
On Distance Rationalizability of Some Voting Rules

Edith Elkind

School of ECS
University of Southampton, UK
and
Division of Mathematical Sciences
Nanyang Technological University
Singapore

Piotr Faliszewski

Department of Computer Science
AGH University of Science
and Technology, Kraków, Poland

Arkadii Slinko

Department of Mathematics
University of Auckland
New Zealand

Abstract

The concept of *distance rationalizability* has several applications within social choice. In the context of voting, it allows one to define (“rationalize”) voting rules via a consensus class (roughly, a set of elections in which it is obvious who should win) and a distance function: namely, a candidate is said to be an election winner if it is ranked first in one of the nearest (with respect to the given distance) consensus elections. It is known that many classic voting rules can be represented in this manner. In this paper, we provide new results on distance rationalizability of several well-known voting rules such as all scoring rules, Approval, Young’s rule and Maximin. We also show that a previously published proof of distance rationalizability of Young’s rule is incorrect: the consensus notion and the distance function used in that proof give rise to a voting rule that is similar to—but distinct from—the Young’s rule. Finally, we demonstrate that some voting rules cannot be rationalized via certain notions of consensus. To the best of our knowledge, these are the first non-distance-rationalizability results for voting rules.

1 Introduction

Preference aggregation is an important task both for human societies and for multi-agent systems. Indeed, it is often the case that a group of agents has to make a joint decision, i.e., to select a unique alternative from a space of available solutions, even though the agents may have different opinions about the relative merits of these solutions. A standard method of preference aggregation is voting: the agents submit ballots, which are usually rankings (total orders) of the alternatives

(candidates), and a *voting rule* is used to select the best alternative. While in such settings the goal is usually to select the alternative that best matches the collective preferences, there is no universal agreement on how to reach this goal. As a consequence, there is a multitude of voting rules, each of which purports to aggregate the voters’ preferences as faithfully as possible. These voting rules are remarkably diverse. For example, in Plurality voting, each voter assigns a point to her most preferred alternative, and the alternative with the largest number of points wins. In contrast, in Copeland voting, for each pair of candidates (c_1, c_2) we determine if the majority of voters prefers c_1 to c_2 , and the winner is the candidate that wins the largest number of such pairwise contests. While both of these rules select the alternative that is, in some sense, the most preferred by the voters, it is well-known that their outputs can be very different. Moreover, there are many other voting rules, such as Borda, STV, Maximin, Bucklin, Dodgson, Ranked Pairs, etc., which behave differently from both Plurality and Copeland, as well as from each other; all of these rules are well-studied and used in practice (see, e.g., [Brams and Fishburn, 2002]).

Why cannot we settle on a single voting rule, which will aggregate the preferences optimally? One answer to this question is provided by the famous Arrow’s impossibility theorem [Arrow, 1951 revised edition 1963]: there is no voting rule (or, more precisely, no social welfare function) that simultaneously satisfies several natural desiderata, so in each real-life scenario we have to decide which of these desiderata we are willing to sacrifice. Another perspective on this issue is offered by viewing each collection of preferences as an imperfect approximation to some kind of consensus. Under this view, the winner for a given collection of preferences, or a *preference profile*, is the most preferred candidate in the “closest” consensus preference profile. The differences among voting rules can then be explained by the fact that there are several ways of defining consensus, as well as closeness.

In more detail, there are certain situations where it appears obvious which candidate is best liked by the voters. For example, if all voters have identical preferences (this situation is usually referred to as *strong unanimity*), then the candidate that is ranked first by all voters is obviously the best choice. A more realistic scenario, known as *unanimity*, is where all voters rank the same candidate first; again, it is natural to assume that this candidate should be elected. Somewhat more controversially, if there is a candidate that beats any other candidate in a pairwise election (such a candidate is known as the *Condorcet winner*, and the situation is known as *Condorcet consensus*), it can be assumed that this candidate is the best choice. Now, for many preference profiles, even the least restrictive of these conditions does not hold, i.e., there is no obvious consensus. One can deal with this issue by trying to tweak the voters’ opinions or the structure of the election. Then a plausible outcome of the election is a consensus profile (or, more precisely, the most preferred alternative in this profile) that can be obtained from the given profile by making as few changes as possible. In many situations, this approach can be formalized using the notion of *distance*: given a preference profile, we identify the closest consensus preference profile and let the winner be the top alternative in this profile. The voting rules that can be described in such terms are called *distance-rationalizable*: they can be “rationalized” by describing an appropriate notion of consensus and a distance function. For instance, we can rationalize the Plurality rule via unanimity consensus and the Hamming distance (where preference profiles are interpreted as $|V|$ -dimensional vectors); for details, see Section 2.

This approach is not limited to Plurality. In fact, it turns out that many well-known voting rules are distance-rationalizable. To date, the most complete list of such rules is provided by Meskanen and Nurmi [2008] (but see also [Baigent, 1987; Klamler, 2005b; 2005a]). There, the authors show how to distance-rationalize many voting rules, including, among others, Plurality, Borda, Veto, Copeland, Dodgson, Kemeny, Slater, and STV¹. However, the treatment provided by Meskanen and Nurmi [2008] is not exhaustive: there are well-known voting rules (and, in fact, entire families of rules) whose distance rationalizability was left open.

In this paper, we fill in some of the gaps in our understanding of distance rationalizability. In Section 3, we focus on scoring rules—a natural class of voting

rules that includes such important rules as Plurality, Borda, Veto, and k -approval for any fixed value of k . Each scoring rule can be described by a vector $\alpha = (\alpha_1, \dots, \alpha_m)$. The interpretation is that each voter grants α_1 points to his top candidate, α_2 points to his second most favorite candidate, etc., and the candidate with the largest number of points wins. We start by showing that all such rules can be rationalized with respect to the unanimity consensus via a *pseudodistance*—a function that satisfies all distance axioms except one that requires that the distance between two non-identical objects must be greater than 0. Moreover, if all entries of α are different, this pseudodistance is, in fact, a distance, so all scoring rules that correspond to such values of α are distance-rationalizable. We then show that for a large class of scoring rules that includes Veto and k -approval this result cannot be improved: all rules in this class are not distance-rationalizable with respect to the unanimity consensus. We conclude Section 3 by showing that, when proving distance-rationalizability, it is important to choose an appropriate notion of consensus: we demonstrate that no scoring rule can be rationalized with respect to Condorcet consensus. Our proofs that some voting rules cannot be distance-rationalized with respect to a particular notion of consensus are quite simple. However, we find them interesting as, to the best of our knowledge, these are the very first negative results in the general area of distance rationalizability of voting rules.

