



HAL
open science

The Borda Class

Zoi Terzopoulou, Ulle Endriss

► **To cite this version:**

Zoi Terzopoulou, Ulle Endriss. The Borda Class. Journal of Mathematical Economics, 2021, 92, pp.31-40. hal-04809479

HAL Id: hal-04809479

<https://hal.science/hal-04809479v1>

Submitted on 28 Nov 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

The Borda Class

An Axiomatic Study of the Borda Rule on Top-Truncated Preferences

Zoi Terzopoulou*, Ulle Endriss

*Institute for Logic, Language and Computation
University of Amsterdam*

Abstract

The Borda rule, originally defined on profiles of individual preferences modelled as linear orders over the set of alternatives, is one of the most important voting rules. But voting rules often need to be used on preferences of a different format as well, such as top-truncated orders, where agents rank just their most preferred alternatives. What is the right generalisation of the Borda rule to such richer models of preference? Several suggestions have been made in the literature, typically considering specific contexts where the rule is to be applied. In this work, taking an axiomatic perspective, we conduct a principled analysis of the different options for defining the Borda rule on top-truncated preferences.

Keywords: Social Choice Theory, Voting, Borda, Axiomatic Method

JEL Classification Code: D71. Microeconomics – Analysis of Collective Decision Making – Social Choice

1. Introduction

The Borda rule, introduced by Jean-Charles de Borda in 1784 for profiles of ballots that are linear orders over a set of alternatives, is a very well-established voting rule. Borda's rule is very intuitive when the agents submit a linear order over m alternatives: it prescribes that, for each voter, $m - 1$ points are to be given to her top alternative, $m - 2$ points to her second-to-top alternative, and so forth, with 0 points given to the alternative ranked last. The alternatives with the largest sum of points across all voters are then announced the winners of the election. But not every voter can be expected to always rank all alternatives she may be presented with. For example, in an election with hundreds of candidates belonging to many different parties, voters may be able to fully rank only a subset of those candidates, possibly those that come from parties to which the voters are sympathetic and have paid more attention. So, what is the right generalisation of the Borda rule to richer models of preference, beyond linear orders?

*Corresponding author
Email address: z.terzopoulou@uva.nl (Zoi Terzopoulou)

In this paper we address this question for domains of preferences that are *top-truncated*. A preference is called top-truncated if it consists of a linear order over a subset of all given alternatives, with the implicit assumption that all non-ranked alternatives are inferior to all ranked alternatives. Top-truncated preferences provide a sensible model for many real-world applications. From choosing the members of a parliament to selecting favourite movies to add to a watch-list, an agent is likely to recognise her most preferred alternatives more easily, and be willing to put effort into ranking them. On the other hand, it is safe to assume that a voter, when overloaded with an abundance of options, will leave her least preferred alternatives unranked in order to escape some mental burden—this can be interpreted either as the voter being indifferent between these alternatives or as not having compared them at all.

Some suggestions regarding appropriate generalisations of the Borda rule for top-truncated preferences have already been made in the literature. All of them are reasonable at first sight, but heavily depend on the interpretation of the preference domain, and of the rule, we have in mind. For instance, given a preference with k ranked alternatives out of a total of m alternatives overall, should the points assigned to the unranked alternatives at the bottom be $m - k - 1$ (as if they all were ranked at level $k + 1$ from the top), or should the points be 0 (as if all unranked alternatives were ranked at the very lowest level m)? Dummett (1997), Saari (2008), Baumeister et al. (2012), Cullinan et al. (2014), Caragiannis et al. (2015), and Terzopoulou and Endriss (2019) have presented their own versions of the Borda rule (some of which coincide), for several variants of domains of preorders.¹ Emerson (2013) has informally discussed the advantages and disadvantages of different such generalisations of the Borda rule, concentrating on issues related to strategic behaviour. Nonetheless, no systematic analysis has been conducted so far regarding the various versions of the Borda rule for top-truncated preferences that appear in the literature to date. In this paper we attempt to close this gap, by identifying axioms (both established and original ones) that characterise each specific rule of interest. Clearly, within the model of top-truncated preferences, agents still have the freedom to rank all alternatives and submit a profile of linear orders—in such a case, we simply apply the traditional Borda rule.

Our axioms provide a principled way of understanding not only the differences, but also the similarities of the suggested generalisations of the Borda rule. On the one hand, we see that all these rules exhibit an analogous structure, which locates them within a class of positional scoring rules that conform to a “Borda style”. We call this class the *Borda class*. The axiomatic characterisation of the Borda class is based on the original characterisation of the Borda rule for domains of linear orders by Young (1974). In his characterisation, Young used the following four axioms: *neutrality*, *reinforcement*, *faithfulness*, and *cancellation*, with the last being the most critical one for the identification of the Borda rule. Analogously, essential for the characterisation of the Borda class is our axiom of *top-cancellation*. Top-cancellation extends Young’s cancellation axiom to top-truncated orders, by requiring that the voting rule should not distinguish between alternatives that are strictly ranked by all agents and tie in pairwise comparisons. On the other hand, each specific rule within the Borda class satisfies

¹A preorder is a reflexive and transitive binary relation.

distinctive properties, which are brought to the surface when expressed formally as axioms. For instance, we show that by enforcing Young’s original cancellation axiom within the Borda class, we can specify the rule that assigns to each non-ranked alternative in a top-truncated order the average of the Borda scores that this alternative could be assigned with if the given top-truncated order were to be extended to a linear order—this is a method that Ackerman et al. (2013) call *bucket averaging*.² Two additional concrete rules in the Borda class are obtained when we impose two axioms that are reminiscent of monotonicity conditions.

Although our work is tightly connected to the characterisation of the Borda rule by Young (1974), other characterisations of the Borda rule in the same formal setting have also been produced by Farkas and Nitzan (1979) and by Saari (1990). The former have used the axiom of Pareto stability based on a notion of relative unanimity, while the latter has employed weaker versions of Young’s axioms and has incorporated the axiom of anonymity as well. We also note that aggregation processes based on the Borda scores, together with their corresponding axiomatic properties, have received much attention in several settings beyond voting as well. Nitzan and Rubinstein (1981) have characterised the Borda rule as a social welfare function (i.e., a function that outputs social rankings instead of winning alternatives). Duddy et al. (2016) and Brandl and Peters (2019) have focused on aggregation mechanisms that produce collective dichotomous preferences and are inspired by Borda’s form of scoring. Lastly, Dietrich (2014) has introduced a judgment aggregation rule that reduces to Borda’s voting rule when applied to the appropriate domain.

The remainder of this paper is organised as follows. Section 2 introduces our basic voting model, together with our notation and terminology. Section 3 reviews the relevant definitions of previous literature for rules that generalise the Borda rule to top-truncated preferences. It also builds important technical connections between these rules. Section 4 contains our main results, namely the axiomatic characterisations of three specific rules that extend Borda’s rule to top-truncated domains, together with the characterisation of the larger class to which all these rules belong. Section 5 concludes.

2. The Model

We have a finite set of *alternatives* A with $\#A = m \geq 3$ and a set of potential (infinite, but countably many) *agents* \mathbb{N} , denoting all agents that may ever participate in an election. Then, in every concrete election, a finite group of agents $N \subset \mathbb{N}$ express their preferences over the set of alternatives A . Such a *preference* \succsim can in general take the form of any preorder on A . We write \mathcal{D} for a *domain* of preferences over A (for example, \mathcal{D} may be the set of all preorders, or a subset of those).

