Aggregating Opinions

Salvador Barberà
Walter Bossert

May2021

Barcelona GSE Working Paper Series

Working Paper nº 1262

Aggregating opinions*

Salvador Barberà
MOVE, Universitat Autònoma de Barcelona
and Barcelona GSE
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@videotron.ca

This version: May 26, 2021

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 a social relation on the set of alternatives to every possible multiset of opinions. Our primary motivation is to extend some basic results of social choice theory to a more general model in which no specific reference to agents generating or holding opinions is needed. It turns out that, although our analysis has some bearing on those cases where opinions reflect the preferences of agents in a society, it is not limited to them. In addition to the preference interpretation, opinions can also be used to represent other forms of comparative assessments. The main results of the paper provide characterizations of suitably defined versions of the Borda rule and the majority rule. *Journal of Economic Literature* Classification Nos.: D71, D72.

* Salvador Barberà acknowledges financial support from the Spanish Ministry of Economy and Competitiveness, through the Severo Ochoa Programme for Centers of Excellence in R&D (SEV-2015-0563) and grant ECO2017-83534-P and FEDER, and from the Generalitat de Catalunya, through grant 2017SGR0711. We thank Susumu Cato, Anke Gerber, 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 (x, y) of alternatives. Such a pair expresses the idea that x is superior to y in some sense. Our primary motivation is to extend some basic results of social choice theory to a more general model in which no specific reference to agents generating or holding opinions is needed. It turns out that, although our analysis has some bearing on those cases where opinions reflect the preferences of agents in a society, it is not limited to them. In addition to the preference interpretation, opinions can also be used to represent other forms of comparative assessments. 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. 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 finite multiset, each of whose elements is an opinion that may appear in the same state any number of times. The special case of the empty state of opinion—a multiset without any elements—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 preference is covered, and so is the possibility of each agent merely passing a judgment on a single pair of alternatives.

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 indifference 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.

We focus on (and characterize) two specific opinion aggregators. The first of these, the Borda opinion aggregator, is in the spirit of Borda (1781); the second, the majority opinion aggregator, in that of Condorcet (1785). It is well-known that the positions of these two pioneers of social choice theory diverged considerably, and 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).

The Borda opinion aggregator ranks any two alternatives based on the total number of wins minus the total number of losses that each one accumulates across all opinions in which it appears. The second opinion aggregator we consider consists of the majority rule formulated in our environment. Thus, according to this opinion aggregator, an alternative x is socially at least as good as an alternative y if and only if the number of times the pair (x, y) appears in a state of opinion is greater than or equal to the number of times (y, x) is present in that state. Herrero and Villar (2021) examine a rule that represents a combination of the proposals by Borda and Condorcet.

We have several reasons to concentrate on the analysis of these two basic methods, which have competed for primacy among many other conceivable aggregators for over two centuries.

One motivation for their use is that both of them 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. This is in contrast to most other aggregators that need additional information before they can be defined. For example, scoring rules require the states of opinion to be rich enough to properly allow for the definition of (strict) orderings of the alternatives held by the members of a previously specified set of agents. This is not required to define the Borda opinion aggregator or the majority opinion aggregator.

A second important reason for our focus on these aggregators is that, in natural contexts where different well-known rules can be defined, our two aggregators coincide with these rules, thereby even becoming equivalent themselves. 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 preferred alternative, agents express their opinions between that alternative and any other. We show that if attention is restricted to such states, the social ranking according to the plurality rule winner coincides with the social ranking according to the Borda opinion aggregator and 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 set of opinions that consists of all opinions (x, y), where x is approved 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 the majority opinion aggregator. Note that, in particular, for the restricted sets of states of opinion generated by these two types of ballots, the majority rule actually generates a transitive social relation. This is the case because Condorcet cycles cannot occur on these domains. The above-described equivalences are established formally once the requisite definitions have been introduced.

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. According to Morales, the term opinion is used in the same way as in our model represents exactly the same as in our model but, in addition, Morales strongly defends this notion as a natural unit of measurement. Borda's presentation also refers to we 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 preferences remains the same independently of the rank they occupy. Of course, we do not maintain this cardinal point of view as a starting point, but our axiomatic characterization can be interpreted as a foundation for its ex-post defense.