In Section 4, we consider elections in which voters can approve of one or more candidates, and each voter chooses how many candidates she approves of. In such settings, the winner is usually determined using the Approval rule: the winners are the candidates with the largest number of approvals. As such elections cannot be represented in the traditional model, in which the voters’ preferences correspond to total orders, the three notions of consensus outlined above are not appropriate for this setting. To remedy this, we propose two alternative notions of consensus for this model, and show that Approval can be rationalized with respect to both of them.

We then move on to another notion of consensus, namely, Condorcet consensus, and consider three related operations that can be used to define a distance between preference profiles, namely, adding, deleting, or replacing voters. We show that all three of these transformations can be formalized as distances. We then study the voting rules obtained by combining these three distances with Condorcet consensus. It turns out that deleting voters corresponds to the well-known Young’s rule, while adding voters corresponds to the equally well-known Maximin rule. However, for

¹Meskanen and Nurmi [2008] also claim that one can distance-rationalize the well-known Young’s rule. However, as we show in Section 5, the construction provided in [Meskanen and Nurmi, 2008] is incorrect.

replacing voters the situation is somewhat more complicated. Indeed, Meskanen and Nurmi [2008] claim that the corresponding voting rule is equivalent to Young’s rule. However, we show that this is not the case. Specifically, we construct a preference profile on which Young’s rule and the voter replacement-based rule produce different results. Since in [Meskanen and Nurmi, 2008] this equivalence claim was used to prove that Young’s rule is distance-rationalizable, it follows that our Proposition 7 is the first proof that Young’s rule is distance-rationalizable. We then proceed to study the voter replacement-based rule in more detail. While this rule has a very natural definition via distance-rationalizability, it appears to be distinct from all known voting rules. We conclude Section 5 by showing that the winner determination problem under this rule is computationally hard. We conclude the paper by summarizing our results and presenting several directions for future work.

2 Preliminaries

An *election* E is a pair (C, V) , where $C = \{c_1, \dots, c_m\}$ is a set of *candidates* and $V = (v_1, \dots, v_n)$ is a collection of *voters*. Each voter v_i is usually represented via her *preference order* \succ_i , which is a total order over the candidates in C (see, however, Section 4). For example, given a candidate set $C = \{c_1, c_2, c_3\}$, a voter v_i that likes c_2 best, then c_1 , and then c_3 is represented as $c_2 \succ_i c_1 \succ_i c_3$. To simplify notation, we will sometimes denote the position of a candidate c_j in the preference order of a voter v by v^j ; thus, in the example above we have $v_i^1 = 2$, $v_i^2 = 1$, $v_i^3 = 3$.

A *voting rule* \mathcal{R} is a function that given an election $E = (C, V)$ outputs a set $\mathcal{R}(E) \subseteq C$ of *winners* of the election. Note that we do not require $|\mathcal{R}(E)| = 1$. Indeed, there are cases where, e.g., due to symmetry, it is impossible to declare a single winner, in which case we may have $\mathcal{R}(E) = \emptyset$ or $|\mathcal{R}(E)| > 1$. In practice, one may then need to use a *draw resolution rule*, which can be either deterministic (e.g., lexicographic) or randomized (e.g., a fair coin toss); however, in the rest of this paper we will ignore this issue. Perhaps the best known voting rule is the Plurality rule $\mathcal{R}_{\text{plur}}$ defined in the introduction; note that there can be two or more candidates that have the largest number of first-place votes, so we may have $|\mathcal{R}_{\text{plur}}(E)| > 1$. Another important rule is that of Condorcet [1785]: Given an election $E = (C, V)$, a candidate c_i is a *Condorcet winner* (we write $c_i \in \mathcal{R}_{\text{Cond}}(E)$) if for each $c_j \in C$, $c_i \neq c_j$, a strict majority of voters prefers c_i to c_j . It is easy to see that for any election E we have $|\mathcal{R}_{\text{Cond}}(E)| \leq 1$, and it can be the case that $\mathcal{R}_{\text{Cond}}(E) = \emptyset$. We will define several other voting rules later in the paper.

Intuitively, a preference profile corresponds to a *consensus* among the voters when there exists an alternative that is clearly better from the collective point of view than any other alternative. There are several ways of formalizing this intuition, which correspond to different classes of *consensus elections*. Specifically, fix an election $E = (C, V)$. We say that an election $E = (C, V)$ is *strongly unanimous* if $\succ_i = \succ_j$ for all $v_i, v_j \in V$. We denote the set of all strongly unanimous elections by \mathcal{S} .² Further, we say that $E = (C, V)$ is *unanimous* if there exists a candidate $c_i \in C$ such that all voters in V rank c_i first. The set of all unanimous elections is denoted by \mathcal{U} . The elections in \mathcal{S} and \mathcal{U} are clearly consensus elections: in both cases, the top alternative of all voters should be elected. Another class of consensus elections is given by elections that have a Condorcet winner; we denote the set of all such elections by \mathcal{C} . Clearly, each strongly unanimous election is also unanimous, and each unanimous election has a Condorcet winner, but the converse is not true. Our list of consensus classes is not meant to be exhaustive. Indeed, depending on the setting, one may want to define other types of consensus elections. However, any such definition should have the property that any “consensus” election has an obvious winner.

Given a set X , we say that a function $d: X \times X \rightarrow \mathbb{R} \cup \{+\infty\}$ is a *distance* (or *metric*) over X if for each $x, y \in X$ it satisfies the following four axioms:

- (1) $d(x, y) \geq 0$ (non-negativity),
- (2) $d(x, y) = 0$ if and only if $x = y$ (identity of indiscernibles),
- (3) $d(x, y) = d(y, x)$ (symmetry), and
- (4) for each $z \in X$, $d(x, y) \leq d(x, z) + d(z, y)$ (triangle inequality).

A function that satisfies axioms (1), (3), and (4) is called a *pseudodistance*, and a function that satisfies axioms (1), (2), and (4) is called a *quasidistance*. In what follows, the elements of the set X will usually be either voters (total orders over candidates) or elections. Note that any distance $d(v, w)$ over voters with preferences over a candidate set C can be extended to a distance $\hat{d}(E, E')$ over elections $E = (C, (v_1, \dots, v_n))$, $E' = (C, (v'_1, \dots, v'_n))$ by setting $\hat{d}(E, E') = \sum_{i=1}^n d(v_i, v'_i)$. Clearly, \hat{d} satisfies all distance axioms as long as d does. We now provide two examples of distances defined over pairs of voters with preferences over a set of candidates C . Our first example is the *discrete distance* $d_{\text{discr}}(v, w)$, given

²In this paper we do not treat strongly unanimous elections directly, but we mention this notion due to its naturalness and the fact that it is very useful in distance-rationalizing some natural voting rules [Meskanen and Nurmi, 2008].

by $d_{\text{discr}}(v, w) = 1$ if $v \neq w$ and $d_{\text{discr}}(v, w) = 0$ otherwise. Clearly, the corresponding distance over elections $\widehat{d_{\text{discr}}}(E, E')$ is equivalent to the *Hamming distance* $d_H(E, E')$, which is defined as $d_H(E, E') = |\{i \mid v_i \neq v'_i\}|$. Our second example is the *Dodgson distance*, or *swap distance*, $d_{\text{swap}}(v, w)$, defined as $d_{\text{swap}}(v, w) = |\{(c, c') \in C^2 \mid c \succ_v c', c' \succ_w c\}|$. It is not hard to check that both Dodgson distance and the discrete distance (and hence the Hamming distance) satisfy the distance axioms listed above. We are now ready to define distance rationalizability.

