and $o'(y, x) = o(y, x) - 1$ and $o'(v, w) = o(v, w)$ for all $(v, w) \in D \setminus \{(x, y), (y, x)\}$, then

$$f(o') = f(o).$$

We obtain

Theorem 1 *An opinion aggregator f satisfies pairwise opinion cancellation if and only if f is based on net wins.*

Proof. Suppose first that f is based on net wins. Let $o, o' \in \mathcal{O}$ and $(x, y) \in D$ be such that $o'(x, y) = o(x, y) - 1$ and $o'(y, x) = o(y, x) - 1$ and $o'(v, w) = o(v, w)$ for all $(v, w) \in D \setminus \{(x, y), (y, x)\}$. It follows immediately that

$$o'(x, y) - o'(y, x) = o(x, y) - 1 - o(y, x) + 1 = o(x, y) - o(y, x)$$

and

$$o'(v, w) - o'(w, v) = o(v, w) - o(w, v) \text{ for all } (v, w) \in D \setminus \{(x, y), (y, x)\}$$

and, because f is based on net wins, we obtain $f(o') = f(o)$ so that pairwise opinion cancellation is satisfied.

Conversely, suppose that f satisfies pairwise opinion cancellation and that

$$o'(v, w) - o'(w, v) = o(v, w) - o(w, v) \text{ for all } (v, w) \in D. \tag{1}$$

Define the state of opinion $\overline{o} \in \mathcal{O}$ by letting

$$\overline{o}(v, w) = o(v, w) - o(w, v) \text{ and } \overline{o}(w, v) = o(w, v) - o(w, v) = 0$$

for all $(v, w) \in D$ such that $o(v, w) - o(w, v) \geq 0$. By (1), it follows that

$$\overline{o}(v, w) = o'(v, w) - o'(w, v) \text{ and } \overline{o}(w, v) = o'(w, v) - o'(w, v) = 0$$

for all $(v, w) \in D$ such that $o'(v, w) - o'(w, v) = o(v, w) - o(w, v) \geq 0$. Applying pairwise opinion cancellation $o(w, v)$ times for each of the requisite pairs (w, v) in the case of o, and $o'(w, v)$ times for each of the requisite pairs (w, v) in the case of o', it follows that

$$f(\overline{o}) = f(o) \text{ and } f(\overline{o}) = f(o')$$

so that

$$f(o') = f(o)$$

and it follows that f is based on net wins. ■

Another special case of apparent interest is obtained if the criterion employed in the comparison of the members of X is additively separable, without going all the way to the fully additive structure of the Borda opinion aggregator f^B. However, the objective of obtaining a sound characterization of such a class is likely to prove elusive, owing to the discrete nature of the variables involved; see the seminal contribution of Kraft, Pratt, and Seidenberg (1959). Their observations illustrate why properties such as separability or independence that can successfully be applied in a continuum fail to generate additively separable structures if the domain under consideration is a finite set. The difficulties encountered in the present setting are of an analogous nature, which is why we do not provide an axiomatization in this spirit.

3 Borda and Condorcet

This section is devoted to axiomatizations of the Borda opinion aggregator and of the majority opinion aggregator. These results are of interest in their own right and serve to illustrate how our proposed framework of opinion aggregation can successfully be employed.

We begin with the Borda aggregator. Although the Borda method has been characterized in numerous previous contributions, our axiomatization distinguishes itself in that it relies on axioms that are specifically relevant in the context of states of opinion.

The first property that we use in our characterization of the Borda opinion aggregator is largely motivated by Morales's (1797) maxim that all opinions count equally, no matter who holds them. To begin with, this view is reflected in the axiom of pairwise opinion cancellation defined in the previous section—which reappears as a special case of the stronger axiom of opinion cancellation. But there also is a second (more subtle) possibility of expressing the notion of equally-valued opinions. Consider a situation in which an alternative y loses to an alternative x in one opinion and y wins against an alternative z in another opinion but, unlike in the pairwise case, the alternatives x and z differ. Reducing the numbers of occurrences of the two opinions (x, y) and (y, z) by one leaves x with one win less and the same number of losses as before and, analogously, z ends up with the number of losses reduced by one and an unchanged number of wins. This situation can be dealt with by adding one instance of the opinion (x, z). In this manner, the numbers of wins and losses for x and z are preserved so that equality of the two resulting aggregate relations can again be required. The second type of cancellation property just described represents a distinguishing feature of the Borda opinion aggregator, whereas the pairwise version is of more universal appeal.

Opinion cancellation. For all $o, o' \in \mathcal{O}$ and for all $x, y, z \in X$, if $o'(x, y) = o(x, y) - 1$ and $o'(y, z) = o(y, z) - 1$ and $o'(v, w) = o(v, w)$ for all $(v, w) \in D \setminus \{(x, y), (y, x), (x, z)\}$ and $(o'(x, z) = o(x, z) + 1$ if $x \neq z)$, then

$$f(o') = f(o).$$

If $x = z$, the axiom of pairwise opinion cancellation familiar from the previous section results. We note that the formulation of opinion cancellation relies on the notion of a state of opinion. In the traditional profile-based setting, it may not be possible to treat instances of betterness in isolation, as we are able to do if the entities under consideration are opinions.

The second property employed in our characterization of the Borda opinion aggregator is a natural monotonicity requirement the scope of which is restricted to an important subclass of states of opinion. The members of this class distinguish themselves from other states of opinion in that they allow us to partition the set of alternatives into three natural groups. A state of opinion $o \in \mathcal{O}$ is trichotomous if each alternative in X is either (i) a winner in at least one opinion in o and not a loser in any opinion in o; or (ii) not a winner in any opinion in o and a loser in at least one opinion in o; or (iii) not a winner and not a loser in any opinion in o. Thus, for a trichotomous state of opinion o, the set of alternatives can be partitioned into three groups, namely, (i) those who sometimes win and never lose; (ii) those who never win and sometimes lose; and (iii) those who never win and never lose. To provide a precise definition of the set $\mathcal{T} \subseteq \mathcal{O}$ of trichotomous states of opinion, let $o \in \mathcal{O}$ be a state of opinion. Define $W(o) \subseteq X$ as the set of all alternatives $x \in X$ such that

$$\{z \in X \setminus \{x\} \mid o(x,z) > 0\} \neq \emptyset \ \text{ and } \ \{z \in X \setminus \{x\} \mid o(z,x) > 0\} = \emptyset.$$