Given two alternatives $a, b \in A$, we write $a \sim b$ when $a \succsim b$ and $b \succsim a$, and $a > b$ when $a \succ b$ and it is not the case that $b \succ a$. When $a \sim b$, then a and b are said to be *indistinguishable*. When $a > b$ (or $a \succ b$), then a is said to be *strongly* (or *weakly*) *preferred to* b . When none of the above holds, then a and b are said to be *incomparable*.

²This result is also in line with the work of Cullinan et al. (2014), who have characterised this specific version of the Borda rule for domains of partial orders, relying on the four classical axioms of Young.




Figure 1: Types of top-truncated preferences in two domains: \mathcal{D}_1 and \mathcal{D}_2 .

For a preference \succsim and two alternatives $a, b \in A$, we write $\succsim^{(ab)}$ for the new preference that is identical to \succsim except for having the positions of a and b switched.

Let us denote by \succsim_i the individual preference of agent i . We are particularly interested in *top-truncated* preferences. A preference \succsim is top-truncated if it strictly ranks a subset $A' \subseteq A$ of the alternatives and requires that all other alternatives are less desirable than those in A' . Note that every agent i is allowed to rank a subset of the alternatives that is possibly different in size from the relevant sets ranked by her peers, and agents may also rank subsets of alternatives with an empty intersection.

In this paper we specifically consider two domains of top-truncated preferences, \mathcal{D}_1 and \mathcal{D}_2 . In each of our two domains, all preferences take the same form, which is one of those depicted in Figure 1 (where transitive arrows are omitted for simplicity). In Figure 1 an arrow captures the fact that the alternative appearing in the position that the arrow starts from is strictly preferred to the alternative in the position that the arrow ends in. Specifically, \mathcal{D}_1 contains top-truncated preferences where k alternatives are ranked, for any $0 \leq k \leq m$, and the alternatives that are not ranked are indistinguishable from each other; \mathcal{D}_2 also contains top-truncated preferences where k alternatives are ranked, for any $0 \leq k \leq m$, but the alternatives that are not ranked are incomparable to each other. Note that both \mathcal{D}_1 and \mathcal{D}_2 are restricted domains of preorders.

A top-truncated preference \succsim consists of two parts: the *top* part and the *bottom* part. Let us define $TOPsets(\succsim)$ as the collection of all subsets $A' \subseteq A$ that contain strictly ordered alternatives in \succsim that are superior, according to \succsim , to all alternatives not in A' . Formally,

$$TOPsets(\succsim) = \{A' \subseteq A \mid (i) x \succ z \text{ for all } x \in A', z \in A \setminus A' \text{ and} \\ (ii) x \succ y \text{ or } y \succ x \text{ for all } x, y \in A'\}.$$

Then, $TOP(\succsim)$ is the unique largest set in $TOPsets(\succsim)$, and $BOT(\succsim) = A \setminus TOP(\succsim)$.

Next, given a group of agents N of size $\#N = n$, we suppose without loss of generality that $N = \{1, \dots, n\}$. Now a *profile* of preferences for the group N can be captured by a vector $\succsim = (\succsim_1, \dots, \succsim_n) \in \mathcal{D}^n$. Given such a profile \succsim , we denote by \succsim_{-i} the new profile where all agents besides i submit the same preferences as in \succsim .




Figure 2: The top alternatives of each single preference are in red. From left to right, these are the sets $\{a_1, a_2, a_3\}$, $\{a_1, \dots, a_5\}$, and $\{a_1, a_2\}$, respectively. The $TOPsets(\succsim)$ of the leftmost preference \succsim are $\{\{a_1\}, \{a_1, a_2\}, \{a_1, a_2, a_3\}\}$, and similarly for the other preferences. The intersection of these sets is then the set $\{a_1, a_2\}$, which includes the top alternatives of the whole profile.

Also, given a different profile $\succsim' = (\succsim_{n+1}, \dots, \succsim_{n+\ell}) \in \mathcal{D}^\ell$ with $\ell \geq 1$, that concerns a disjoint set of voters than those of \succsim , we write

$$(\succsim, \succsim') = (\succsim_1, \dots, \succsim_n, \succsim_{n+1}, \dots, \succsim_{n+\ell}) \in \mathcal{D}^{n+\ell}.$$

The set of top alternatives in a profile \succsim includes those alternatives that are on top for all agents. That is, $TOP(\succsim)$ is defined as the unique largest set in $\bigcap_{i \in N} TOPsets(\succsim_i)$ if that intersection is non-empty; otherwise, $TOP(\succsim)$ is the empty set. We also define $BOT(\succsim)$ as the set $A \setminus TOP(\succsim)$, but note that this definition is intuitively meaningful (in the sense of characterising a set of “bottom alternatives”) only when $TOP(\succsim) \neq \emptyset$. See Figure 2 for an example with $m = 5$ and $n = 3$.

Given a domain of preferences \mathcal{D} over the set of alternatives A , a *voting rule* is a function that maps every possible preference profile (for any group of agents $N \subset \mathbb{N}$) to a nonempty subset of the alternatives, which is the set of (tied) winners. A *positional scoring rule* F_s is a voting rule associated with some positional scoring function s . Here a *scoring function* $s : A \times \mathcal{D} \rightarrow \mathbb{R}$ maps every alternative $a \in A$ in a preference $\succsim \in \mathcal{D}$ to a score, which is a real number that we denote by $s_{\succsim}(a)$. Moreover, a scoring function s is *positional* if, for all permutations $\sigma : A \rightarrow A$, all preferences $\succsim \in \mathcal{D}$, and all alternatives $x \in A$, it holds that $s_{\succsim}(x) = s_{\sigma(\succsim)}(\sigma(x))$, where $\sigma(\succsim) = \{(\sigma(x), \sigma(y)) \mid x \succ y\}$. Intuitively, a positional scoring function assigns scores to *positions* in a graph, rather than to specific alternatives in specific preferences. Finally, the corresponding rule F_s is defined as follows: for every profile $\succsim = (\succsim_1, \dots, \succsim_n)$,

$$F_s(\succsim) = \operatorname{argmax}_{x \in A} \sum_{i \in N} s_{\succsim_i}(x).$$

In words, a positional scoring rule selects as winning alternatives those with the largest score over all individual preferences. We will often also use the abbreviation $s_{\succsim}(a) = \sum_{i \in N} s_{\succsim_i}(a)$.

3. Generalising the Borda Rule

The Borda rule is commonly defined on domains of linear orders in one of two ways. First, as a positional scoring rule with score-vector $(m-1, m-2, \dots, 0)$, where

$$\begin{array}{ll}
B_i(a) = \sum_{y \in A} \mathbb{1}_{a >_i y} - \mathbb{1}_{y >_i a}. & B_i(a) = \sum_{y \in A} \mathbb{1}_{a >_i y}. \\
\text{(a) the symmetric Borda score as for linear orders} & \text{(b) based on the alternatives strictly dominated by } a \\
\\
B_i(a) = \sum_{y \in A} \mathbb{1}_{a \succeq_i y}. & B_i(a) = - \sum_{y \in A} \mathbb{1}_{y >_i a}. \\
\text{(c) based on the alternatives weakly dominated by } a & \text{(d) based on the alternatives not strictly dominating } a
\end{array}$$

Figure 3: Possible domination-based scores for defining the Borda rule on preorders.

the first position in the vector corresponds to the score assigned to the top alternative in a linear order, the second position to the second-to-top alternative, and so forth, until the last position in the vector that is associated with the score of the bottom alternative. A second way of defining the Borda rule is in terms of the *weighted majority graph*, where the winning alternatives are those that maximise the following function:³

$$\begin{aligned}
B(a) &= \sum_{y \in A} \#\{i \in N \mid a >_i y\} - \#\{i \in N \mid y >_i a\}. \\
&= \sum_{i \in N} \sum_{y \in A} \mathbb{1}_{a >_i y} - \mathbb{1}_{y >_i a}.
\end{aligned}$$

We can think of $B(a)$, the *symmetric Borda score* of a , as $B(a) = \sum_{i \in N} B_i(a)$ with $B_i(a) = \sum_{y \in A} \mathbb{1}_{a >_i y} - \mathbb{1}_{y >_i a}$.