Definition 1. Let d be a distance over elections, let \mathcal{E} be a set of elections, and let \mathcal{W} be a voting rule that for each election $E \in \mathcal{E}$ satisfies $\mathcal{W}(E) \neq \emptyset$. We define the $(\mathcal{E}, \mathcal{W}, d)$ -score of a candidate c_i in an election E to be the distance (according to d) between E and a closest election $E' \in \mathcal{E}$ such that $c_i \in \mathcal{W}(E')$. The set of $(\mathcal{E}, \mathcal{W}, d)$ -winners of an election $E = (C, V)$ consists of those candidates in C whose $(\mathcal{E}, \mathcal{W}, d)$ -score in E is the smallest.

Definition 2. A voting rule \mathcal{R} is distance-rationalizable via a set of elections \mathcal{E} , a voting rule \mathcal{W} , and a distance d , or $(\mathcal{E}, \mathcal{W}, d)$ -rationalizable, if for each election E , a candidate $c \in C$ is an \mathcal{R} -winner of E if and only if he is a $(\mathcal{E}, \mathcal{W}, d)$ -winner of E .

Throughout this paper, we only use Definitions 1 and 2 in settings where \mathcal{E} is a consensus class (i.e., $\mathcal{E} \in \{\mathcal{S}, \mathcal{U}, \mathcal{C}\}$, or \mathcal{E} is one of the two classes defined in Section 4). In such cases, we assume that the rule \mathcal{W} outputs the consensus winner(s), and omit \mathcal{W} from the notation. Further, when a voting rule \mathcal{R} is distance-rationalizable via one of the consensus classes defined above, i.e., \mathcal{R} is (\mathcal{E}, d) -rationalizable for some distance d and $\mathcal{E} \in \{\mathcal{S}, \mathcal{U}, \mathcal{C}\}$, we will simply say that \mathcal{R} is distance-rationalizable. One can define pseudodistance-rationalizable rules and quasidistance-rationalizable rules in the same manner; we omit the details to avoid repetition.

One of the most natural examples of a distance-rationalizable rule is Dodgson’s rule [Dodgson, 1876], defined as follows. Fix an election $E = (C, V)$. The *Dodgson score* of a candidate $c_i \in C$ is the minimum number of swaps of adjacent candidates in the preference lists of the voters in V after which c_i becomes the Condorcet winner. A candidate c_i is a *Dodgson winner* if she has the minimum Dodgson score (naturally, there may be multiple Dodgson winners). Thus, in the nomenclature of Definition 2, Dodgson’s rule is $(\mathcal{C}, \widehat{d_{\text{swap}}})$ -rationalizable. Similarly, it is not hard to see that the Plurality rule is (\mathcal{U}, d_H) -rationalizable.

3 Scoring Rules

In this section, we will show that all scoring rules—an important class of voting rules that includes such famous rules as Plurality, Borda, Veto, and k -approval—are very close to being distance-rationalizable. More precisely, we prove that all such rules are *pseudodistance-rationalizable*, i.e., can be defined via a consensus class and a pseudodistance.

We start by formally defining scoring rules. For any m -element vector $\alpha = (\alpha_1, \dots, \alpha_m)$ of nonnegative integers, where $\alpha_1 \geq \alpha_2 \geq \dots \geq \alpha_m$, a *scoring rule* \mathcal{R}_α is defined as follows. Fix an election $E = (C, V)$, where $C = \{c_1, \dots, c_m\}$ and $V = (v_1, \dots, v_n)$. The α -score of a candidate $c_j \in C$ is given by $\sum_{i=1}^n \alpha_{v_i^j}$. That is, candidate c_j receives α_k points from each voter that puts her in the k th position. The winners of E under \mathcal{R}_α are the candidates with the maximum α -score.

Perhaps the best-known voting procedure that is traditionally defined via a scoring rule is the *Borda rule*, which corresponds to the scoring vector $(m-1, m-2, \dots, 0)$. Under this rule, the number of points a candidate c receives from a voter v is equal to the number of candidates that v ranks below c . The Plurality rule can also be represented as a scoring rule: the corresponding scoring vector is given by $(1, 0, \dots, 0)$. Another prominent rule in this class is *Veto*, which can be described by the scoring vector $(1, \dots, 1, 0)$. For k -approval, the corresponding scoring vector is given by $\alpha_1 = \dots = \alpha_k = 1$, $\alpha_{k+1} = \dots = \alpha_m = 0$.

Any vector $\alpha = (\alpha_1, \dots, \alpha_m)$ can be used to define a pseudodistance d_α over voters with preferences over candidates in C , $|C| = m$, as follows. Fix two voters v and w . Set $d_\alpha(v, w) = |\alpha_{v^1} - \alpha_{w^1}| + |\alpha_{v^2} - \alpha_{w^2}| + \dots + |\alpha_{v^m} - \alpha_{w^m}|$. It is not hard to see that d_α satisfies all pseudodistance axioms. Moreover, if $\alpha_i \neq \alpha_{i+1}$ for all $i = 1, \dots, m-1$ (we will call such vectors, and the corresponding distances and scoring rules *non-degenerate*), then d_α also satisfies axiom (2), i.e., it is a distance. On the other hand, if $\alpha_i = \alpha_{i+1}$ for some $i = 1, \dots, m$, then we have $d_\alpha(v, w) = 0$ for the two voters v, w such that w is obtained from v by swapping the i th and the $(i+1)$ st candidate in v ’s preference ordering, i.e., d_α is a pseudodistance, but not a distance. Note that for $\alpha = (m-1, \dots, 1, 0)$ (the vector that corresponds to the Borda rule) we have $d_\alpha(v, w) = 2d_{\text{swap}}(v, w)$. More generally, $d_\alpha(v, w)$ can be interpreted as the cost of transforming v into w by a sequence of swaps of adjacent candidates, where the cost of swapping the candidate in the k th position with the one just below him is given by $2(\alpha_k - \alpha_{k+1})$.