Analogously, let $L(o) \subseteq X$ be the set of all alternatives $x \in X$ such that

$$\{z \in X \setminus \{x\} \mid o(x,z) > 0\} = \emptyset \ \text{ and } \ \{z \in X \setminus \{x\} \mid o(z,x) > 0\} \neq \emptyset.$$

Finally, let $U(o) \subseteq X$ be the set of all alternatives $x \in X$ such that

$$\{z \in X \setminus \{x\} \mid o(x,z) > 0\} = \{z \in X \setminus \{x\} \mid o(z,x) > 0\} = \emptyset.$$

The set $\mathcal{T}$ of trichotomous states of opinion is defined as the set of all states of opinion $o \in \mathcal{O}$ such that

$$X = W(o) \cup L(o) \cup U(o).$$

Therefore, in a trichotomous state of opinion, there are no alternatives that sometimes win and sometimes lose. This implies that, for such a state of opinion $o \in \mathcal{T}$, the set of alternatives can be partitioned unambiguously into the set of winners given by $W(o)$, the set of losers given by $L(o)$, and the set of unclassified alternatives given by $U(o)$. The empty state of opinion $o_\emptyset$ is a trichotomous state of opinion. This is the case because, by definition, $W(o_\emptyset) = L(o_\emptyset) = \emptyset$ and $U(o_\emptyset) = X$, which immediately implies

$$X = \emptyset \cup \emptyset \cup X = W(o_\emptyset) \cup L(o_\emptyset) \cup U(o_\emptyset).$$

To illustrate the first part of the following axiom, consider any trichotomous state of opinion o and the aggregate relation $f(o)$ that corresponds to o according to an opinion aggregator f. Furthermore, consider any set of alternatives $Y \subseteq W(o) \cup U(o)$ such that all members of Y are best in $W(o) \cup U(o)$ according to the relation $f(o)$. We note that elements of $U(o)$ need not be included in Y; in fact, unless $o = o_\emptyset$, the entire set Y will be composed of alternatives in $W(o)$ when we apply the axiom in our proof. The reason why we include elements of $U(o)$ in the set of possible alternatives in Y is to cover the case in which $o = o_\emptyset$ because, for the (trichotomous) empty state of opinion $o_\emptyset$, the set $W(o_\emptyset)$ is empty. Now define a new trichotomous state of opinion o' by adding one win for each of the members of Y. Intuitively, the members of Y are alternatives that are best according to $f(o)$ among those in $W(o) \cup U(o)$, and the change from the state of opinion o to the state of opinion o' further improves their standing by adding another win for each of them.

We require such a move to have three consequences that we consider very intuitive. First, note that the members of Y are equally good according to $f(o)$ because all of them are best elements in $W(o) \cup U(o)$ according to this relation. Adding one win for each of them is assumed not to change their relative positions and, therefore, it seems highly plausible to require that all alternatives in Y be equally good according to $f(o')$ as well. Second, an additional win is assigned to the members of Y and no one else receives any additional wins in the new state of opinion. Thus, it appears only natural to require that the position of those in Y relative to those who are winners or unclassified alternatives in the new state of opinion o' has improved. Third, the winners and the unclassified alternatives in the new state of opinion o' who are not in Y did not experience any change in their numbers of wins or losses. We therefore require that their relative rankings in the aggregate relation remain unchanged as a consequence of the move from o to o'.

The second part of the monotonicity axiom is dual. Instead of adding a win to a set of alternatives that are best elements in $W(o) \cup U(o)$ according to the aggregate relation $f(o)$ that corresponds to a trichotomous state of opinion, we add one loss to each member of a subset of the worst elements in $L(o) \cup U(o)$ and require the three consequences that are analogous to those detailed above.

Opinion monotonicity. (a) For all $o, o' \in \mathcal{T}$ and for all $Y \subseteq W(o) \cup U(o)$ such that $(y, x) \in f(o)$ for all $y \in Y$ and for all $x \in W(o) \cup U(o)$, if for all $y \in Y$, there exists $z_y \in L(o) \cup U(o)$ with $o'(y, z_y) = o(y, z_y) + 1$ and $o'(v, w) = o(v, w)$ for all $(v, w) \in D \setminus \{(y, z_y) \mid y \in Y\}$, then

$$(y, y') \in f(o') \;\text{ for all } y, y' \in Y$$

and

$$(y, x) \in P(f(o')) \;\text{ for all } y \in Y \text{ and for all } x \in W(o') \cup U(o') \setminus Y$$

and

$$(x, x') \in f(o') \;\Leftrightarrow\; (x, x') \in f(o) \;\text{ for all } x, x' \in W(o') \cup U(o') \setminus Y.$$

(b) For all $o, o' \in \mathcal{T}$ and for all $Z \subseteq L(o) \cup U(o)$ such that $(x, z) \in f(o)$ for all $z \in Z$ and for all $x \in L(o) \cup U(o)$, if for all $z \in Z$, there exists $y_z \in W(o) \cup U(o)$ with $o'(y_z, z) = o(y_z, z) + 1$ and $o'(v, w) = o(v, w)$ for all $(v, w) \in D \setminus \{(y_z, z) \mid z \in Z\}$, then

$$(z, z') \in f(o') \;\text{ for all } z, z' \in Z$$

and

$$(x, z) \in P(f(o')) \;\text{ for all } z \in Z \text{ and for all } x \in L(o') \cup U(o') \setminus Z$$

and

$$(x, x') \in f(o') \;\Leftrightarrow\; (x, x') \in f(o) \;\text{ for all } x, x' \in L(o') \cup U(o') \setminus Z.$$

Although the above axiom may appear to be somewhat complex, its underlying idea as outlined prior to its formal statement is very simple and intuitive. The relatively heavy notation results from the necessity of having to precisely identify the scopes of the premise and of each implication. As is the case for the opinion-cancellation axiom, it is the flexibility of the state-of-opinion setting that allows us to phrase the monotonicity property in terms of trichotomous states. This is a feature that we think is worth emphasizing when it comes to comparing our approach with earlier contributions.

To motivate the final axiom that we employ, we note that the relation $f^B(o)$ is complete and transitive for all $o \in \mathcal{O}$. It turns out that completeness follows from our remaining axioms but transitivity does not so that we have to impose the latter property explicitly. Because the universal equal-goodness relation is transitive, the empty state of opinion is already taken care of by the definition of an opinion aggregator and can thus be excluded in the following definition.

Opinion transitivity. For all $o \in \mathcal{O} \setminus \{o_\emptyset\}$, $f(o)$ is transitive.

The only opinion aggregator that satisfies all of the three axioms introduced above is the Borda opinion aggregator f^B.