It seems sensible to presuppose that any interesting generalisation of the Borda rule will also be defined in terms of a scoring function B with $B(x) = \sum_{i \in N} B_i(x)$ for some scoring functions B_i that each only makes reference to the preference of one agent i . In Figure 3 we present some options for how one could define such a function on general domains of preorders—all of them reduce to the standard Borda rule when we restrict ourselves to profiles of linear orders. Note also that any positive affine transformation of a function $B_i(x)$ induces the same rule as $B_i(x)$ itself.

When indistinguishability is not materialised in a domain, definition (b) coincides with definition (c); when incomparability is not materialised, definition (c) coincides with definition (d).

Specifically regarding top-truncated preferences, some generalisations of the Borda rule have already been discussed in previous work. These all follow the definition of a positional scoring rule in terms of score-vectors. Note that for domains of top-truncated preferences we actually need $m - 1$ such vectors, one for each possible number of the top alternatives in a preference (when this number is 0, all alternatives will always be assigned the same score, by the definition of positional scoring functions). Now, the last $m - k$ positions in a vector correspond to the scores associated with the bottom alternatives in a top-truncated preference with k top alternatives. Note that all bottom alternatives must have the same score, by the definition of positional scoring functions. Generally, given a top-truncated preference \succeq with $\#TOP(\succeq) = k \leq m - 1$, we will write s_j for the score of the alternative ranked in the j th position within the top part of \succeq and

³The definition of the Borda rule based on the weighted majority graph can also be considered a definition of a positional scoring rule on linear orders, with score-vector $(m, m - 2, \dots, -m)$.

s_{k+1} for the score of all alternatives in the bottom part of \succsim . These are the versions of the Borda rule for top-truncated preferences that can be found in the literature:

- *Pessimistic Borda*⁴ (Baumeister et al., 2012):

$$(m-1, m-2, \dots, m-k, 0, \dots, 0), \text{ for all } 1 \leq k \leq m.$$

- *Optimistic Borda*⁵ (Saari, 2008; Baumeister et al., 2012):

$$(m-1, m-2, \dots, m-k, m-k-1, \dots, m-k-1), \text{ for all } 1 \leq k \leq m.$$

- *Averaged Borda* (Dummett, 1997):

$$(m-1, m-2, \dots, m-k, \frac{m-k-1}{2}, \dots, \frac{m-k-1}{2}), \text{ for all } 1 \leq k \leq m.$$

We have reviewed two different—yet equally natural—directions one could follow to generalise the Borda rule on preorders (and specifically on truncated preferences). We will next see that defining the Borda rule using a domination-based score or a scoring vector can lead to exactly the same outcome. What plays a crucial role here is the particular domain we consider. Specifically, by combining domination-based scores with different domains, we obtain already existing rules (consult Table 1).⁶ The observations included in Table 1 are quite straightforward, except perhaps for the one concerning the averaged Borda rule and the symmetric Borda scores, made explicit in Proposition 1.⁷

Proposition 1. *The averaged Borda rule for top-truncated preferences (on \mathcal{D}_1 or \mathcal{D}_2) is the positional scoring rule with a corresponding scoring function based on the symmetric Borda scores.*

Proof. Take the symmetric Borda scores of the alternatives in an arbitrary top-truncated preference. Divide all these scores by 2 and add to all of them the positive constant $\frac{m}{2}$. The scores we obtain are exactly the ones associated with the averaged Borda rule. Since the new scores are obtained through an affine transformation of the old ones, the two corresponding rules are equivalent. \square

The pessimistic, the optimistic, and the averaged Borda rules, all belong to a wider class of natural generalisations of the classical Borda rule for top-truncated preferences, to

⁴This rule can also be found under the name *truncated Borda* (Ackerman et al., 2013).

⁵This rule has also been referred to as *modified Borda* by Caragiannis et al. (2015).

⁶In our earlier work (Terzopoulou and Endriss, 2019) we defined yet another generalisation of the Borda rule, called *shortsighted Borda*, which is based on the scoring vector $(m-1, m-2, \dots, m-k, m-k, \dots, m-k)$. Note that this rule cannot be obtained through domination-based scores, for these particular domains. Indeed, shortsighted Borda is better suited to domains where the unranked alternatives are taken to be incomparable (rather than inferior) to the ranked alternatives.

⁷Proposition 1 also holds for larger preference domains, like the one of weak orders, as discussed for instance by Duddy et al. (2016).

Scores \ Domains	\mathcal{D}_1	\mathcal{D}_2
(a) symmetric	averaged	averaged
(b) domination	pessimistic	pessimistic
(c) weak domination	optimistic	pessimistic
(d) non-domination	optimistic	optimistic

Table 1: Borda rules for domains of top-truncated preferences.




Figure 4: A rule in the Borda class on the domain \mathcal{D}_1 , for $s_{k+1} < m - k$.

which we will refer as the *Borda class*. Formally, a voting rule on domains of top-truncated preferences is in the Borda class if it is induced by some positional scoring function that gives rise to the scoring vectors

$$(m-1, m-2, \dots, m-k, s_{k+1}, \dots, s_{k+1}),$$

for some number $s_{k+1} < m - k$, and for $1 \leq k \leq m$ (see Figure 4 for an illustration). A further reasonable choice for s_{k+1} would be that it is any number between 0 and $m - k - 1$. In fact, the pessimistic, the optimistic, and the averaged Borda rules also belong to this smaller class.

Note also that the three aforementioned rules produce different outcomes even for very simple profiles. Figure 5 provides an example with three agents and four alternatives.

4. Axiomatic Characterisations

In this section we introduce the axiomatic properties that characterise the rules in the Borda class and we discuss how these relate to the properties in the classical axiomatisation of the Borda rule for domains of linear orders by Young (1974). We then provide axiomatic characterisations of the three specific generalisations of the Borda rule for top-truncated preferences we reviewed in Section 3: the optimistic, the pessimistic, and the averaged Borda rule.

Our analysis will proceed in steps, each of which will concern a smaller class of rules than the previous one: First, we are going to see an axiomatic characterisation of




Figure 5: A profile where the pessimistic, the optimistic, and the average Borda rule output different sets of winners: $\{b\}$, $\{a, c\}$, and $\{a, b\}$, respectively.

all positional scoring rules for top-truncated preferences. Then, we are going to add some new axioms to that characterisation and obtain all positional scoring rules in the Borda class. Finally, by considering a few further axioms, we will be able to restrict attention to our specific rules of interest within the Borda class.