We will now prove that any scoring rule \mathcal{R}_α is $(\mathcal{U}, \widehat{d_\alpha})$ -rationalizable, where $\widehat{d_\alpha}$ is the pseudodistance over

elections that corresponds to \widehat{d}_α . Note that this implies that all non-degenerate scoring rules are distance rationalizable. While distance rationalizability of the Borda rule was proven by Meskanen and Nurmi [2008], no such result was previously known for other scoring rules.

Theorem 3. *For each scoring vector $\alpha = (\alpha_1, \dots, \alpha_m)$ and each election $E = (C, V)$ with $|C| = m$, a candidate $c_j \in C$ is a winner of E according to \mathcal{R}_α if and only if c_j is a $(\mathcal{U}, \widehat{d}_\alpha)$ -winner of E .*

Proof. Fix a candidate $c_j \in C$, and consider a voter $v \in V$ that ranks c_j in the k th position, i.e., $v^j = k$. Consider an arbitrary preference order w in which c_j is ranked first. We have $\sum_{\ell=1}^m \alpha_{v^\ell} = \sum_{\ell=1}^m \alpha_{w^\ell} = m(m+1)/2$, so $\sum_{\ell=1}^m (\alpha_{w^\ell} - \alpha_{v^\ell}) = 0$. On the other hand, we have $\alpha_{w^j} - \alpha_{v^j} = \alpha_1 - \alpha_k$. Clearly, for any m real numbers a_1, \dots, a_m such that $\sum_{i=1}^m a_i = 0$ and $a_1 = x \geq 0$, we have $\sum_{i=1}^m |a_i| \geq 2x$. Thus, for $d_\alpha(v, w) = \sum_{\ell=1}^m |\alpha_{w^\ell} - \alpha_{v^\ell}|$ we have $d_\alpha(v, w) \geq 2(\alpha_1 - \alpha_k)$. On the other hand, for the preference order w' that is obtained from v by swapping c_j with the top candidate in v , we have $d_\alpha(v, w') = 2(\alpha_1 - \alpha_k)$. Hence, the d_α -distance from v to the nearest vote that ranks c_j first is exactly $2(\alpha_1 - \alpha_k)$, and, consequently, the \widehat{d}_α -distance from E to the nearest unanimous election in which all voters rank c_j first is exactly $\sum_{i=1}^n 2(\alpha_1 - \alpha_{v_i^j})$. On the other hand, the score that c_j receives in \mathcal{R}_α is equal to $\sum_{i=1}^n \alpha_{v_i^j}$. Thus, any candidate in E with the highest score under \mathcal{R}_α is a $(\mathcal{U}, \widehat{d}_\alpha)$ -winner of E and vice versa. \square

For some scoring rules the statement of Theorem 3 cannot be strengthened to distance rationalizability.

Proposition 4. *Any scoring rule defined by a vector α with $\alpha_1 = \alpha_2$ is not distance-rationalizable with respect to consensus class \mathcal{U} .*

Proof. Fix a scoring rule \mathcal{R}_α with $\alpha_1 = \alpha_2$. Consider an election E in which all voters rank some candidate c_1 first and another candidate c_2 second. Under \mathcal{R}_α , both c_1 and c_2 are winners. On the other hand, E is clearly unanimous, with c_1 being the consensus winner. Thus, for any distance d , the distance between E and the closest election in \mathcal{U} is 0. Therefore, for c_2 to be a (\mathcal{U}, d) -winner, there must exist a unanimous election E' in which c_2 is ranked first such that $d(E, E') = 0$. However, as E' is necessarily different from E , this is impossible for any distance d . \square

Observe that the condition of Proposition 4 is satisfied by Veto and k -approval for $k > 1$, so these rules are not distance-rationalizable with respect to \mathcal{U} .

What about scoring rules with $\alpha_1 \neq \alpha_2$, but $\alpha_j = \alpha_{j+1}$ for some $j = 2, \dots, m-1$? As argued above, for such rules d_α is a pseudodistance, but not a distance. It is tempting to conjecture that we can extend the proof of Proposition 4 to this case in order to show that no such rule is distance-rationalizable with respect to \mathcal{U} . However, this conjecture is easy to refute: indeed, the Plurality rule for m candidates is a scoring rule with $\alpha_2 = \dots = \alpha_m = 0$, but we have seen that Plurality can be distance-rationalized with respect to \mathcal{U} . Observe that the distance d_{discr} that we have used for this purpose is different from $d_{(1,0,\dots,0)}$: for example, for $v = a \succ b \succ c$ and $w = a \succ c \succ b$ we have $d_{\text{discr}}(v, w) = 1$, $d_{(1,0,\dots,0)}(v, w) = 0$. However, it is not clear how to generalize this construction to other scoring rules with $\alpha_1 > \dots > \alpha_j = \alpha_{j+1}$. Thus, distance rationalizability of such rules with respect to \mathcal{U} remains an interesting direction for future work.

We have shown that any non-degenerate scoring rule can be rationalized with respect to the consensus class \mathcal{U} . We will now argue that there is *no* scoring rule that can be rationalized with respect to Condorcet consensus; the proof is similar to that of Proposition 4.

Proposition 5. *No scoring rule is distance-rationalizable with respect to the consensus class \mathcal{C} .*

Proof. It is known that no scoring rule is Condorcet-consistent [Moulin, 1991]. That is, for any scoring rule \mathcal{R} there exists an election $E = (C, V)$ such that $\mathcal{R}(E) = \{c\}$, E has a Condorcet winner c' , but $c \neq c'$. Now, consider any distance d . Since d satisfies axiom (2) (identity of indiscernibles), the election E has exactly one candidate with the (\mathcal{C}, d) -score of 0. Hence, the set of (\mathcal{C}, d) -winners of E consists of c' only, and is therefore different from the set of \mathcal{R} -winners of E . \square

4 Approval Voting

Under Approval voting, each voter may approve of (give a point to) one or more candidates; the winners are the candidates with the largest number of points. In this setting, a voter's actual vote is not completely determined by her preference order, as we also need to know how many of her top candidates she approves of. We can ask the voter to provide this number in addition to the ordering, or, instead, require her to simply list the approved candidates. In this section, we will work with the latter model. Thus, we identify each voter v with a subset of the candidate set C .