Interestingly, our two particular rules 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 preference orderings. The Borda count is often presented as a special case of a scoring rule. A scoring rule assigns a weight to each position of an alternative in an agent's strict preference ordering and then ranks any two alternatives socially by comparing their total scores. The weights that correspond to the Borda rule are given by the number of wins minus the number of losses experienced in each position. Sometimes the number of wins is used as a weight instead of this difference. As is well-known, the two formulations are equivalent if the individual preferences are complete; if this is not the case, the two versions differ and the difference is to be used. In this context, we note that the Borda rule sets itself apart from other scoring rules because it is perfectly well-defined even if preferences are not complete or not transitive; see also Young (1974) for this observation. Analogously, the majority rule is well-defined without the need of assuming that it operates on profiles of individual orderings—again, no specific requirements regarding where opinions come from are required. This illustrates that such rules are not only attractive per se, but also based on more universal principles than others.

We conclude this introduction with an informal explanation of the properties that we employ to characterize the two opinion aggregators under consideration.

Let us begin with the Borda opinion aggregator. As mentioned earlier, we assume throughout that an opinion aggregator assigns the universal indifference relation to the empty state of opinion. Clearly, the universal indifference relation is complete and transitive but, a priori, none of these coherence properties need to be satisfied by other relations that emerge by applying the rule. Because the Borda opinion aggregator generates only complete and transitive social 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 first property we impose is that all relations generated by an opinion aggregator be transitive.

The axiom of opinion cancellation is analogous to Young's (1974) cancellation property. Intuitively, it requires that opposing opinions cancel each other out when determining the social relation.

The third and final property we employ in axiomatizing the Borda opinion aggregator is an opinion-monotonicity requirement. Loosely speaking, if, under certain circumstances, the situation of some alternatives that are among the best for a specific state of opinion improves by adding a favorable opinion for each of them, then they improve in the social ranking compared to those who were previously best along with them but no longer are—and all else is unchanged. A dual requirement is imposed on the response of the opinion aggregator to the deterioration of worst alternatives.

Moving on to the majority opinion aggregator, we note first that it always generates complete but not necessarily transitive social relations. By virtue of the definition of an opinion aggregator, completeness is guaranteed for the empty state of opinion because the universal indifference relation is complete. For other states of opinion, completeness has to be imposed explicitly, and this is what the axiom of opinion completeness does.

The next axiom is opinion neutrality, which requires that the relative ranking of any two alternatives according to the social relation can only depend on the numbers of times the corresponding pairs appear in a state of opinion.

Finally, opinion responsiveness is an alternative monotonicity axiom that requires the social relation to be sensitive to the improvement of an alternative with respect to the addition of specific favorable opinions.

Although the last two axioms bear a family resemblance to those employed by May (1952) and Sen (1970, Theorem 5*1), our characterization has a novel aspect to it because it illustrates how these classical results can be extended to the considerably more general framework of opinion aggregation.

2 Opinions and opinion aggregators

We model a collective decision process by means of individual opinions. 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 a strict preference for x over y expressed by a voter. 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 voter or of two distinct voters. It is irrelevant which of these two options applies—all opinions are treated impartially, no matter who holds them.

A state of opinion o is a finite multiset of opinions on pairs $(x, y) \in D$. The set of all possible states of opinion is denoted by $\mathcal{O}$. The multiplicity $m((x, y); o) \in \mathbb{N}_0$ of an element $(x, y) \in D$ in a state of opinion $o \in \mathcal{O}$ is the number of times this element (x, y) appears in the state of opinion o. Note that the multiplicity $m((x, y); o)$ of $(x, y) \in D$ in $o \in \mathcal{O}$ is well-defined even if (x, y) does not appear in o; in that case, we have $m((x, y); o) = 0$. Thus, a multiset $o \in \mathcal{O}$ is unambiguously identified by specifying the multiplicity $m((x, y); o)$ of every pair $(x, y) \in D$. The cardinality $|o|$ of a state of opinion $o \in \mathcal{O}$ is the sum of the multiplicities of the elements in the state of opinion o, that is, $|o| = \sum_{(x,y)\in D} m((x, y); o)$. The empty state of opinion $o_\emptyset \in \mathcal{O}$ is the empty multiset and its cardinality is equal to zero. Because we restrict attention to finite multisets, all multiplicities and all cardinalities considered here are finite.

A set is a special case of a multiset in which all elements have a multiplicity of one but a multiset need not be a set. For example, for the set of alternatives $X = \{x, y, z\}$, the state of opinion $o = \{(x, y), (z, x), (x, y), (y, x)\}$ is a multiset but not a set. In this example, the multiplicity of (x, y) is equal to two, and the multiplicities of (z, x) and of (y, x) are equal to one. The multiplicities of all other pairs in D are equal to zero and, therefore, the cardinality $|o|$ of this multiset is equal to four.