Theorem 2 *An opinion aggregator f satisfies opinion cancellation, opinion monotonicity, and opinion transitivity if and only if $f = f^B$.*

Proof. 'If.' That opinion cancellation is satisfied is a consequence of the definition of the Borda opinion scores as the difference between the total numbers of wins and losses in a state of opinion.

To see that opinion monotonicity is satisfied, note first that the addition of a win (loss) to the members of a set of best alternatives in $W(o) \cup U(o)$ (worst alternatives in $L(o) \cup U(o)$) increases (decreases) their Borda opinion scores by one and, because they were equally good before this change as a consequence of all of them being best (worst) elements in the respective set, they continue to be equally good after the change. Moreover, any winners (losers) whose Borda opinion scores do not increase (decrease) and all unclassified alternatives must be worse (better) than those whose scores increase (decrease); again, this follows from the definition of best (worst) elements. Finally, all winners (losers) and unclassified alternatives whose Borda opinion scores are unchanged must be ranked in the same way they are ranked according to the original profile.

That the Borda opinion aggregator satisfies opinion transitivity is immediate.

'Only if.' Suppose that f is an opinion aggregator that satisfies the axioms of the theorem statement. We proceed in three steps, each of which corresponds to one of our three axioms.

First, we employ opinion cancellation to establish that we can, without loss of generality, focus on trichotomous states of opinion. This allows us to restrict attention to trichotomous states in the remaining two steps.

The second step applies the axiom of opinion monotonicity (which is defined for trichotomous states of opinion $o \in \mathcal{T}$) to identify the restriction of the aggregate ranking to the alternatives in $W(o) \cup U(o)$. Starting from the (trichotomous) empty state of opinion, we iteratively increase the number of wins for each alternative in $W(o)$ until we reach the number that corresponds to its Borda opinion score in o. In accordance with the requirements of opinion monotonicity, the alternatives that receive an additional win in a given round are best within the union of the set of winners and the set of unclassified alternatives according to the previous round. Moreover, the corresponding losses are successively assigned to the elements of $L(o)$ in a way such that each member of $L(o)$ experiences at least one loss when the iteration terminates. Upon conclusion of the process, the Borda opinion scores are obtained for the alternatives in $W(o) \cup U(o)$ and, by invoking opinion monotonicity in each round, it follows that these alternatives must be ranked according to the Borda opinion aggregator. This procedure takes care of the alternatives in $W(o) \cup U(o)$ but not necessarily of those in $L(o)$ because the one-by-one augmentation may not be in accordance with the distribution of losses within $L(o)$. For this reason, a dual argument that applies to the alternatives in $L(o) \cup U(o)$ is invoked.

The third and final step then applies opinion transitivity to ensure that the alternatives in $W(o)$ are indeed better than those in $L(o)$.

Step 1. Let $o \in \mathcal{O}$ be a state of opinion. We show that there exists a trichotomous state of opinion $o' \in \mathcal{T}$ such that all Borda opinion scores are the same in o and in o' and, in addition, $f(o') = f(o)$.

If o is itself trichotomous, it follows trivially that $f(o') = f(o)$ with $o' = o$, and all Borda opinion scores are identical for o and for o'.

If o is not trichotomous, there exist three alternatives $x, y, z \in X$ such that $o(x, y) > 0$ and $o(y, z) > 0$.

If $x = z$, we can apply opinion cancellation to conclude that the aggregate ranking $f(o)$ is the same as that corresponding to the state of opinion o' that is obtained from o by reducing $o(x, y)$ and $o(y, x)$ by one. This leaves the Borda opinion scores of all alternatives unchanged.

If $x \neq z$, opinion cancellation allows us to conclude that the aggregate ranking $f(o)$ is the same as that obtained for the state of opinion o' that is obtained from o by reducing $o(x, y)$ and $o(y, z)$ by one, and increasing $o(x, z)$ by one. Again, this leaves the Borda opinion scores of all alternatives unchanged.

If the state of opinion o' is trichotomous, we are done. If not, we can apply the above-described procedure as many times as required to arrive at a trichotomous state of opinion; this iterative process is well-defined and converges because all states of opinion are finite and, in each step, the cardinality of the requisite state of opinion is reduced. Thus, for any state of opinion o, there exists a trichotomous state of opinion o' such that $b(x; o') = b(x; o)$ for all $x \in X$ and $f(o') = f(o)$. Therefore, once it is established that $f(o') = f^B(o')$, it follows immediately that $f(o) = f^B(o)$. As a consequence, it is sufficient to establish in the remainder of the proof that the Borda opinion aggregator must apply to all trichotomous states of opinion.

Step 2. Our next step proceeds by using opinion monotonicity to prove that, for any trichotomous state of opinion $o \in \mathcal{T}$, any two alternatives in the set $W(o) \cup U(o)$ must be ranked by $f(o)$ according to their Borda opinion scores. Note that, for any $x \in W(o)$, this score is given by the number $o(x, z)$ for all $z \in L(o)$ such that $(x, z) \in o$; this follows immediately because no opinions of the form (z, x) can be in a trichotomous state of opinion

o if $x \in W(o)$.

First, note that, by definition of an opinion aggregator, we have $f(o_\emptyset) = X^2 = f^B(o_\emptyset)$. This immediately establishes the claim for the case $o = o_\emptyset$.

Now suppose that $o \in \mathcal{T} \setminus \{o_\emptyset\}$. Because o is a trichotomous state of opinion that is not equal to the empty state of opinion, it follows that $W(o)$ is non-empty. Starting out with the empty state of opinion $o^0 = o_\emptyset \in \mathcal{T}$, we first construct a trichotomous state of opinion o^1 by setting $o^1(y, z_y^1) = o^0(y, z_y^1) + 1$ for all $y \in Y^0 = W(o)$. The alternatives $z_y^1 \in L(o)$ are chosen so that as many elements in $L(o)$ as possible are assigned at least one loss; clearly, this assignment need not be unique but any one of them will do for our purpose because we focus on the elements of $W(o) \cup U(o)$ in this iteration. No other changes are made when moving from $o^0 = o_\emptyset$ to o^1. By definition, $W(o^1) = W(o)$. We note that all elements in $Y^0 = W(o^1) = W(o)$ are best elements in $W(o_\emptyset) \cup U(o_\emptyset)$ according to $f(o_\emptyset) = X^2$ by definition of the universal equal-goodness relation. Part (a) of opinion monotonicity implies that

$$(y, y') \in f(o^1) \text{ for all } y, y' \in Y^0 = W(o^1) = W(o)$$

and

$$(y, x) \in P(f(o^1)) \text{ for all } y \in Y^0 \text{ and for all } x \in W(o^1) \cup U(o^1) \setminus Y^0$$

and

$$(x, x') \in f(o^1) \Leftrightarrow (x, x') \in f(o) \text{ for all } x, x' \in W(o^1) \cup U(o^1) \setminus Y^0.$$

By definition, this means that all alternatives in $W(o^1) \cup U(o^1)$ are ranked according to their Borda opinion scores for the state of opinion o^1.