As should be clear from the previous paragraph, to rationalize Approval voting, we need a notion of consensus and a metric that are defined on sets of candidates rather than preference orderings. A natural class of consensus elections in this setting corresponds

to situations where there is a single candidate that all voters approve of. Formally, we let \mathcal{A} be the set of all elections (C, V) such that $|\cap_{v \in V} v| = 1$. Alternatively, we can allow the voters to agree on several candidates; the corresponding consensus class $m\mathcal{A}$ consists of all elections that satisfy $|\cap_{v \in V} v| \geq 1$.

Furthermore, we define the distance between two voters as the size of the symmetric difference of the corresponding candidate sets, i.e., we set $d_A(v, w) = |(v \setminus w) \cup (w \setminus v)|$. We extend this definition to a distance \widehat{d}_A on entire elections in the standard manner. We are now ready to present the main result of this section.

Theorem 6. *Approval voting is $(\mathcal{A}, \widehat{d}_A)$ -rationalizable, as well as $(m\mathcal{A}, \widehat{d}_A)$ -rationalizable.*

Proof. We will prove the theorem for \mathcal{A} ; the argument for $m\mathcal{A}$ is similar. Consider an election $E = (C, V)$, $|V| = n$. For each $c_i \in C$, set $n_i = |\{v \in V \mid c_i \in v\}|$. Now, fix a candidate c_i . First, suppose that there is no candidate that is approved by all voters. Then, to obtain an election in \mathcal{A} in which c_i is the unique consensus winner, we need to make $n - n_i$ voters approve of c_i , so c_i 's $(\mathcal{A}, \widehat{d}_A)$ -score is $n - n_i$. Thus, the candidates with the lowest $(\mathcal{A}, \widehat{d}_A)$ -score are exactly the candidates with the highest Approval score. Now, suppose that in E there is exactly one candidate c_j that all voters approve of. In this case, c_j is the unique Approval winner, and $E \in \mathcal{A}$. Clearly, c_j is also the unique candidate with the $(\mathcal{A}, \widehat{d}_A)$ -score of 0, so c_i is a $(\mathcal{A}, \widehat{d}_A)$ -winner if and only if $c_i = c_j$. Finally, suppose that in E there are several candidates that are approved by all voters, i.e., for $W = \cap_{v \in V} v$ we have $|W| > 1$. Now, to turn a candidate $c_j \in W$ into the unique candidate approved by all voters, we can ask some voter to withdraw his approval from all other candidates in W ; thus, the $(\mathcal{A}, \widehat{d}_A)$ -score of c_j is at most $|W| - 1$. On the other hand, to make a candidate $c_k \notin W$ the unique candidate approved by all voters, we need to take away at least one point from each candidate in W as well as to ask some additional voters to approve of c_k , so the $(\mathcal{A}, \widehat{d}_A)$ -score of c_k is at least $|W|$. Thus, in this case, too, the set of Approval winners coincides with the set of candidates with the smallest $(\mathcal{A}, \widehat{d}_A)$ -score. \square

On the other hand, it is not hard to see that Approval voting cannot be rationalized with respect to any consensus class that is defined on preference profiles only, such as unanimity consensus. Indeed, there can be two (unanimous) approval elections E, E' that correspond to the same preference profile, but have different sets of winners, because the number of candidates approved by each voter in E and E' is different.

5 Rationalizability via Condorcet Consensus

In this section, we focus on voting rules that can be rationalized with respect to Condorcet consensus via distances that correspond to adding, deleting, or replacing voters.

To begin, observe that, given an election $E = (C, V)$ with $|V| = n$, we can make any candidate $c \in C$ the Condorcet winner by adding at most $n + 1$ voters that rank c first. Similarly, we can make c the Condorcet winner by replacing at most $\lfloor n/2 \rfloor + 1$ voters in V with voters that rank c first. While not every candidate can be made the Condorcet winner by voter deletion—for example, if a candidate is ranked last by all voters, he will not become the Condorcet winner no matter how many voters we delete—it is still the case that, if at least one voter ranks a given candidate first, this candidate can be made the Condorcet winner by removing at most $n - 1$ voters. Thus, for each candidate c we can define her score with respect to each of these operations as the number of voters that need to be inserted, replaced, or removed, respectively, to make c the Condorcet winner (for deletion, some candidates will have a score of $+\infty$). We will refer to these scores as the *insertion score*, the *replacement score* and the *deletion score*, respectively. Intuitively, for each of these scores, the candidates with a lower score are closer to being the consensus winners than the candidates with a higher score, so each of these scores can be used to define a voting rule.

In fact, there is a well-known voting rule that is defined in these terms, namely, Young's rule, which elects the candidates with the lowest deletion score. Thus, it is natural to ask if the two other scores defined above, i.e., the replacement score and the insertion score, also correspond to well-known voting rules. Another interesting question is whether all three of these scores can be transformed into distances, i.e., whether the corresponding voting rules are distance-rationalizable with respect to Condorcet consensus; observe that this issue is more complicated than might appear at the first sight, since we have to satisfy the symmetry axiom. In the rest of this section, we will provide answers to these two questions.

We will first answer the second question by showing how to transform each of our three scores into a distance. The easiest case is that of the replacement score. Formally, given an election $E = (C, V)$, the *replacement score* $s_r(c)$ of a candidate $c \in C$ is the smallest value of k such that there exists an election $E' = (C, V')$ obtained by changing the preferences of exactly k voters in V in which c is the Condorcet winner; as argued above, $s_r(c) \leq \lfloor n/2 \rfloor + 1$ for all $c \in C$. It

is immediate that the replacement score of any $c \in C$ is exactly the Hamming distance from E to the closest election over the set of candidates C in which c is the Condorcet winner. Thus, the corresponding voting rule is (C, d_H) -rationalizable. In the rest of this section, we will refer to this rule as *the voter replacement rule*. We postpone the discussion of whether this rule is equivalent to any voting rule considered in the literature till the end of the section.

The *insertion score* $s_i(c)$ of a candidate $c \in C$ in an election $E = (C, V)$ is defined as the smallest number $k \geq 0$ such that there exists a list of voters V' , $|V'| = k$, such that c is the Condorcet winner in $E' = (C, V \cup V')$. Similarly, the *deletion score* $s_d(c)$ of a candidate $c \in C$ in an election $E = (C, V)$ is defined as the smallest number $k \geq 0$ such that there exists a list of voters $V' \subseteq V$, $|V'| = k$, such that c is the Condorcet winner in $E' = (C, V \setminus V')$, and $+\infty$ if c cannot be made the Condorcet winner in this manner.