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 indifference 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 $f(o)$.

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

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

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 $x, y \in X$. This definition is analogous to the definition of the traditional Borda (1781) rule for more structured voting rules that depend on individual strict orderings. We note that the Borda rule is a special case of a scoring rule (Young, 1975), obtained by assigning the difference between the number of wins and the number of losses in a profile as a weight to each alternative; of course, any increasing affine transformation of these weights leads to the same rule. The Borda rule sets itself apart from other scoring rules because it is well-defined even if individual preferences are neither complete nor transitive. Closely related to the scoring rules are the positional voting functions analyzed in Gärdenfors (1973).

The majority opinion aggregator f^M is defined by letting

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

for all $o \in \mathcal{O}$ and for all $x, y \in X$. As is the case for the Borda opinion aggregator, the analogue of the majority rule formulated as a voting rule that depends on individual preferences is well-defined even if these preferences are neither complete nor transitive.

The state-of-opinion framework that we employ is quite 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, as illustrated below, the two opinion aggregators that we focus on coincide with the requisite voting rules 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).

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 an alternative $x \in X$ 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 $g^P\colon \mathcal{S} \to \mathcal{R}$ 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 $\hat{o}(s)$ 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 $\hat{o}(s)$ from s by defining the multiplicities as

$$\begin{aligned} m((x, y); \hat{o}(s)) &= |\{i \in N \mid x = s_i\}| \\ &= p(x; s) \end{aligned}$$

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 $\hat{o}(s)$ 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 $\hat{o}(s)$ 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 x is not on top—that is, the number $|N| - p(x; s)$. Thus, using the definition of the Borda opinion scores, it follows that

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

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

$$\begin{aligned} (x, y) \in f^B(\hat{o}(s)) &\Leftrightarrow |X| p(x; s) - |N| \geq |X| p(y; s) - |N| \\ &\Leftrightarrow |X| p(x; s) \geq |X| p(y; s) \\ &\Leftrightarrow p(x; s) \geq p(y; s) \\ &\Leftrightarrow (x, y) \in g^P(s) \end{aligned}$$

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

According to the majority opinion aggregator, it follows that

$$\begin{aligned} (x, y) \in f^M(\hat{o}(s)) &\Leftrightarrow m((x, y); \hat{o}(s)) \geq m((y, x); \hat{o}(s)) \\ &\Leftrightarrow p(x; s) \geq p(y; s) \\ &\Leftrightarrow (x, y) \in g^P(s) \end{aligned}$$

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 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 $h^A: \mathcal{C} \to \mathcal{R}$ 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 $x \in X$.

We can assign a state of opinion $\bar{o}(C)$ 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 $\bar{o}(C)$ from C by defining the multiplicities as

$$\begin{aligned} m((x, y); \bar{o}(C)) &= |\{i \in N \mid x \in C_i \text{ and } y \notin C_i\}| \\ &= a(x; C) - |\{i \in N \mid x \in C_i \text{ and } y \in C_i\}| \end{aligned}$$

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 $\bar{o}(C)$ 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 $\bar{o}(C)$. This means that the multiplicity of a pair (x, y) in the state of opinion $\bar{o}(C)$ 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 $m((x, y); \bar{o}(C))$. It follows that

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

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

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

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

For the majority opinion aggregator f^M, we obtain

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

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.

3 Characterizations

Our first result provides a characterization the Borda opinion aggregator. To motivate our first axiom, 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 indifference 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 first part of the next axiom 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 social relation for the state of opinion under consideration. That is, if the multiplicities of (x, y) and of (y, x) are reduced by one and no other changes occur, the social 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 is employed by Young (1974) in his characterization of the Borda rule.

There also is a second (more subtle) possibility of eliminating opinions that involve a win and a loss. This occurs if 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 first case, the alternatives x and z differ. Reducing the multiplicities 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 increasing the multiplicity of (x, z) by one. In this manner, the numbers of wins and losses for x and z are preserved so that equality of the two resulting social relations can again be required. The axiom that deals with these two types of cancellation represents a distinguishing feature of the Borda opinion aggregator.