Smith (1973) and Young (1975) initiated the axiomatic analysis of positional scoring rules for the special case of profiles of linear orders. Two decades later, Myerson (1995) generalised the previous results to profiles of ballots that could take any form over the set of alternatives. Recently, Kruger and Terzopoulou (2020) provided an analogous characterisation of positional scoring rules for acyclic preferences without indistinguishability (but it is not hard to see that their relevant proofs go through for our domains of top-truncated preferences as well). The axioms shared by all these characterisations for positional scoring rules (here stated for domains of top-truncated preferences) are *anonymity*, *neutrality*, *reinforcement*, and *continuity*.


Anonymity requires that the outcome of the aggregation should not depend on the names of the agents but only on the preferences they contribute to a given profile; neutrality says that all alternatives should be treated symmetrically; reinforcement prescribes that if two groups unite and vote together, then the alternatives that win should be the alternatives that would win for both elections if each group were to vote separately (unless there are no such alternatives); finally, continuity states that a sufficiently large number of agents should be able to change the outcome in accordance with their preference. Let us now define formally what it means for a rule F to satisfy each of these four axioms.

- *Anonymity*: For any permutation $\sigma : \mathcal{N} \rightarrow \mathcal{N}$, it is the case that $F(\succsim_1, \dots, \succsim_n) = F(\succsim_{\sigma(1)}, \dots, \succsim_{\sigma(n)})$.
- *Neutrality*: For any permutation $\sigma : A \rightarrow A$, it is the case that $F(\succsim_1, \dots, \succsim_n) = F(\sigma(\succsim_1), \dots, \sigma(\succsim_n))$, where $\sigma(\succsim) = \{(\sigma(x), \sigma(y)) \mid x \succ y\}$.
- *Reinforcement*: For any two profiles \succsim, \succsim' with $F(\succsim) \cap F(\succsim') \neq \emptyset$, it is the case that $F(\succsim, \succsim') = F(\succsim) \cap F(\succsim')$.
- *Continuity*: For any two profiles \succsim and \succsim' , there exists a natural number K such that, for every natural number k that is greater than K , it is the case that $F(\underbrace{\succsim, \dots, \succsim, \succsim'}_k) \subseteq F(\succsim)$.

Theorem 1 is an easy adaptation of the result of Kruger and Terzopoulou (2020), mentioned above.

Theorem 1. *A voting rule for top-truncated preferences (on \mathcal{D}_1 or \mathcal{D}_2) is a positional scoring rule if and only if it satisfies anonymity, neutrality, reinforcement, and continuity.*

Although positional scoring rules are nicely defined for top-truncated preferences, they are not the only natural rules for these domains. For instance, consider a majority-based rule that selects as winners those alternatives that are considered superior to the largest number of other alternatives by a strict majority of agents, which may be regarded as a generalisation of the well-known Copeland rule (Zwicker, 2016). Then, suppose we have the following profile with three agents:



In the above profile the winning set is $\{a, b, c\}$, since for each of these alternatives there exist two agents that rank it higher than exactly three other alternatives.

It is not hard to see that the reinforcement axiom is violated by this rule: add to the given profile the single-agent profile consisting of the complete preference $a > b > c > d > e$, where our rule would output alternative a as the unique winner. In the combined profile, reinforcement demands that a is again the only winner. However, a and b are both considered superior to three other alternatives by a strict majority of agents and thus must both win.

4.1. Characterising the Borda Class

We are now going to provide two distinct characterisations of the Borda class of rules in our top-truncated preference domains. Recall that the Borda class only includes positional scoring rules for which it is the case that, if a is strongly preferred to b (i.e., $a > b$), then the score of a is larger than the score of b in \succeq (i.e., $s_{\succeq}(a) > s_{\succeq}(b)$). This is formally imposed by a *monotonicity* property, intuitively prescribing that moving an alternative to a “higher” position in a preference is beneficial to that alternative. Formally, we can define monotonicity of a rule F as follows.

- *Monotonicity:* Consider any preference \succeq_i within any profile \succeq and two alternatives a, b such that $a >_i b$. If $b \in F(\succeq)$, then $F(\succeq_{-i}, \succeq_i^{(ab)}) = \{b\}$.

Lemma 1 makes the link between monotonicity of a positional scoring rule F_s and the scoring function s defining that rule precise. Note also that this statement (here only described for the two top-truncated domains) holds for wider preference domains as well, like any kind of preorders.

Lemma 1. *A positional scoring rule F_s for top-truncated preferences (on either \mathcal{D}_1 or \mathcal{D}_2) satisfies monotonicity if and only if*

$$s_{\succ}(a) > s_{\succ}(b) \text{ whenever } a > b.$$

Proof. To see that monotonicity is satisfied by any positional scoring rule for which the required condition on the scores holds, note that by flipping the positions of a and b in \succ , the score of b increases, the score of a decreases, and all other scores remain the same (by definition of the positional scoring rule). Thus, if b was among the winners before the flipping, then afterwards it will be the unique winner with the highest score.

For the other direction, given a preference \succ and alternatives a, b with $a > b$, we can construct a profile \succsim where a and b are amongst the winners under the rule F_s , with arbitrarily large score:

Consider a preference \succ^* and two positions in \succ^* such that the score of an alternative in the first position is at least as large as the score of an alternative in the second position, which is at least as large as the score of the alternatives in all other positions. Then, consider a preference \succ' where a takes the first position in \succ^* while b takes the second position in \succ^* , and a preference \succ'' where b takes the first position in \succ^* while a takes the second position in \succ^* . With sufficiently many (and equally many) copies of these preferences, we ensure that we have a profile \succsim^* where a and b win with arbitrarily large score. Define the profile $\succsim = (\succ^*, \succ, \succ^{(ab)})$ and call i the agent that submits the preference \succ in \succsim . Then, for alternative b to be the unique winner in the profile $(\succ_{-i}, \succ^{(ab)})$, the required condition on the scores must be satisfied. \square

Our first characterisation theorem is in line with the characterisation of the Borda rule by Young (1974), who—informally speaking—identified the Borda rule as the unique scoring rule that satisfies a *cancellation* property. In this paper, we examine an axiom that is similar in flavour to Young’s cancellation, but applies specifically to domains of top-truncated preferences. *Top-cancellation* concerns preference profiles \succsim with $TOP(\succsim) \neq \emptyset$. In such profiles, if the preferences of the agents between the top alternatives “cancel” each other, then no alternative can be considered better than the others in the top set, and hence all alternatives in that set should be treated the same by the voting rule. Note that for domains of linear orders (that are special domains of top-truncated preferences), the top-cancellation axiom reduces to the standard cancellation axiom of Young.

- *Top-cancellation:* Consider any profile \succsim with $TOP(\succsim) \neq \emptyset$. If for all alternatives $x, y \in TOP(\succsim)$ with $x \neq y$ it is the case that

$$\#\{i \in N \mid x >_i y\} = \#\{i \in N \mid y >_i x\},$$

then $TOP(\succsim) \subseteq F(\succsim)$ or $TOP(\succsim) \cap F(\succsim) = \emptyset$.

Note that top-cancellation allows for the case where no top alternative of a profile belongs to the winning set (i.e., where $TOP(\succsim) \cap F(\succsim) = \emptyset$). Indeed, top-cancellation is a weak axiom, only ensuring that all top alternatives will be treated symmetrically when appropriate conditions hold—that these alternatives should also be among the winners is a separate intuitive requirement, which we will later take care of with the axiom of monotonicity.