Now, it is easy to see that both the insertion score and the deletion score naturally correspond to quasidistances, i.e., mappings that satisfy non-negativity, identity of indiscernibles and the triangle inequality, but not symmetry. Indeed, given two elections $E = (C, V)$ and $E = (C, V')$ over the same set of candidates C , we can define a function $d'_i(E, E')$ by setting $d'_i(E, E') = k$ if V is a sublist of V' and $|V' \setminus V| = k$, and $d'_i(E, E') = +\infty$ otherwise. Similarly, we can define $d'_d(E, E')$ by setting $d'_d(E, E') = k$ if V' is a sublist of V and $|V \setminus V'| = k$, and $d'_d(E, E') = +\infty$ otherwise. It is not hard to verify that both d'_i and d'_d are quasidistances. Moreover, for each candidate in C his insertion score $s_i(c)$ is equal to the d'_i -distance from E to the nearest (with respect to d'_i) election in \mathcal{C} in which c is the Condorcet winner. Similarly, c 's deletion score $s_d(c)$ is equal to the d'_d -distance from E to the nearest (with respect to d'_d) election in \mathcal{C} in which c is the Condorcet winner. We will now show that we can replace both of these quasidistances with true distances.

For d'_i the solution is simple: we can make d'_i symmetric by allowing ourselves to delete voters as well as to add voters, as, intuitively, deleting a voter is never more useful than adding a voter. Formally, given two elections $E = (C, V)$ and $E = (C, V')$ over the same set of candidates C , we set $d_i(E, E') = |V \setminus V'| + |V' \setminus V|$ (recall that V and V' are lists rather than sets, so by $V \setminus V'$ we mean the list obtained from V by deleting the voters in V'). Clearly, d_i is a distance. Moreover, we will now show that for our purposes it is indistinguishable from d'_i .

Proposition 7. *Consider an election $E = (C, V)$, a candidate $c \in C$, and a $k > 0$. Then there exists an election $E_1 = (C, V_1) \in \mathcal{C}$ such that c is the Condorcet*

winner of E_1 and $d'_i(E, E_1) \leq k$ if and only if there exists an election $E_2 = (C, V_2) \in \mathcal{C}$ such that c is the Condorcet winner of E_2 and $d_i(E, E_2) \leq k$.

Proof. The “only if” direction is immediate: if $d'_i(E, E_1) \leq k$, then $d_i(E, E_1) \leq k$, so we can set $E_2 = E_1$. For the “if” direction, suppose that E_2 has been obtained from E by deleting a sublist of voters $V' \subseteq V$, $|V'| = k_1$, and adding a list of voters V'' , $|V''| = k_2$. Now, consider an election E_3 obtained from E by first adding the voters in V'' and then adding another k_1 voters that rank c first. Clearly, $d'_i(E, E_3) = d_i(E, E_2) \leq k$. We will now show that c is the Condorcet winner in E_3 . Indeed, fix an arbitrary voter $c' \in C$. Suppose that in $(C, V \cup V'')$ there are x voters that prefer c to c' and y voters that prefer c' to c . Then in E_2 there are at most x voters that prefer c to c' and at least $y - k_1$ voters that prefer c' to c . Since c is the Condorcet winner of E_2 , we have $x > y - k_1$. Now, in E_3 there are $x + k_1$ voters that prefer c to c' and y voters that prefer c' to c . As we have argued that $x + k_1 > y$, it follows that the majority of voters in E_3 prefer c to c' . As this is true for any $c' \neq c$, it follows that c is the Condorcet winner in E_3 . Moreover, E_3 has been obtained from E by candidate insertion only, so we can set $E_1 = E_3$. \square

Clearly, we cannot use the same solution for d'_d . Indeed, the argument above demonstrates that adding voters is more useful than deleting voters. Thus, we need to construct a metric that makes it expensive to add voters. As this metric has to be symmetric, a natural approach would be to make the distance between two elections depend on the number of voters in the larger of them, as well as on the difference in the number of voters. For example, we could try to set $d((C, V), (C, V')) = ||V| - |V'|| + (\max\{|V|, |V'|\})^2$. However, it turns out that this approach does not quite work: under this metric, deleting $s_d(c)$ voters may still be more expensive than first deleting some $s' < s_d(c)$ voters and then adding a few voters that rank c first. To overcome this difficulty, we construct a metric that makes it prohibitively difficult to do insertion and deletion at the same time.

Formally, for any pair of elections $E = (C, V)$, $E' = (C, V')$ over the same set of candidates C , we set $k = ||V| - |V'||$, $M = \max\{|V|, |V'|\}$, and let $\bar{d}_d(E, E') = 0$ if $V = V'$, $\bar{d}_d(E, E') = 2 - \frac{1}{k+M^2+1}$ if $V \subset V'$ or $V' \subset V$, and $\bar{d}_d(E, E') = +\infty$ otherwise. The function $\bar{d}_d(E, E')$ is not a metric, as it does not satisfy the triangle inequality. However, we can use it to construct a metric d_d by setting $d_d(E, E') = \min\{\bar{d}_d(E, E_1) + \bar{d}_d(E_1, E_2) + \dots + \bar{d}_d(E_\ell, E') \mid \ell \in \mathbb{N}, E_1, \dots, E_\ell \in \mathcal{E}_C\}$, where \mathcal{E}_C denotes the set of all elections with the set of candidates C . Intuitively, $d_d(E, E')$ is the

shortest path distance in the graph whose vertices are elections in \mathcal{E}_C , and the edge lengths are given by \bar{d}_d . It is well known that for any graph with non-negative edge lengths the shortest path distance satisfies the triangle inequality; it should be clear that d_d satisfies all other axioms of a metric as well. Observe that for any two elections $E, E' \in \mathcal{E}_C$ such that $E = (C, V)$, $E' = (C, V')$ we have $d_d(E, E') < 2$ if $V \subseteq V'$ or $V' \subseteq V$ and $d_d(E, E') > 2$ otherwise.

We will now show that d_d can be used to rationalize Young's rule with respect to Condorcet consensus.