If $b(y; o) = 1$ for all $y \in W(o^1) = W(o)$, our choice of the z_y^1 guarantees that $L(o^1) = L(o)$ and, together with the equality $W(o^1) = W(o)$, it follows that $U(o^1) = U(o)$ by definition of a trichotomous state of opinion. By construction, the Borda opinion score of each alternative in $W(o^1) \cup U(o^1) = W(o) \cup U(o)$ for o^1 is identical to its Borda opinion score for o. Thus, it follows that the alternatives in $W(o^1) \cup U(o^1) = W(o) \cup U(o)$ are ranked according to their Borda opinion scores not only for the state of opinion o^1 but also for o, as was to be established.

If there exists a set $Y^1 \subseteq W(o^1) = W(o^0)$ of alternatives y such that $b(y; o) \geq 2$, the above procedure can be repeated with a trichotomous state o^2 in place of o^1, with o^1 in place of o^0, with Y^1 in place of Y^0, and with $o^2(y, z_y^2) = o^1(y, z_y^2) + 1$ for all $y \in Y^1$, where the z_y^2 are elements of $L(o)$ chosen so that as many elements of $L(o)$ as possible are assigned at least one loss when combining these losing alternatives with those selected in the previous round of the iteration. Again, it follows that $W(o^2) = W(o^1) = W(o)$ and that the trichotomous state of opinion o^2 that we reach is such that all alternatives in $Y^1 \subseteq W(o^2) = W(o^1) = W(o)$ are best elements in $W(o^1) \cup U(o^1)$ according to $f(o^1)$. Using opinion monotonicity, it follows that all alternatives in $W(o^2) \cup U(o^2)$ are ranked according to their Borda opinion scores for the state of opinion o^2.

If $b(y; o) \leq 2$ for all $y \in W(o^2) = W(o^1) = W(o)$, the choice of the z_y^1 and the z_y^2 guarantees that $L(o^2) = L(o)$ and, again, $U(o^2) = U(o)$. As in the previous round, the Borda opinion score of each alternative in $W(o^2) \cup U(o^2) = W(o) \cup U(o)$ for o^2 is identical to its Borda opinion score for o, and it follows that the alternatives in $W(o^2) \cup U(o^2) =$

$W(o) \cup U(o)$ are ranked according to their Borda opinion scores not only for the state of opinion o^2 but also for o.

If there exists a set $Y^2 \subseteq W(o^2) = W(o^1) = W(o)$ of alternatives y such that $b(y; o) \geq 3$, our finiteness assumption implies that this iteration can be continued as many times as required until we reach a trichotomous state of opinion o^K with $K \geq 3$ such that each alternative in $W(o^K) \cup U(o^K) = W(o) \cup U(o)$ has the same Borda opinion score for o^K and for o, and these alternatives are ranked according to their Borda opinion scores for the state of opinion o.

Using part (b) of opinion monotonicity instead of part (a), the above argument can be applied to conclude that any two alternatives in $L(o) \cup U(o)$ must be ranked by $f(o)$ according to their Borda opinion scores for the state of opinion o.

Step 3. Because $f(o)$ is transitive by opinion transitivity, it follows that all alternatives in $W(o)$ are ranked as better than all alternatives in $L(o)$ which, because the Borda opinion scores of all elements of $W(o)$ are positive and the scores of all elements of $L(o)$ are negative, corresponds to the ranking according to the Borda opinion aggregator as well. This completes the proof. ■

To see that the axioms employed in the above theorem are independent, consider the following examples.

Example 1 *The opinion aggregator of the first example assigns different weights to wins and to losses in calculating opinion scores, an unequal treatment that leads to violations of opinion cancellation. Let* $b'(x; o) = 2 \sum_{y \in X \setminus \{x\}} o(x, y) - \sum_{z \in X \setminus \{x\}} o(z, x)$ *for all* $o \in \mathcal{O}$ *and for all* $x \in X$*. Define* f^1 *by letting*

$$(x, y) \in f^1(o) \;\Leftrightarrow\; b'(x; o) \geq b'(y; o)$$

for all $o \in \mathcal{O}$ *and for all* $x, y \in X$*. This opinion aggregator violates opinion cancellation and satisfies opinion monotonicity and opinion transitivity.*

Example 2 *The opinion aggregator of the second example declares all alternatives to be equally good, independent of the state of opinion; this leads to violations of opinion monotonicity. Let* $f^2(o) = X^2$ *for all* $o \in \mathcal{O}$*, that is,* f^2 *assigns the universal equal-goodness relation to all states of opinion. The opinion aggregator* f^2 *violates opinion monotonicity and satisfies opinion cancellation and opinion transitivity.*

Example 3 *The opinion aggregator of the final example ranks alternatives with a negative Borda opinion score as better than alternatives with a positive score and performs all other comparisons according to the Borda opinion aggregator, thereby generating violations of transitivity. Formally, define the opinion aggregator* f^3 *by letting, for all* $o \in \mathcal{O}$*,* $(x, y) \in P(f^3(o))$ *for all* $x, y \in X$ *such that* $b(x; o) < 0$ *and* $b(y; o) > 0$*, and*

$$(x, y) \in f^3(o) \;\Leftrightarrow\; b(x; o) \geq b(y; o)$$

for all remaining $x, y \in X$*. The opinion aggregator* f^3 *violates opinion transitivity and satisfies opinion cancellation and opinion monotonicity.*