Theorem 2. *A voting rule for top-truncated preferences (on either \mathcal{D}_1 or \mathcal{D}_2) is in the Borda class if and only if it satisfies anonymity, neutrality, reinforcement, continuity, monotonicity, and top-cancellation.*

Proof. That all rules in the Borda class satisfy these properties is easy to see. For the other direction, we know that a voting rule F that satisfies anonymity, neutrality, reinforcement, and continuity must be a positional scoring rule by Theorem 1. Let us call this rule F_s . Moreover, F_s , satisfying neutrality, reinforcement, monotonicity, and top-cancellation, when restricted to profiles of preferences that are linear orders, must reduce to the Borda rule (Young, 1974).⁸

Now, consider an arbitrary top-truncated preference \succsim (that is not a linear order), and a linear order L that extends \succsim' (i.e., $\succsim' \subset L$), where \succsim' is the same as \succsim besides having its top alternatives placed in the reverse order. Figure 6 presents an illustration for the domain \mathcal{D}_1 , where $TOP(\succsim) = \{a_1, \dots, a_{m-k}\}$ and $BOT(\succsim) = \{a_{m-k+1}, \dots, a_m\}$. Construct the profile $\mathbf{z} = (\succsim, L)$. Observe that $TOP(\mathbf{z}) = TOP(\succsim) = TOP(L)$. From monotonicity and Lemma 1, we get $F_s(\mathbf{z}) \subseteq TOP(\succsim)$, while top-cancellation entails that $TOP(\succsim) \subseteq F(\mathbf{z})$ or $TOP(\succsim) \cap F(\mathbf{z}) = \emptyset$. Thus, $F_s(\mathbf{z}) = TOP(\succsim)$.

We know that the scores of the alternatives in L will be Borda-like. Moreover, the scores of all alternatives in $TOP(\mathbf{z})$ have to be the same. So, we have that

$$s_1 + m - k = s_2 + m - k + 1 = \dots = s_{k-1} + m - 2 = s_k + m - 1.$$

Equivalently,

$$s_{\succsim}(a_1) = s_{\succsim}(a_2) + 1 = s_{\succsim}(a_3) + 2 = \dots = s_{\succsim}(a_k) + k - 1.$$

We conclude that F_s is in the Borda class. □

Theorem 2—together with the characterisation of positional scoring rules presented in Theorem 1—implies the following corollary.

Corollary 1. *A positional scoring rule for top-truncated preferences (on either \mathcal{D}_1 or \mathcal{D}_2) is in the Borda class if and only if it satisfies monotonicity and top-cancellation.*

An immediate question that arises is whether the axioms appearing in Theorem 2 are all necessary for the characterisation result, i.e., whether they are independent. We know from the characterisation of positional scoring rules that anonymity, neutrality, reinforcement, and continuity are independent, and it is easy to see that monotonicity does not break this fact. Adding top-cancellation also preserves independence between these axioms. Proposition 2 states exactly this, and a proof for the most interesting case (i.e., that top-cancellation, together with the rest of our relevant axioms, does not imply anonymity) is provided. Note that this is a rather unexpected result, since—as we will see later too—in the original proof of Young (1974) the analogous cancellation axiom (in combination with the other axioms) ends up implying anonymity.

⁸More specifically, Young's axiomatisation of the Borda rule relies on a *faithfulness* axiom, demanding that in any single-agent profile the winning alternative is the one on the top of that agent's preference. Straightforwardly, all monotonic positional scoring rules satisfy this property.




Figure 6: A top-truncated preference \succsim and a linearisation L of it with the order of its top alternatives reversed. We have that $s_{\succsim}(a_1) = s_1, \dots, s_{\succsim}(a_k) = s_k, s_{\succsim}(a_{k+1}) = \dots = s_{\succsim}(a_m) = s_{k+1}$ and $s_L(a_k) = m-1, \dots, s_L(a_1) = m-k, s_L(a_{k+1}) = m-k-1, \dots, s_L(a_m) = 0$.

Proposition 2. *The axioms of anonymity, neutrality, reinforcement, continuity, monotonicity, and top-cancellation are logically independent on domains of top-truncated preferences (\mathcal{D}_1 and \mathcal{D}_2).*

Proof. Let us show the most interesting case, i.e., that there exists a voting rule on top-truncated preferences that satisfies neutrality, reinforcement, continuity, monotonicity, and top-cancellation, but violates anonymity (analogous counterexamples can be easily found for all other combinations of our axioms as well).

Consider the voting rule F on top-truncated preferences that works just like the optimistic Borda rule, but with a small exception: when agent 1 reports a preference \succsim^* that identifies some alternative $a \in A$ (whichever that a is) as the unique top one (that is, a preference $\succsim^* = \{(a, x) \mid x \in A \setminus \{a\}\}$), then agent 1 gets assigned double the weight of the other agents (i.e., the scores associated with her preference are twice the standard scores of the optimistic Borda rule). Obviously, this rule is not anonymous.

Neutrality is satisfied by F because the definition of the rule does not distinguish between the names of the alternatives. Monotonicity is also satisfied, because it is always better for an alternative to appear in a higher position in a preference relation. Continuity holds too, since by adding a sufficiently large number of copies of the same profile \succsim' to a given profile \succsim , we can arbitrarily increase the Borda scores of the alternatives that win in \succsim' , and thus obtain the result prescribed by the axiom. To see that top-cancellation is satisfied as well, note that for any profile \succsim with $TOP(\succsim) \neq \emptyset$, if the preference \succsim^* defined above (for some given alternative a) appears in \succsim , it must be the case that $TOP(\succsim) = \{a\}$, and thus $F(\succsim) = \{a\}$ (which means that top-cancellation is vacuously satisfied). Otherwise, F functions as the standard optimistic Borda rule, and hence top-cancellation holds. Finally, it is also easy to see that F satisfies reinforcement: for all profiles \succsim and \succsim' and for any alternative $a \in A$, we have that the score that a receives by the rule F in the joint profile (\succsim, \succsim') will always be the sum of the scores that a receives in \succsim and in \succsim' . \square

Our second characterisation of the Borda class relies on a result of Fishburn and Gehrlein (1976) for domains of linear orders (based on a proof sketch by Smith, 1973), namely that the Borda rule is the only positional scoring rule for which the Condorcet loser (CL) of a profile is never the winner (let us call this property CL-consistency). We remind the reader that the Condorcet loser of a preference profile is an alternative that loses in pairwise comparisons to all other alternatives, where “losing” means that a majority of agents considers that alternative inferior to the one it is compared to. Let us extend this fundamental principle for profiles of linear preferences to profiles of top-truncated preferences by stipulating no alternative that is a Condorcet loser relative to the top part of a profile should ever be amongst the winners.⁹

- *Top-CL-consistency*: Consider any profile \succsim with $TOP(\succsim) \neq \emptyset$. For any $b \in TOP(\succsim)$ such that for all $x \in TOP(\succsim) \setminus \{b\}$,

$$\#\{i \in N \mid x \succ_i b\} > \#\{i \in N \mid b \succ_i x\},$$

it is the case that $b \notin F(\succsim)$.

Theorem 3. *A voting rule for top-truncated preferences (on either \mathcal{D}_1 or \mathcal{D}_2) is in the Borda class if and only if it satisfies anonymity, neutrality, reinforcement, continuity, monotonicity, and top-CL-consistency.*

Proof. The proof uses an analogous argument to that of Smith (1973). □

Theorem 3, together with the characterisation of positional scoring rules for top-truncated preferences (Theorem 1), implies the following corollary.