Opinion cancellation. (a) For all $o, o' \in \mathcal{O}$ and for all $x, y \in X$, if $m((x, y); o) > 0$ and $m((y, x); o) > 0$ and $m((x, y); o') = m((x, y); o) - 1$ and $m((y, x); o') = m((y, x); o) - 1$ and $m((v, w); o') = m((v, w); o)$ for all $(v, w) \in D \setminus \{(x, y), (y, x)\}$, then

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

(b) For all $o, o' \in \mathcal{O}$ and for all $x, y, z \in X$, if $m((x, y); o) > 0$ and $m((y, z); o) > 0$ and $m((x, y); o') = m((x, y); o) - 1$ and $m((y, z); o') = m((y, z); o) - 1$ and $m((x, z); o') =$

$m((x,z);o)+1$ and $m((v,w);o') = m((v,w);o)$ for all $(v,w) \in D \setminus \{(x,y),(y,x),(x,z)\}$, then

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

The final 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

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

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

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

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

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

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 social relation $f(o)$ that corresponds to o according to an opinion aggregator f. Furthermore, consider any set of alternatives $Y \subseteq X$ the members of which are best elements according to $f(o)$; note that not all best elements need to be included in Y. Now define a new trichotomous state of opinion o' by adding one win for each of the members of Y.

We require such a move to have three consequences that we consider very intuitive. First, note that the members of Y are pairwise indifferent according to $f(o)$ because all of them are best elements. 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 pairwise indifferent 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 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 according to the social relation that corresponds to a trichotomous state of opinion, we add one loss to each member of a subset of the worst elements 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 $Y \subseteq X$ such that $(y, x) \in f(o)$ for all $y \in Y$ and for all $x \in X$, if there exists $Z \subseteq U(o) \cup L(o)$ such that, for all $y \in Y$, there exists $z_y \in Z$ with $m((y, z_y); o') = m((y, z_y); o) + 1$ for all $y \in Y$ and $m((v, w); o') = m((v, w); o)$ 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 $Z \subseteq X$ such that $(x, z) \in f(o)$ for all $z \in Z$ and for all $x \in X$, if there exists $Y \subseteq U(o) \cup W(o)$ such that, for all $z \in Z$, there exists $y_z \in Y$ with $m((y_z, z); o') = m((y_z, z); o) + 1$ for all $z \in Z$ and $m((v, w); o') = m((v, w); o)$ 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.

The only opinion aggregator that satisfies all of the three axioms introduced above is the Borda opinion aggregator f^B. See, in addition to Young (1974), Hansson and Sahlquist (1976) and Nitzan and Rubinstein (1981) for alternative axiomatizations of the Borda rule in the traditional setting.

Theorem 1 *An opinion aggregator f satisfies opinion transitivity, opinion cancellation, and opinion monotonicity if and only if $f = f^B$.*

Proof. 'If.' That the Borda opinion aggregator satisfies opinion transitivity is immediate.

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 (worst) alternatives increases (decreases) their Borda opinion scores by one and, because they were pairwise indifferent before this change as a consequence of all of them being best (worst) elements, they continue to be pairwise indifferent 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.

'Only if.' Suppose that f is an opinion aggregator that satisfies the axioms of the theorem statement.

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 $m((x, y); o) > 0$ and $m((y, z); o) > 0$.

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

If $x \neq z$, part (b) of opinion cancellation allows us to conclude that the social ranking $f(o)$ is the same as that obtained for the state of opinion o' that is obtained from o by reducing $m((x, y); o)$ and $m((y, z); o)$ by one, and increasing $m((x, z); o)$ 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 because all states of opinion are finite. 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 $x \in X$. 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 is to use opinion monotonicity to prove that, for any trichotomous state of opinion $o \in \mathcal{T}$, any two alternatives in $W(o) \cup U(o)$ 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 sum of the multiplicities $m((x,z);o)$ 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 o because this state of opinion is assumed to be trichotomous.

By definition of an opinion aggregator, we have

$$f(o_\emptyset) = X^2 = f^B(o_\emptyset). \tag{1}$$

Now suppose that $o \in \mathcal{T} \setminus \{o_\emptyset\}$. 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 $m((y,z_y);o^1) = m((y,z_y);o_\emptyset)+1 = 0+1 = 1$ for all $y \in Y^0 = W(o)$, where each z_y can be any alternative in $L(o)$ such that $m((y,z_y);o) > m((y,z_y);o^0) = 0$. No other changes are made when moving from $o^0 = o_\emptyset$ to o^1. Note that, by definition, $W(o^1) = W(o)$. Using (1), part (a) of opinion monotonicity implies that

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

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 by the relation $f(o^1)$.

If $m((y,z);o) = 1$ for all $y \in W(o)$, it follows that $U(o^1) = U(o)$ and $L(o^1) = L(o)$ in addition to $W(o^1) = W(o)$ and, therefore, $o^1 = o$. Thus, all alternatives in $W(o) \cup U(o)$ are ranked according to their Borda opinion scores by the relation $f(o) = f(o^1)$.