A few remarks on some earlier characterization of Borda's method are in order at this point. Young's (1974) seminal axiomatization employs a variant of pairwise cancellation, in addition to the well-known neutrality axiom, a consistency condition that ensures that different sets of voters who are in agreement can be merged without changing the chosen alternatives, and a faithfulness property requiring that social and individual choice are identical in one-person societies. The setup is that of profile aggregation with a variable population—a social choice function must be capable to select candidates for different sets of voters. Nitzan and Rubinstein (1981) work within an analogous framework, except that they consider social aggregation mechanisms that generate aggregate rankings rather than chosen alternatives. Nitzan and Rubinstein (1981) mostly employ axioms that are analogous to those of Young (1974) but they are formulated in terms of aggregate rankings instead of choices. An exception is their monotonicity property, a positive-responsiveness requirement that they use in place of Young's (1974) faithfulness axiom. Nitzan and Rubinstein (1981) do not impose transitivity on individual goodness relations. It is worth mentioning in this context that, although Young (1974) proves his result for individual antisymmetric orderings, he explicitly says on page 43 of his paper that this assumption can be relaxed. In particular, Young (1974, p. 51) states that his results "can also be proved in much the same way if individual voters are allowed to express their assessment of the alternatives by weak orders, or, even more generally, by partial, antisymmetric relations." Among several others, a recent axiomatization of the Borda method is established by Mihara (2017) who also focuses on rules that generate aggregate orderings from profiles of individual orderings. In his model, the population is fixed and, therefore, axioms such as Young's (1974) consistency and faithfulness do not apply. The axioms employed by Mihara (2017) are a weakening of neutrality, a positional cancellation condition, and positive responsiveness. All of the above results are phrased in terms of individual goodness relations. We note that the monotonicity axiom employed in our characterization as well as the part of cancellation that goes beyond the pairwise restriction are specifically formulated for the opinion-aggregation framework and, therefore, there is a considerable difference between our result and these earlier axiomatizations. Of particular note is the observation that our opinion cancellation property permits us to phrase the monotonicity restriction exclusively in terms of trichotomous states of opinion, an attribute of our characterization that does not have a parallel in the traditional approaches.

We now move on to a characterization of the opinion aggregator that is based on Condorcet's method of majority decision. Compared to our axiomatization of Borda's method, this result is more closely related to a well-known earlier result provided by May (1952). A possible explanation of this difference between the two opinion aggregators considered here is that the majority rule discards much of the information available in a state of opinion (or, in the traditional setting, the information available in a profile of individual goodness relations) and, as a consequence, cannot but rely on a strong notion of neutrality.

As alluded to earlier, we have to impose completeness of the aggregate relation because it does not follow from the remaining axioms (except in the case of the empty state of opinion) and, therefore, the first property used in our characterization of the majority opinion aggregator is the axiom of opinion completeness.

Opinion completeness. For all $o \in \mathcal{O} \setminus \{o_\emptyset\}$, $f(o)$ is complete.

The remaining two axioms parallel the properties of neutrality and responsiveness used by May (1952) in his characterization of the majority rule.

Opinion neutrality. For all $o, o' \in \mathcal{O}$ and for all $x, y, z, w \in X$, if $o(x, y) = o'(z, w)$ and $o(y, x) = o'(w, z)$, then

$$(x, y) \in f(o) \;\Leftrightarrow\; (z, w) \in f(o') \;\text{ and }\; (y, x) \in f(o) \;\Leftrightarrow\; (w, z) \in f(o').$$

Opinion responsiveness. For all $o, o' \in \mathcal{O}$ and for all $x, y \in X$, if $(x, y) \in f(o)$ and $o'(x, y) = o(x, y) + 1$ and $o'(z, w) = o(z, w)$ for all $(z, w) \in D \setminus \{(x, y)\}$, then

$$(x, y) \in P(f(o')).$$

The following theorem characterizes the majority opinion aggregator f^M. Clearly, the result is inspired by and quite similar to May's (1952) original axiomatization and, thus, we claim a lesser degree of originality than that associated with our characterization of the Borda opinion aggregator. We note, however, that our observation indeed provides the first characterization of the majority rule formulated as an opinion aggregator.

Theorem 3 *An opinion aggregator f satisfies opinion completeness, opinion neutrality, and opinion responsiveness if and only if $f = f^M$.*

Proof. 'If.' That the majority opinion aggregator f^M satisfies the three axioms is straightforward to verify.

'Only if.' Suppose that f is an opinion aggregator that satisfies opinion completeness, opinion neutrality, and opinion responsiveness. Because f^M satisfies opinion completeness, it is sufficient to show that, for all $o \in \mathcal{O}$ and for all $x, y \in X$,

$$(x, y) \in I(f^M(o)) \;\Rightarrow\; (x, y) \in I(f(o)) \tag{2}$$

and

$$(x, y) \in P^M(f(o)) \;\Rightarrow\; (x, y) \in P(f(o)). \tag{3}$$

To prove (2), note first that the case $o = o_\emptyset$ is immediate by definition of an opinion aggregator. Now suppose that $o \in \mathcal{O} \setminus \{o_\emptyset\}$ and $(x, y) \in I(f^M(o))$ for some $x, y \in X$ which, by definition, is equivalent to

$$o(x, y) = o(y, x). \tag{4}$$

Suppose that, by way of contradiction, $(x, y) \notin I(f(o))$. By opinion completeness, it follows that $(x, y) \in P(f(o))$ or $(y, x) \in P(f(o))$. Without loss of generality, suppose the former betterness relationship applies; the proof is the same for the latter. Let $o' \in \mathcal{O}$ be such that $o'(x, y) = o(y, x)$ and $o'(y, x) = o(x, y)$. Moreover, for all other pairs $(z, w) \in D \setminus \{(x, y), (y, x)\}$, let $o'(z, w) = o(z, w)$. By construction, it follows that $o'(u, v) = o(u, v)$ for all $(u, v) \in D$ and, therefore, $o' = o$. Thus, $f(o') = f(o)$ and hence $(x, y) \in P(f(o'))$.

Setting $z = y$ and $w = x$, (4) implies $(y, x) \in P(f(o'))$ by opinion neutrality. This is a contradiction and hence (2) must be true.

To establish (3), suppose that $(x, y) \in P(f^M(o))$ and hence $o(x, y) > o(y, x)$ by definition. Let $o' \in \mathcal{O}$ be such that $o'(x, y) = o(y, x)$ and $o'(z, w) = o(z, w)$ for all $(z, w) \in D \setminus \{(x, y)\}$. By (2), it follows that $(x, y) \in I(f(o'))$ and hence $(x, y) \in f(o')$. Repeated application of opinion responsiveness implies $(x, y) \in P(f(o))$, as was to be shown. ■

The independence of the axioms used in Theorem 3 is established by means of the following three examples.

Example 4 *The opinion aggregator of this example replaces some instances of equal goodness with non-comparability, leading to violations of opinion completeness. Let $f^4(o_\emptyset) = X^2$ and, for all $o \in \mathcal{O} \setminus \{o_\emptyset\}$ and for all $x, y \in X$,*

$$(x, y) \in f^4(o) \;\Leftrightarrow\; (x, y) \in P(f^M(o)).$$

The opinion aggregator f^4 violates opinion completeness and satisfies opinion neutrality and opinion responsiveness.