Corollary 2. *A positional scoring rule for top-truncated preferences (on either \mathcal{D}_1 or \mathcal{D}_2) is in the Borda class if and only if it satisfies monotonicity and top-CL-consistency.*

At this point we also need to examine whether the axioms of Theorem 3 are independent. For example, can we find a positional scoring rule for top-truncated preferences that satisfies top-CL-consistency but is not monotonic? Proposition 3 answers this question in the affirmative.

Proposition 3. *The axioms of anonymity, neutrality, reinforcement, continuity, monotonicity, and top-CL-consistency are logically independent on domains of top-truncated preferences (\mathcal{D}_1 and \mathcal{D}_2).*

Proof. We will show that there exists a positional scoring rule (that is, a rule satisfying anonymity, neutrality, reinforcement, and continuity) for top-truncated preferences that satisfies top-CL-consistency but violates monotonicity. Showing independence for the remaining sets of axioms can be done in an analogous manner.

Consider the positional scoring rule for which, for an any given top-truncated preference \succsim , the Borda-like scores are assigned to the alternatives in $TOP(\succsim)$, and

⁹Observe that this axiom does not require that alternatives in the bottom part of a profile must be barred from winning as well (but this of course would be enforced by imposing monotonicity).

score $m + 1$ is assigned to the alternatives in $BOT(\succsim)$. Clearly, this rule is not monotonic, but top-CL-consistency holds: no Condorcet loser within the top alternatives of a profile can ever win (either because an alternative that Pareto dominates it wins, or because some bottom alternative wins instead). \square

Intuitively, top-CL-consistency is the axiom ensuring that the scores of the top alternatives in an agent's top-truncated preference will be distributed in a linear manner, as required for rules in the Borda class. Note that classical CL-consistency—although applicable to domains of top-truncated preferences as well—is not appropriate for our purposes. In particular, not all rules in the Borda class satisfy CL-consistency.¹⁰ For example, consider the pessimistic Borda rule, and a profile with nine agents and four alternatives such that: four agents rank alternative a on top and every other alternative directly below, and the remaining five agents have preferences as follows: $b \succ_1 c \succ_1 d \succ_1 a$, $b \succ_2 c \succ_2 d \succ_2 a$, $c \succ_3 d \succ_3 b \succ_3 a$, $d \succ_4 b \succ_4 c \succ_4 a$, $d \succ_5 c \succ_5 b \succ_5 a$. Alternative a is the Condorcet loser of this profile, but it will be the winner according to the pessimistic Borda rule (it will receive 12 points, while all other alternatives will only get 10 points).

To sum up, the axioms of top-cancellation and of top-CL-consistency (together with monotonicity) are the ones that distinguish rules in the Borda class from all other positional scoring rules. Interestingly, these two axioms only bite for profiles of top-truncated preferences with a non-empty set of top alternatives. On the one hand, such profiles are rare in general. On the other hand, every preference can appear in some profile of that form. The key idea behind our proofs is that the rules with which we work are positional scoring rules. So, the score assigned to an alternative a in a given preference \succsim will be fixed, and can be deduced by applying the relevant axioms in profiles with a non-empty set of top alternatives that \succsim is part of.

4.2. Characterising Specific Rules in the Borda Class

After having characterised the Borda class of rules via a number of normative axiomatic properties, in the remainder of this section we proceed with identifying those properties that characterise each one of our specific rules of interest, within the Borda class.

We observe that the pessimistic Borda rule is the only rule in the Borda class for which the scores of the bottom alternatives in the top-truncated preferences do not depend on *how many* of these alternatives there are. Loosely speaking, this translates into the following slogan:

The number of other alternatives with which some alternative a shares the bottom position does not affect a 's performance.

The axiom of *bot-indifference* formally captures this idea:

- *Bot-indifference*: Consider any two profiles \succsim and $\succsim' = (\succsim_{-i}, \succsim'_i)$ for some agent i such that the preference \succsim'_i is obtained from the preference \succsim_i by having one of the bottom alternatives of \succsim_i , say alternative a , moved to the last position of the

¹⁰We are grateful to an anonymous reviewer of the JME for this observation.

ranked alternatives in the top part. If $a \notin F(\succsim')$, then for any $b \in \text{BOT}(\succsim_i) \setminus \{a\}$ it is the case that

$$b \in F(\succsim) \text{ if and only if } b \in F(\succsim').$$

Thus, by moving a we create two profiles in which the number of alternatives that b shares a bottom position with changes, and we stipulate that this should not affect whether or not b will be amongst the winners (at least not in case a is not winning in the second profile, the one where it intuitively is put in a better position). Now, as suggested by our earlier observation, bot-indifference characterises the pessimistic Borda rule:

Theorem 4. *The only voting rule for top-truncated preferences (on either \mathcal{D}_1 or \mathcal{D}_2) that satisfies anonymity, neutrality, reinforcement, continuity, monotonicity, top-cancellation (or top-CL-consistency), and bot-indifference is the pessimistic Borda rule.*

Proof. We can easily verify that the pessimistic Borda rule satisfies all relevant axioms. For the other direction, suppose that we have a rule F_s , with corresponding scoring function s , in the Borda class (satisfying anonymity, neutrality, reinforcement, continuity, monotonicity, top-cancellation and top-CL-consistency) for which bot-indifference holds. Take two arbitrary preferences \succsim and \succsim' such that \succsim' is obtained from \succsim by having one of the bottom alternatives of \succsim , namely alternative a , moved to the last position of the ranked alternatives above. We will show that $s_{\succsim}(b) = s_{\succsim'}(b)$ for any alternative $b \in \text{BOT}(\succsim) \setminus \{a\}$. Note that in the special case where \succsim' is a linear order, we know that b is the lowest alternative in \succsim' , and $s_{\succsim'}(b) = 0$. Thus, our proof ensures that the bottom alternatives of all preferences \succsim will be assigned score 0, meaning that F_s must be the pessimistic Borda rule.

Take the alternatives $b \in \text{BOT}(\succsim) \setminus \{a\}$ and $c \in \text{TOP}(\succsim)$ and construct a profile \succsim^* where b and c are the only winners. That is:

$$F(\succsim^*) = \{b, c\}.$$

The profile \succsim^* can be easily constructed by taking equally many copies of two preferences, one with b followed by c making up the top part and one with c followed by b making up the top part (and the remaining alternatives making up the bottom part in both cases). By using sufficiently many copies, we can take the difference in score between b and c on the one hand and the next-best alternative on the other to be arbitrarily large.