Proposition 8. *Consider an election $E = (C, V)$, $|V| = n$, and two candidates $c_1, c_2 \in C$ such that $s_d(c_1) < +\infty$ or $s_d(c_2) < +\infty$. For $i = 1, 2$, let d_i be the d_d -distance from c_i to the closest (with respect to d_d) election over C in which c_i is a Condorcet winner. Then $s_d(c_1) < s_d(c_2)$ if and only if $d_1 < d_2$.*

Proof. Suppose first that $s_d(c_1) = k_1 < +\infty$, $s_d(c_2) = k_2 < +\infty$. Then one can obtain an election over C in which c_1 (respectively, c_2) is the Condorcet winner by deleting k_1 (respectively k_2) voters from E ; denote this election by E_1 (respectively, E_2). We have $d_d(E, E_1) = 2 - \frac{1}{k_1 + n^2 + 1}$, $d_d(E, E_2) = 2 - \frac{1}{k_2 + n^2 + 1}$. We claim that $d_1 = d_d(E, E_1)$. Indeed, suppose that this is not the case, i.e., $d_d(E, E_1) > d_1$. This means that there exists an election $E' = (C, V')$ such that c_1 is the Condorcet winner of E' and $d_d(E, E') < d_d(E, E_1)$. As E' cannot be obtained from E by deleting voters, it holds that $V' \not\subseteq V$. Now, if also $V \not\subseteq V'$, we immediately obtain $d_d(E, E') > 2$, a contradiction with $d_d(E, E_1) < 2$. Hence, it must be the case that $V \subset V'$, so $|V'| \geq n + 1$, and we have $d_d(E, E') \geq 2 - \frac{1}{2 + (n+1)^2}$. On the other hand, we have $k_1 \leq n - 1$, which implies $d_d(E, E_1) \leq 2 - \frac{1}{n-1+n^2+1}$. As $2 - \frac{1}{2+(n+1)^2} > 2 - \frac{1}{n-1+n^2+1}$, this gives a contradiction as well. Similarly, we can show that $d_2 = d_d(E, E_2)$. Hence, it follows that $k_1 < k_2$ if and only if $d_1 < d_2$.

Now suppose that $s_d(c_1) < +\infty$, $s_d(c_2) = +\infty$ (the case $s_d(c_1) = +\infty$, $s_d(c_2) < +\infty$ is symmetric). Then we have $d_1 \leq 2 - \frac{1}{n+n^2}$, $d_2 \geq 2 - \frac{1}{2+(n+1)^2}$, since we cannot transform E into an election over C in which c_2 is the Condorcet winner by candidate deletion only. Thus, in this case, too, $s_d(c_1) < s_d(c_2)$ if and only if $d_1 < d_2$. \square

As in any election there is at least one candidate c with $s_d(c) < +\infty$, Proposition 8 immediately implies the following result.

Theorem 9. *Young's rule is (C, d_d) -rationalizable.*

We now turn to the first of the two questions posed in the beginning of this section. We have observed

that the voter deletion-based rule is equivalent to Young's rule; the proof follows immediately from the definitions of both rules. We will now show that the voter insertion-based rule is equivalent to another well-known rule, namely, Maximin. Under Maximin, the score of each voter is the outcome of his worst pairwise election. Formally, given an election $E = (C, V)$, for each $c_j \in C$ we set $s_M(c_j) = \min\{\#\{i : c_j \succ_i c_k\} \mid c_k \in C, c_k \neq c_j\}$. The winners are then the candidates c with the highest Maximin score $s_M(c)$.

Proposition 10. *For any election $E = (C, V)$, $|V| = n$, and any candidate $c \in C$ we have $s_i(c) = n - 2s_M(c) + 1$, where $s_i(c)$ is the insertion score of c and $s_M(c)$ is the Maximin score of c .*

Proof. Fix an election $E = (C, V)$, $|V| = n$, and a candidate $c_j \in C$. Set $t = s_M(c_j)$. Let c_k be one of c_j 's worst pairwise opponents, i.e., $|\{q : c_j \succ_q c_k\}| = t$. Now, if we add $n - 2t + 1$ voters that rank c_j first, for any $c_\ell \neq c_j$ there are at most $n - t$ voters that rank c_ℓ above c_j and at least $t + n - 2t + 1 = n - t + 1$ voters that rank c_j above c_ℓ , so c_j is the Condorcet winner of the resulting election. On the other hand, if we add at most $n - 2t$ new voters to E , in the resulting election there will be at least $n - t$ voters that prefer c_k to c_j and at most $t + n - 2t = n - t$ voters that prefer c_j to c_k , so in this case c_k prevents c_j from becoming the Condorcet winner. \square

Thus, the candidates with the highest Maximin score are exactly the candidates with the lowest insertion score. Together with Proposition 7, this implies the following result.

Theorem 11. *Maximin is (C, d_i) -rationalizable.*

The situation with the voter replacement rule is more complicated. Meskanen and Nurmi [2008] claim that Young's rule is (C, d_H) -rationalizable. As we have argued that the voter replacement rule is (C, d_H) -rationalizable, this would imply that the voter replacement rule is equivalent to Young's rule, or, in other words, deleting voters is equivalent to replacing voters. However, it turns out that this is not true.

Theorem 12. *There exists an election in which the voter replacement rule and Young's rule declare different candidates as winners.*

Proof. We construct an election $E = (C, V)$ with $C = \{a, b, c, d\}$ and $|V| = 29$. Among the first 5 voters in $|V|$, there are 2 voters with preference order $a \succ b \succ c \succ d$, 2 voters with preference order $a \succ c \succ d \succ b$, and 1 voter with preference order $a \succ b \succ d \succ c$.

Further, there are 8 voters with preferences $b \succ c \succ a \succ d$ (b -voters), 8 voters with preferences $c \succ d \succ a \succ b$

b (c -voters), and 8 voters with preferences $d \succ b \succ a \succ c$ (d -voters).

We summarize the numbers of voters that prefer x to y for $x, y \in \{a, b, c, d\}$ in the table below; we write $x > y : t$ to denote the fact that there are t voters that prefer x to y .

$$\begin{array}{llll} a > b : 13, & b > a : 16, & b > c : 19, & c > b : 10 \\ a > c : 13, & c > a : 16, & b > d : 11, & d > b : 18 \\ a > d : 13, & d > a : 16, & c > d : 20, & d > c : 9 \end{array}$$

Let us now compute $s_r(x)$ and $s_d(x)$ for $x \in \{b, c, d\}$. Candidate b wins pairwise elections against a and c , but loses to d by 7 votes. Hence, $s_d(b) \geq 8$. On the other hand, deleting 8 votes is sufficient: indeed, deleting all c -voters makes b the Condorcet winner. Thus, $s_d(b) = 8$. For the same reason, we need to replace at least 4 voters to make b the Condorcet winner (each replacement reduces d 's margin of victory over b by at most 2), and, indeed, replacing 4 of the c -voters with 4 voters that rank b first makes b the Condorcet winner. Hence, $s_d(b) = 4$. Similarly, c loses the pairwise election to b by 9 votes, so we have $s_d(c) \geq 10$, $s_r(c) \geq 5$ (we can show that, in fact, $s_d(c) = 10$ and $s_r(c) = 5$, but this is not needed for our proof), and d loses the pairwise election to c by 11 votes, so we have $s_d(d) \geq 12$, $s_r(d) \geq 6$.