Example 5 *The opinion aggregator defined in this example treats a specific alternative in a manner different from the others, thereby generating violations of opinion neutrality. Let $f^5(o_\emptyset) = X^2$. Fix an alternative $x^0 \in X$ and define, for all $o \in \mathcal{O} \setminus \{o_\emptyset\}$ and for all $x, y \in X$,*

$$(x, y) \in I(f^5(o)) \;\Leftrightarrow\; o(x, y) = o(y, x) \text{ and } x^0 \notin \{x, y\}$$

and

$$\begin{aligned}(x, y) \in P(f^5(o)) \;\Leftrightarrow\;\; & o(x, y) > o(y, x) \text{ or} \\ & x = x^0 \text{ and } o(x, y) = o(y, x).\end{aligned}$$

The opinion aggregator f^5 violates opinion neutrality and satisfies opinion completeness and opinion responsiveness.

Example 6 *Finally, in analogy to the monotonicity property of the previous characterization result, the opinion aggregator f^2 that assigns the universal equal-goodness relation to all states of opinion violates opinion responsiveness and satisfies opinion completeness and opinion neutrality.*

4 Plurality and approval voting

The state-of-opinion framework that we employ is very flexible in the sense that it can easily accommodate traditional forms of voting rules that are familiar from the literature and from real-world applications. Moreover, the Borda opinion aggregator and the majority opinion aggregator coincide with the requisite voting rule under consideration in some important special cases. Analogous observations regarding the equivalence of various voting schemes on specific domains are established by Martínez and Moreno (2017). This section discusses two examples that are of considerable practical relevance—namely, plurality voting and approval voting and some variants of the latter. We emphasize that these are mere illustrations and, as such, they are not intended to provide a general method of moving from an informationally different environment to the opinion-aggregation framework. In fact, a more general approach seems to be difficult to implement because the way we link opinion aggregation to the plurality rule relies on ballots that differ from those utilized in the context of approval voting. While it may be possible to establish some general connections, it is not obvious to us how we can identify a precise definition of an opinion aggregator from a ballot profile without knowledge of the specific structure of the profile that the analysis is based on.

A well-known voting rule is the plurality rule; see, for instance, Richelson (1978), Ching (1996), Goodin and List (2006), Yeh (2008), Sekiguchi (2012), and Kelly and Qi (2016) for characterizations. Consider a non-empty and finite set N of voters. To apply plurality voting, the inputs required from the voters are single-alternative ballots of the form $s_i \in X \cup \{\emptyset\}$ for all $i \in N$, where s_i is assumed to be the (unique) top alternative for voter i if $s_i \in X$, and voter i abstains if $s_i = \emptyset$. Let $s = (s_i)_{i \in N}$ be a single-alternative ballot profile, and let $\mathcal{S}$ be the set of all possible single-alternative ballot profiles. The plurality score of $x \in X \cup \{\emptyset\}$ for the single-alternative ballot profile $s \in \mathcal{S}$ is

$$p(x; s) = |\{i \in N \mid x = s_i\}|,$$

and the plurality rule $g^P : \mathcal{S} \to \mathcal{B}$ is defined by

$$(x, y) \in g^P(s) \;\Leftrightarrow\; p(x; s) \geq p(y; s)$$

for all $s \in \mathcal{S}$ and for all $x, y \in X$.

We can assign a state of opinion o_s to any single-alternative ballot profile $s \in \mathcal{S}$. To do so, we assume that whenever $x = s_i$ for some $i \in N$, x wins against all other alternatives $y \in X \setminus \{x\}$—and these are the only opinions expressed (implicitly) by voter i. Thus, we obtain o_s from s by defining

$$o_s(x, y) = |\{i \in N \mid x = s_i\}| = p(x; s)$$

for all $(x, y) \in D$; this follows immediately because, according to our interpretation, $x = s_i$ means that x wins against all other alternatives. Therefore, every time there is a voter who has x as the top element, the pair (x, y) must be in the state of opinion o_s for each $y \in X \setminus \{x\}$. Because there are $|X| - 1$ alternatives that lose against x, the number of times x wins against some other alternative in the state of opinion o_s is given by the product of $|X| - 1$ and the plurality score $p(x; s)$. Likewise, the number of times x loses against another alternative is given by the number of voters for whom an alternative $y \in X$ other than x is on top—that is, the number $|N| - p(x; s) - p(\emptyset; s)$. Thus, using the definition of the Borda opinion scores, it follows that

$$\begin{aligned} b(x; o_s) &= \sum_{y \in X \setminus \{x\}} o_s(x, y) - \sum_{y \in X \setminus \{x\}} o_s(y, x) \\ &= (|X| - 1) p(x; s) - (|N| - p(x; s) - p(\emptyset; s)) \\ &= |X| p(x; s) - |N| + p(\emptyset; s) \end{aligned}$$

for all $s \in \mathcal{S}$ and for all $x \in X$ and, therefore,

$$\begin{array}{rcl}(x, y) \in f^B(o_s) & \Leftrightarrow & |X|p(x; s) - |N| + p(\emptyset; s) \geq |X|p(y; s) - |N| + p(\emptyset; s) \\ & \Leftrightarrow & |X|p(x; s) \geq |X|p(y; s) \\ & \Leftrightarrow & p(x; s) \geq p(y; s) \\ & \Leftrightarrow & (x, y) \in g^P(s)\end{array}$$

for all $s \in \mathcal{S}$ and for all $x \in X$.

According to the majority opinion aggregator, it follows that

$$\begin{array}{rcl}(x, y) \in f^M(o_s) & \Leftrightarrow & o_s(x, y) \geq o_s(y, x) \\ & \Leftrightarrow & p(x; s) \geq p(y; s) \\ & \Leftrightarrow & (x, y) \in g^P(s)\end{array}$$

for all $s \in \mathcal{S}$ and for all $x \in X$.

Therefore, the Borda opinion aggregator and the majority opinion aggregator agree with the plurality rule on the set of states of opinion that are generated by single-alternative ballot profiles.

As a second example, consider the method of approval voting analyzed by Brams and Fishburn (1978, 1983); see also Brams (1975). To employ this voting rule, each voter $i \in N$ gets to submit a set-valued ballot with a strict subset $C_i \subsetneq X$ of approved-of alternatives. A set-valued ballot profile is given by an $|N|$-tuple $C = (C_i)_{i \in N}$ of strict subsets of X. The set of all possible set-valued ballot profiles is denoted by $\mathcal{C}$. The approval score of an alternative $x \in X$ for a set-valued ballot profile $C \in \mathcal{C}$ is given by

$$a(x; C) = |\{i \in N \mid x \in C_i\}|.$$

The approval-voting rule $h^A: \mathcal{C} \to \mathcal{B}$ is defined by

$$(x, y) \in h^A(C) \Leftrightarrow a(x; C) \geq a(y; C)$$

for all $C \in \mathcal{C}$ and for all $x, y \in X$.