Then consider the following profiles:

$$\succsim = (\succsim^*, \succsim, \succsim^{(bc)}) \quad \text{and} \quad \succsim' = (\succsim^*, \succsim', \succsim^{(bc)}),$$

where $F(\succsim) = \{b, c\}$ and $F(\succsim') \subseteq \{b, c\}$. So $a \notin F(\succsim')$, and $\succsim' = (\succsim_{-i}, \succsim'_i)$, with $\succsim_i = \succsim$ and $\succsim'_i = \succsim'$. Hence, we have that

$$s_{\succsim}(b) = s_{\succsim}(c) = s_{\succsim'}(c).$$

But $b \in F(\succsim)$ implies that $b \in F(\succsim')$ by bot-indifference. It must then be the case that $s_{\succsim'}(b) \geq s_{\succsim'}(c) = s_{\succsim}(b)$. So, $s_{\succsim_{-i}}(b) + s_{\succsim'}(b) \geq s_{\succsim_{-i}}(b) + s_{\succsim}(b)$. We conclude that $s_{\succsim'}(b) \geq s_{\succsim}(b)$. Proving the inverse inequality in a symmetric manner, we have that $s_{\succsim'}(b) = s_{\succsim}(b)$. \square

We obtain an immediate corollary:

Corollary 3. *The only voting rule for top-truncated preferences (on either \mathcal{D}_1 or \mathcal{D}_2) in the Borda class that satisfies bot-indifference is the pessimistic Borda rule.*

We next define a new axiomatic property, building on the basic idea that if the dominance relationships between different winning alternatives remain unaltered, then no tie between these alternatives can be broken. In words, *dom-power* suggests that a winning alternative a can *only* break a tie between itself and a different winning alternative b by having its support against b strictly increased.

Note that the optimistic Borda rule is the only rule in the Borda class for which, in any top-truncated preference, the score of the last ranked alternative on top remains the same if that alternative “moves” to the bottom instead. The axiom ensuring this is precisely *dom-power*.

- *Dom-power:* Consider any two profiles \succsim and $\succsim' = (\succsim_{-i}, \succsim'_i)$ such that the preference \succsim'_i is obtained from the preference \succsim_i by having one of the bottom alternatives of \succsim_i moved to the last position of the ranked alternatives in the top part. Then, for any $a \in TOP(\succsim_i)$, it is the case that

$$a \in F(\succsim) \text{ if and only if } a \in F(\succsim').$$

Theorem 5. *The only voting rule for top-truncated preferences (on either \mathcal{D}_1 or \mathcal{D}_2) that satisfies anonymity, neutrality, reinforcement, continuity, monotonicity, top-cancellation (or top-CL-consistency), and dom-power is the optimistic Borda rule.*

Proof. We can easily verify that the optimistic Borda rule satisfies all relevant axioms. For the other direction, suppose that we have a rule F_s in the Borda class, induced by a suitable scoring function s (and satisfying anonymity, neutrality, reinforcement, continuity, monotonicity, top-cancellation and top-CL-consistency) for which dom-power holds. Take two arbitrary preferences \succsim and \succsim' such that \succsim' is obtained from the preference \succsim by having one of the bottom alternatives of \succsim , namely alternative b , moved to the last position amongst the ranked alternatives above. We will show that $s_{\succsim}(b) = s_{\succsim'}(b)$, meaning that F must be the optimistic Borda rule.

Take an alternative $a \in TOP(\succsim)$ and construct a profile \succsim^* , such that

$$F(\succsim^*) = \{a, b\}.$$

The profile \succsim^* can be easily constructed by taking equally many copies of two preferences, one with a followed by b making up the top part and one with b followed by a making up the top part (and the remaining alternatives making up the bottom part in both cases). By using sufficiently many copies, we can take the difference in score between a and b on the one hand and the next-best alternative on the other to be arbitrarily large.

Then consider the following profiles:

$$\succsim = (\succsim^*, \succsim, \succsim^{(ab)}) \quad \text{and} \quad \succsim' = (\succsim^*, \succsim', \succsim^{(ab)}),$$

where $F(\succsim) = \{a, b\}$, $F(\succsim') \subseteq \{a, b\}$, $\succsim_i = \succsim$, $\succsim'_i = \succsim'$, and $\succsim' = (\succsim_{-i}, \succsim'_i)$. So, we have that

$$s_{\succsim}(b) = s_{\succsim}(a) = s_{\succsim'}(a).$$

In addition, since $a \in F(\succsim)$, dom-power implies that $a \in F(\succsim')$, which means that the score of a in the profile \succsim' must be at least as high as the score of b . We hence have that

$$s_{\succsim}(b) = s_{\succsim'}(a) \geq s_{\succsim'}(b).$$

But we know that $s_{\succsim}(b) = s_{\succsim_{-i}}(b) + s_{\succsim}(b)$ and $s_{\succsim'}(b) = s_{\succsim_{-i}}(b) + s_{\succsim'}(b)$, with which we can deduce that $s_{\succsim}(b) \geq s_{\succsim'}(b)$. The case with the inverse inequality can be proven symmetrically, and we conclude that $s_{\succsim}(b) = s_{\succsim'}(b)$. \square

Corollary 4. *The only voting rule for top-truncated preferences (on either \mathcal{D}_1 or \mathcal{D}_2) in the Borda class that satisfies dom-power is the optimistic Borda rule.*

After realising that both the axioms of bot-indifference and of dom-power take the form of monotonicity-like properties, we easily see that they are independent of all other axioms in the characterisation of the Borda class.

We continue with the averaged Borda rule, which we are going to link to the property of *full-cancellation*. This axiom, in the spirit of top-cancellation, prescribes the equal status of all alternatives as far as the outcome of the aggregation process is concerned, and applies in cases where for all pairs of alternatives a, b the same number of agents prefers a to b and b to a .

- *Full-cancellation:* Consider any profile \succsim . If for all $x, y \in A$ it is the case that

$$\#\{i \in N \mid x \succ_i y\} = \#\{i \in N \mid y \succ_i x\},$$

then $F(\succsim) = A$.

Note that full-cancellation reduces to the standard cancellation axiom for the special case of profiles of linear orders, and is in general logically independent of top-cancellation. Interestingly, when combined with other axioms that appear in the characterisation of the Borda class, full-cancellation becomes very strong:

Lemma 2. *Neutrality, reinforcement, monotonicity, and full-cancellation together imply anonymity, continuity, top-cancellation, and top-CL-consistency.*

Proof. First, it is easy to see that monotonicity (together with neutrality) implies that in every single-agent profile, that agent's top alternative is going to be the unique winner, which is the property of *faithfulness*. Then, in a similar manner to the one that Hansson and Sahlquist (1976) have used to prove that neutrality, reinforcement, faithfulness and (full-)cancellation characterise the Borda rule on domains of linear orders, we can show that these axioms directly characterise the rule represented by the symmetric Borda scores on top-truncated preferences, and therefore imply a rule that is in the Borda class, satisfying anonymity, continuity, top-cancellation, and top-CL-consistency. \square

Using Lemma 2, we obtain a proof for the characterisation of the averaged Borda rule on top-truncated preferences (in Theorem 6) that explicitly hinges on the effect of full-cancellation within the Borda class. Although Theorem 6 could also be proven without any reference to the Borda class,¹¹ our proof sheds light on the particular way in which the averaged Borda rule differs from the other two rules in the Borda class (the optimistic and pessimistic Borda rules), by taking advantage of structurally analogous proof techniques.

Theorem 6. *The only voting rule for top-truncated preferences (on either \mathcal{D}_1 or \mathcal{D}_2) that satisfies neutrality, reinforcement, monotonicity, and full-cancellation is the averaged Borda rule.*




Figure 7: The construction of the profiles appearing in the proof of Theorem 6, for $m = 5$ and $k = 2$.

Proof. After recalling that the averaged Borda rule on top-truncated preferences corresponds to the symmetric way of defining domination scores, it is not hard to see that this rule satisfies full-cancellation (and the other axioms of the statement).