If there exists a set $Y^1 \subseteq W(o^1)$ of alternatives y such that $m((y,z);o) \geq 2$, the above procedure can be repeated with o^2 in place of o^1 and with o^1 in place of o^0, with Y^1 in place of Y^0, and with $m((y,z_y);o^2) = m((y,z_y);o^1) + 1 = 1 + 1 = 2$ for all $y \in Y^1$, where each z_y can be any alternative in $L(o)$ such that $m((y,z_y);o) > m((y,z_y);o^1)$. 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 $W(o^2) \cup U(o^2)$ are ranked according to their Borda opinion scores by $f(o^2)$. Because of our finiteness assumption, this iteration can be continued as many times as required until we reach the original trichotomous state of opinion o such that all alternatives in $W(o) \cup U(o)$ are ranked according to their Borda opinion scores by the relation $f(o)$.

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 by the relation $f(o)$.

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 f^B 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 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^1 by letting, for all $o \in \mathcal{O}$, $(x, y) \in P(f^1(o))$ for all $x, y \in X$ such that $b(x; o) < 0$ and $b(y; o) > 0$, and*

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

for all remaining $x, y \in X$. The opinion aggregator f^1 violates opinion transitivity and satisfies the remaining three axioms.

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

$$(x, y) \in f^2(o) \ \Leftrightarrow \ b'(x; o) \geq b'(y; o)$$

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

Example 3 *The opinion aggregator of the final example declares all alternatives to be pairwise indifferent, independent of the state of opinion; this leads to violations of opinion monotonicity. Let $f^3(o) = X^2$ for all $o \in \mathcal{O}$, that is, f^3 assigns the universal indifference relation to all states of opinion. The opinion aggregator f^3 violates opinion monotonicity and satisfies the remaining axioms.*

We now turn to a characterization of the majority opinion aggregator. First, note that the relation $f^M(o)$ is complete for all $o \in \mathcal{O}$ but it is not always transitive. To characterize the majority opinion aggregator, we have to impose completeness of the social relation because it does not follow from the remaining axioms. The empty state of opinion is already covered by the definition of an opinion aggregator so that it can be excluded in the following definition.

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

The remaining two axioms parallel the corresponding properties used by May (1952) in his characterization of the majority rule; see also May (1953) and Sen (1970, Theorem

5*1). The first of these two is an adaptation of the well-established neutrality property, a strengthening of Arrow's (1951/1963/2012) independence of irrelevant alternatives.

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

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

Our final axiom is the responsiveness property phrased in terms of an opinion aggregator.

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

Theorem 2 *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

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

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 strict preference applies; the proof is the same for the latter. Let $o' \in \mathcal{O}$ be such that $(x,y) \in o'$ if and only if $(y,x) \in o$ and $(y,x) \in o'$ if and only if $(x,y) \in o$. Moreover, for all other pairs $(u,v) \in D \setminus \{(x,y),(y,x)\}$, let $(u,v) \in o'$ if and only if $(u,v) \in o$. By construction, it follows that $m((u,v);o') = m((u,v);o)$ 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 $m((x, y); o) > m((y, x); o)$ by definition. Let $o' \in \mathcal{O}$ be such that $m((x, y); o') = m((y, x); o)$ and $m((u, v); o') = m((u, v); o)$ for all $(u, v) \in D \setminus \{(x, y)\}$. 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 2 is established by means of the following three examples.

Example 4 *The opinion aggregator of this example replaces some instances of indifference 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 the remaining two axioms.

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 $x, y \in X$,*

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

and

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

The opinion aggregator f^5 violates opinion neutrality and satisfies the other two axioms.

Example 6 *Finally, in analogy to the monotonicity property of the previous characterization result, the opinion aggregator f^3 that assigns the universal indifference relation to all states of opinion violates opinion responsiveness and satisfies the two remaining axioms.*

4 Concluding remarks

A common feature of the two 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. As a consequence, these rules are well-suited to be analyzed in the context of opinion aggregation.

Our two results 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.

References

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

Arrow, K.J. (1951/1963/2012), Social Choice and Individual Values. 1^{st} ed. (1951), Wiley, New York. 2^{nd} ed. (1963), with Notes on the theory of social choice, Wiley, New York. 3^{rd} ed. (2012), with Foreword to the third edition by Eric Maskin, Yale University Press, New Haven.

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.

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.

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.

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.

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

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

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.

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.

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.

May, K.O. (1953), A note on the complete independence of the conditions for simple majority decision. Econometrica, 21, 172–173.

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.