We can assign a state of opinion o_C to any set-valued ballot profile $C \in \mathcal{C}$. To do so, we assume that whenever $x \in C_i$ for some $i \in N$, x wins against all alternatives $y \in X \setminus C_i$—and these are the only opinions expressed (implicitly) by voter i. Thus, we obtain o_C from C by defining

$$\begin{array}{rcl}o_C(x, y) & = & |\{i \in N \mid x \in C_i \text{ and } y \notin C_i\}| \\ & = & a(x; C) - |\{i \in N \mid x \in C_i \text{ and } y \in C_i\}|\end{array}$$

for all $(x, y) \in D$; this follows because, according to our interpretation, $x \in C_i$ means that x wins against all alternatives y that are not approved of by voter i. Therefore, if there is a voter $i \in N$ who has x as an approved-of alternative, the pair (x, y) must be in the state of opinion o_C for every $y \in X \setminus C_i$; on the other hand, if y is also approved of by i, the pair (x, y) is not in o_C. This means that the number of instances a pair (x, y) appears in the state of opinion o_C is not equal to the approval score of x because the number of instances in which both x and y are approved of must be subtracted from $a(x;C)$ to arrive at $o_C(x,y)$. It follows that

$$\begin{aligned} b(x;o_C) &= \sum_{y\in X\setminus\{x\}} [a(x;C) - |\{i \in N \mid x \in C_i \text{ and } y \in C_i\}|] \\ &\quad - \sum_{y\in X\setminus\{x\}} [a(y;C) - |\{i \in N \mid x \in C_i \text{ and } y \in C_i\}|] \\ &= \sum_{y\in X\setminus\{x\}} [a(x;C) - a(y;C)] \\ &= (|X|-1)a(x;C) - \sum_{y\in X\setminus\{x\}} a(y;C) \end{aligned}$$

for all $C \in \mathcal{C}$ and for all $x \in X$ and, therefore,

$$\begin{aligned} (x,y) \in f^B(o_C) &\Leftrightarrow (|X|-1)a(x;C) - \sum_{z\in X\setminus\{x\}} a(z;C) \\ &\qquad \geq (|X|-1)a(y;C) - \sum_{z\in X\setminus\{y\}} a(z;C) \\ &\Leftrightarrow (|X|-1)a(x;C) - a(y;C) \geq (|X|-1)a(y;C) - a(x;C) \\ &\Leftrightarrow |X|a(x;C) \geq |X|a(y;C) \\ &\Leftrightarrow a(x;C) \geq a(y;C) \\ &\Leftrightarrow (x,y) \in h^A(C) \end{aligned}$$

for all $C \in \mathcal{C}$ and for all $x, y \in X$.

For the majority opinion aggregator f^M, we obtain

$$\begin{aligned} (x,y) \in f^M(o_C) &\Leftrightarrow a(x;C) - |\{i \in N \mid x \in C_i \text{ and } y \in C_i\}| \\ &\qquad \geq a(y;C) - |\{i \in N \mid x \in C_i \text{ and } y \in C_i\}| \\ &\Leftrightarrow a(x;C) \geq a(y;C) \\ &\Leftrightarrow (x,y) \in h^A(C) \end{aligned}$$

for all $C \in \mathcal{C}$ and for all $x, y \in X$.

Thus, the Borda opinion aggregator and the majority opinion aggregator agree with the approval-voting rule on the set of states of opinion that are generated by set-valued ballot profiles.

There are several variants of approval voting that have been discussed in the literature; see, for instance, Alcantud and Laruelle (2014) and Gonzalez, Laruelle, and Solal (2019).

If, for example, disapproval voting is used instead of approval voting, each voter again gets to submit a set-valued ballot. In this case, the set identifies the alternatives that the voter disapproves of, and the alternatives are ranked inversely with respect to their disapproval scores. Again, we can define a state of opinion that corresponds to each set-valued ballot profile and the observations for approval voting translate directly to disapproval voting. This is the case because we can reinterpret the alternatives that are not disapproved of as alternatives that are approved of and, analogously, the disapproved-of alternatives are those that are not approved of.

A version of mixed approval-disapproval voting proceeds by allowing each voter to submit a ballot that consists of two disjoint strict subsets of the universal set of alternatives—one set of alternatives that are approved of, one set of alternatives that the voter disapproves of. The mixed approval-disapproval rule then ranks the alternatives according to the differences between their approval scores and their disapproval scores. In terms of states of opinion, this can again be translated into the approval-voting framework. To do so, we can use the approved-of sets of each voter as before and assign the disapproved-of sets to a set of voters that is disjoint from the original set of voters. This is perfectly legitimate in our setting because, by assumption, it does not matter who holds which opinion. Now each disapproved-of set can be interpreted as a set of alternatives that are not approved of, and its complement as a set of approved-of alternatives. By using this method, we arrive at a larger population of voters whose ballots consist of single sets of approved-of alternatives. This brings us back to the case of approval voting and, again, the same equivalence results as above are obtained.

A corollary of the above equivalence results is that the majority opinion aggregator f^M always generates transitive aggregate relations if attention is restricted to states of opinion that can be derived from single-alternative ballots or from set-valued ballots. Thus, the domains considered here provide novel examples for scenarios in which Condorcet cycles can be ruled out; see also Black (1948) and Sen (1970, Chapters 10 and 10*), for instance. An intuitive explanation for this phenomenon is the dichotomous nature of these states of opinion.

The observations of this section represent instances of opinion aggregation with restricted domains: depending on the underlying voting rule, only some states of opinion can materialize while others are ruled out. This raises the issue of providing characterization results on limited domains, a topic that is not explicitly addressed in this paper. The axiomatizations of the previous section are established on an unrestricted domain and possible modifications that employ domain conditions constitute, at this stage, an open question. This is undoubtedly an important aspect of the framework that we propose. At the same time, it seems natural to us to focus on the unlimited-domain case in this initial contribution and leave the exploration of domain restrictions for future work. In this respect, our approach parallels that followed in numerous (if not most) analogous advancements of relatively novel topics such as the existence of social welfare functions with desirable properties or the study of strategy-proofness.