For the other direction, take a rule that satisfies all the required axioms. By Lemma 2, we know that this rule also satisfies anonymity, continuity, top-cancellation, and top-CL-consistency, and hence is in the Borda class. Take such a rule F_S . Consider an arbitrary top-truncated preference \succsim with $\#TOP(\succsim) = k$, with $1 \leq k \leq m - 2$

¹¹The fact that the characterisation of the rule represented by the symmetric Borda scores (i.e., the averaged Borda rule on top-truncated preferences) works for other domains of preorders, as well as for linear orders, is also mentioned in informal comments by Young (1974) and by Hansson and Sahlquist (1976).

(otherwise the proof is trivial). Let \succsim be the profile that consists of k copies of that preference \succsim , and let \succsim' be the profile that consists of k copies of the preference that reverses the order of the alternatives in $TOP(\succsim)$ and keeps the bottom alternatives in \succsim unaltered. Then, we construct two profiles \mathbf{L} and \mathbf{L}' of linear orders such that $TOP(\mathbf{L}) = TOP(\mathbf{L}') = BOT(\succsim)$. Profile \mathbf{L} consists of k preferences, all having the alternatives on top ranked in the same order (any arbitrary one); Profile \mathbf{L}' also consists of k preferences, with the alternatives on top ranked in the same order, reversed from the one of \mathbf{L} . Moreover, in both \mathbf{L} and \mathbf{L}' , every alternative on the bottom takes up each of the k positions exactly once. Figure 7 provides an example of this construction, for $m = 5$ and $k = 2$, in the top-truncated preference domain \mathcal{D}_1 .

Consider the following profile \succsim'' , where full-cancellation applies:

$$\succsim'' = (\succsim, \succsim', \mathbf{L}, \mathbf{L}').$$

Suppose that $TOP(\succsim) = \{a_1, \dots, a_k\} \subseteq A$ and $BOT(\succsim) = \{a_{k+1}, \dots, a_m\} \subseteq A$. By construction, we have that

$$s_{\succsim''}(a_1) = \dots = s_{\succsim''}(a_k) = (m - 1 + m - k)k + \frac{k(k - 1)}{2}2,$$

and

$$s_{\succsim''}(a_{k+1}) = \dots = s_{\succsim''}(a_m) = 2ks_{k+1} + (m - 1 + k)k.$$

But by full-cancellation we must have that $F_s(\succsim'') = A$, so all alternatives must have the same score. This means that

$$(m - 1 + m - k) + (k - 1) = 2s_k + (m - 1 + k) \Leftrightarrow s_{k+1} = \frac{m - k - 1}{2}.$$

We conclude that F_s is the averaged Borda rule. \square

Theorem 6 and Lemma 2 imply the following corollary:

Corollary 5. *The only voting rule for top-truncated preferences (on either \mathcal{D}_1 or \mathcal{D}_2) in the Borda class that satisfies full-cancellation is the averaged Borda rule.*

5. Conclusion

We have provided a complete axiomatic analysis of various generalisations of the famous Borda rule previously been discussed in the literature on top-truncated preferences. The three most prominent such rules are the optimistic, the pessimistic, and the averaged Borda rule, which differ on the scores they assign to the bottom alternatives in a top-truncated order. Our axioms clarify the different contexts for which these rules are relevant, and the requirements they implicitly impose on the treatment of the alternatives not ranked by an agent. Figure 8 provides a graphical summary of our characterisation results.

Of course, this paper has not closed all gaps in our understanding regarding suitable generalisations of the classical Borda rule to richer domains of preferences. But by




Figure 8: Borda rules for top-truncated preferences and the axioms that characterise them. The axioms next to an arrow indicate those properties that are required in order to obtain a specific class of rules, or a specific rule, within a larger class.

focusing on top-truncated preferences, we have not only derived a better comprehension of a domain of immediate practical significance but we have also obtained valuable intuitions that could potentially apply to more general domains of preorders as well. Moreover, our work may open up the way for similar investigations with respect to other popular voting rules on domains that go beyond the classical one of linear orders. For instance, do appropriate generalisations of the Kemeny rule come with corresponding desirable axioms that can naturally differentiate between them? Questions of this nature would be intriguing to investigate in future work.

Acknowledgments

We are grateful to Ashley Piggins and Bill Zwicker for the insightful comments received on preliminary versions of this work. We also thank Haris Aziz, Franz Dietrich, Hans Peters, and Marcus Pivato for useful discussions, and two anonymous JME reviewers for their feedback.

References

- Ackerman, M., Choi, S.Y., Coughlin, P., Gottlieb, E., Wood, J., 2013. Elections with partially ordered preferences. *Public Choice* 157, 145–168.
- Baumeister, D., Faliszewski, P., Lang, J., Rothe, J., 2012. Campaigns for lazy voters: Truncated ballots, in: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

- de Borda, J.C., 1784. Memoire sur les elections au scrutin. Histoire de l'Academie Royale des Sciences .
- Brandl, F., Peters, D., 2019. An axiomatic characterization of the Borda mean rule. *Social Choice and Welfare* 52, 685–707.
- Caragiannis, I., Krimpas, G., Voudouris, A., 2015. Aggregating partial rankings with applications to peer grading in massive online open courses, in: Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS).
- Cullinan, J., Hsiao, S.K., Polett, D., 2014. A Borda count for partially ordered ballots. *Social Choice and Welfare* 42, 913–926.
- Dietrich, F., 2014. Scoring rules for judgment aggregation. *Social Choice and Welfare* 42, 873–911.
- Duddy, C., Piggins, A., Zwicker, W.S., 2016. Aggregation of binary evaluations: A Borda-like approach. *Social Choice and Welfare* 46, 301–333.
- Dummett, M., 1997. Principles of Electoral Reform. Oxford University Press.
- Emerson, P., 2013. The original Borda count and partial voting. *Social Choice and Welfare* 40, 353–358.
- Farkas, D., Nitzan, S., 1979. The Borda rule and Pareto stability: A comment. *Econometrica* 47, 1305–1306.
- Fishburn, P.C., Gehrlein, W.V., 1976. Borda's rule, positional voting, and Condorcet's simple majority principle. *Public Choice* 28, 79–88.
- Hansson, B., Sahlquist, H., 1976. A proof technique for social choice with variable electorate. *Journal of Economic Theory* 13, 1027–1041.
- Kruger, J., Terzopoulou, Z., 2020. Strategic manipulation with incomplete preferences: Possibilities and impossibilities for positional scoring rules, in: Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS).
- Myerson, R.B., 1995. Axiomatic derivation of scoring rules without the ordering assumption. *Social Choice and Welfare* 12, 59–74.
- Nitzan, S., Rubinstein, A., 1981. A further characterization of Borda ranking method. *Public Choice* 36, 153–158.
- Saari, D.G., 1990. The Borda dictionary. *Social Choice and Welfare* 7, 279–317.
- Saari, D.G., 2008. *Disposing Dictators, Demystifying Voting Paradoxes*. Cambridge University Press.
- Smith, J.H., 1973. Aggregation of preferences with variable electorate. *Econometrica* 41, 1027–1041.

- Terzopoulou, Z., Endriss, U., 2019. Aggregating incomplete pairwise preferences by weight, in: Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI).
- 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.
- Zwicker, W.S., 2016. Introduction to the theory of voting, in: Brandt, F., Conitzer, V., Endriss, U., Lang, J., Procaccia, A.D. (Eds.), *Handbook of Computational Social Choice*. Cambridge University Press.