5 Concluding remarks

An important feature of the opinion aggregators discussed in this paper is that they, unlike other collective choice mechanisms, do not rely on any properties of the individual inputs. Both the Borda opinion aggregator and the majority opinion aggregator are well-defined without having to assume that there are individual relations that are complete or transitive—or even merely acyclical. As a consequence, the rules are well-suited to be analyzed in the context of opinion aggregation.

The characterization results for the two aggregators exhibit an interesting parallel structure. In each of them, three properties are employed that fall into the same three categories. The transitivity axiom required in the characterization of the Borda opinion aggregator is a coherence requirement imposed on the social relation to be established, as is the completeness property that appears in our axiomatization of the majority opinion aggregator. The axioms of opinion cancellation and opinion neutrality express independence requirements. Finally, opinion monotonicity and opinion responsiveness ensure that an opinion aggregator adjusts suitably to specific changes in the state of opinion under consideration.

The notion of an opinion employed here is based on a betterness interpretation—we think of a pair (x, y) being in a state of opinion o to mean that x is superior to y. Our observations can easily be amended if an at-least-as-good-as interpretation is adopted instead; this follows immediately because all instances of equal goodness cancel out in the definitions of the Borda opinion aggregator and the majority opinion aggregator. The reason why we choose the betterness variant is that it allows us to accommodate applications such as that of opinions representing potentially conflicting impulses in cognitive processes alluded to earlier. To see this point, note that if the pair (x, y) represents a relationship of superiority of x over y according to an individual, it is perfectly possible to have both (x, y) and (y, x) appear in a state of opinion. On the other hand, if (x, y) stands for x being at least as good as y, this is no longer possible. Betterness of x compared to y cannot but be defined as the presence of (x, y) and the absence of (y, x) in an agent's collection of opinions. Thus, under this interpretation, it is logically impossible to have superiority of x over y and superiority of y over x at the same time.

References

Alcantud, J.C.R. and A. Laruelle (2014), Dis&approval voting: a characterization. Social Choice and Welfare, 43, 1–10.

Armstrong, W.E. (1939), The determinateness of the utility function. Economic Journal, 49, 453–467.

Arrow, K.J. (1951, second ed. 1963), Social Choice and Individual Values. Wiley, New York.

Barberà, S., W. Bossert, and K. Suzumura (2021), Daunou's voting rule and the lexicographic assignment of priorities. Social Choice and Welfare, 56, 259–289.

Black, D. (1948), On the rationale of group decision-making. Journal of Political Economy, 56, 23–34.

Borda, J.-C. de (1781), Mémoire sur les élections au scrutin. Mémoires de l'Académie Royale des Sciences année 1781, 657–665. Translated and reprinted in I. McLean and A.B. Urken, eds. (1995), Classics of Social Choice. University of Michigan Press, Ann Arbor, Chapter 5.

Bossert, W. and K. Suzumura (2020), Positionalist voting rules: a general definition and axiomatic characterizations. Social Choice and Welfare, 55, 85–116.

Brams, S.J. (1975), Game Theory and Politics. Free Press, New York.

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

Brams, S.J. and P.C. Fishburn (1983), Approval Voting. Birkhäuser, Boston.

Chebotarev, P.Y. (1994), Aggregation of preferences by the generalized row sum method. Mathematical Social Sciences, 21, 293–320.

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

Condorcet, M.J.A.N. de (1785), Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix. Imprimerie Royale, Paris. Translated and reprinted in part in I. McLean and A.B. Urken, eds. (1995), Classics of Social Choice. University of Michigan Press, Ann Arbor, Chapter 6.

Daunou, P.C.F. (1803), Mémoire sur les élections au scrutin. Baudouin, imprimeur de l'Institut National, Paris. Translated and reprinted in I. McLean and A.B. Urken, eds. (1995), Classics of Social Choice. University of Michigan Press, Ann Arbor, Chapter 11.

Fishburn, P.C. (1970), Intransitive indifference in preference theory: a survey. Operations Research, 18, 207–228.

Fishburn, P.C. (1984), Probabilistic social choice based on simple voting comparisons. Review of Economic Studies, 51, 683–692.

Gärdenfors, P. (1973), Positionalist voting functions. Theory and Decision, 4, 1–24.

Gonzalez, S., A. Laruelle, and P. Solal (2019), Dilemma with approval and disapproval votes. Social Choice and Welfare, 53, 497–517.

Goodin, R.E. and C. List (2006), A conditional defense of plurality rule: generalizing May's theorem in a restricted informational environment. American Journal of Political Science, 50, 940–949.

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

Herrero, C. and A. Villar (2021), Group decisions from individual rankings: the Borda-Condorcet rule. European Journal of Operational Research, 291, 757–765.

Jackson, M.O. and L. Yariv (2015), Collective dynamic choice: the necessity of time inconsistency. American Economic Journal: Microeconomics, 7, 150–178.

Kelly, J.S. and S. Qi (2016), Characterizing plurality rule on a fixed population. Economics Letters, 146, 39–41.

Kraft, C.H., J.W. Pratt, and A. Seidenberg (1959), Intuitive probability on finite sets. Annals of Mathematical Statistics, 30, 408–419.

Luce, R.D. (1956), Semiorders and a theory of utility discrimination. Econometrica, 24, 178–191.

Luce, R.D. and H. Raiffa (1957), Games and Decisions, Wiley, New York.

Martínez, R. and B. Moreno (2017), Qualified voting systems. Mathematical Social Sciences, 88, 49–54.

May, K.O. (1952), A set of independent necessary and sufficient conditions for simple majority decision. Econometrica, 20, 680–684.

Mihara, H.R. (2017), Characterizing the Borda ranking rule for a fixed population. MPRA Paper No. 78093, available at `https://mpra.ub.uni-muenchen.de/78093/`.

Morales, J.I. (1797), Memoria matemática sobre el cálculo de la opinion en las elecciones. Imprenta Real, Madrid. Translated and reprinted in I. McLean and A.B. Urken, eds. (1995), Classics of Social Choice. University of Michigan Press, Ann Arbor, Chapter 10.

Nitzan, S. and A. Rubinstein (1981), A further characterization of Borda ranking method. Public Choice, 36, 153–158.

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

Sekiguchi, Y. (2012), A characterization of the plurality rule. Economics Letters, 116, 330–332.

Sen, A.K. (1970), Collective Choice and Social Welfare. Holden-Day, San Francisco.

Yeh, C.-H. (2008), An efficiency characterization of plurality rule in collective choice problems. Economic Theory, 34, 575–583.

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

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