Now, it is not hard to see that $s_r(a) \leq 3$: after we replace one b -voter, one c -voter and one d -voter with voters that rank a first, for each $x = b, c, d$ we have 15 voters that prefer a to x and 14 voters that prefer x to a . Thus, we have $s_r(a) < s_r(x)$ for $x = b, c, d$. To complete the proof, we will now argue that $s_d(a) > s_d(b)$. Specifically, we will show that $s_d(a) \geq 12$.

Indeed, it is clear that to make a the Condorcet winner, it is never optimal to delete any of the first five voters. Now, suppose that we can make a the Condorcet winner by deleting a set S of voters, $|S| < 12$. Suppose first that S contains at least 4 voters of a particular type (i.e., b -voters, c -voters, or d -voters); without loss of generality, we can assume that S contains 4 b -voters. After these voters have been deleted, a loses to d by at least 7 votes, so we need to delete at least 8 more voters, i.e., at least 12 voters altogether, a contradiction. Hence, we can now assume that S contains at most 3 voters of each type. Next, suppose that S contains exactly 3 voters of some type; again, without loss of generality we can assume that those are b -voters. After these voters have been deleted, a loses to d by 6 votes, so we have to additionally delete at least 7 other voters, i.e., at least 4 voters of some other type, a contradiction. Hence, S contains at most 2 voters of each type. Now, consider an arbitrary voter in S ; without loss of generality we can assume that this is

a b -voter. After this voter has been deleted, a loses to d by 4 votes, so we need to additionally delete at least 5 other voters, i.e., at least 3 voters of some other type, a contradiction. We conclude that $s_d(a) \geq 12$. \square

In fact, the voter replacement rule, despite having a very natural definition in terms of distances and consensus, appears not to be equivalent to any known voting rule and is therefore not studied in the existing literature. As a first step towards understanding the properties of this rule, we will now prove that, similarly to Young's rule, it is unlikely to have an efficient winner determination procedure. We omit the proof of the next theorem due to space limits.

Theorem 13. *Given an election $E = (C, V)$ and a candidate $p \in C$, it is NP-hard to decide if p is a winner of E under the voter replacement rule.*

For Young's rule, the winner determination problem is known to be complete for the complexity class Θ_2^P [Rothe *et al.*, 2003]. It seems likely that this is also the case for the voter replacement rule.

Observe that out of the three voting rules considered in this section, one (Maximin) has an efficient winner determination procedure, while the other two do not (assuming $P \neq NP$). The intuitive reason for this difference is that when we add voters to make a candidate c the Condorcet winner, we only need to add voters that rank c first, and, moreover, it does not matter how these voters rank other candidates. On the other hand, when we delete or replace voters, we have to choose which voters to remove, and this decision is not straightforward.

6 Conclusions and Future Research

We presented a number of new results on distance rationalizability of several well-known voting rules. A few questions suggest themselves for further study. First, while some of our results are negative, they only show that some voting rules cannot be rationalized with respect to a particular notion of consensus, but not that these rules cannot be distance-rationalized at all. It would be interesting to see if one can obtain results of this type. In particular, this would require formalizing the general idea of consensus. Similarly, it seems reasonable to expect that voting rules that are defined via a particular notion of consensus (such as, e.g., Dodgson's rule or Young's rule) cannot be rationalized using a different notion of consensus, such as unanimity; can we confirm this intuition?

A natural research direction is to seek distance-rationalizability results for further voting rules (e.g., it would be interesting to find to what extent gener-

alized scoring rules [Xia and Conitzer, 2008] are distance rationalizable) and to seek connections between distance rationalizability and related notions, such as explaining voting rules via maximum likelihood estimation [Conitzer and Sandholm, 2005; Conitzer *et al.*, 2009].

Another research direction concerns making connections between distance rationalizability and dishonest behavior in elections, such as control [Bartholdi *et al.*, 1992] and bribery [Faliszewski *et al.*, 2006]. Indeed, replacing voters to make a particular candidate the election winner, as we do in Section 5, is very similar to bribery (see [Faliszewski *et al.*, 2006]; Dodgson distance is also very similar to a recently introduced notion of swap bribery [Elkind *et al.*, 2009]), and adding or deleting voters or candidates is reminiscent of election control. We hope that making this intuition more precise will lead to interesting results for both areas.

Finally, distance rationalizability may provide a useful tool for understanding the computational complexity of various voting-related problems. In particular, it is likely that hardness/easiness results for a particular voting rule will often enable us to obtain similar hardness/easiness results for other rules that are defined via the same (or related) metric or the same notion of consensus. Indeed, the hardness result in Section 5 can be seen as an example of this approach.

Acknowledgments We would like to thank the anonymous TARK referees for their thorough and tremendously helpful work. Piotr Faliszewski was supported by AGH University of Science and Technology Grant no. 11.11.120.777. Edith Elkind was supported by ESRC (grant no. ES/F035845/1) and EP-SRC (grant no. GR/T10664/01).

References

- [Arrow, 1951 revised editon 1963] K. Arrow. *Social Choice and Individual Values*. John Wiley and Sons, 1951 (revised editon, 1963).
- [Baigent, 1987] N. Baigent. Metric rationalisation of social choice functions according to principles of social choice. *Mathematical Social Sciences*, 13(1):59–65, 1987.
- [Bartholdi *et al.*, 1992] J. Bartholdi, III, C. Tovey, and M. Trick. How hard is it to control an election? *Mathematical and Computer Modeling*, 16(8/9):27–40, 1992.
- [Brams and Fishburn, 2002] S. Brams and P. Fishburn. Voting procedures. In K. Arrow, A. Sen, and K. Suzumura, editors, *Handbook of Social Choice and Welfare, Volume 1*, pages 173–236, 2002.
- [Conitzer, 1785] J.-A.-N. de Caritat, Marquis de Condorcet. *Essai sur l’Application de L’Analyse à la Probabilité des Décisions Rendues à la Pluralité des Voix*. 1785. Facsimile reprint of original published in Paris, 1972, by the Imprimerie Royale.
- [Conitzer and Sandholm, 2005] V. Conitzer and T. Sandholm. Common voting rules as maximum likelihood estimators. In *Proceedings of the 21st Conference in Uncertainty in Artificial Intelligence*, pages 145–152. AUAI Press, July 2005.
- [Conitzer *et al.*, 2009] V. Conitzer, M. Rognlie, and L. Xia. Preference functions that score rankings and maximum likelihood estimation. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence*. AAAI Press, July 2009. To appear.
- [Dodgson, 1876] C. Dodgson. A method of taking votes on more than two issues. Pamphlet printed by the Clarendon Press, Oxford, and headed “not yet published”, 1876.
- [Elkind *et al.*, 2009] E. Elkind, P. Faliszewski, and A. Slinko. Swap bribery. Technical Report arXiv:0905.3885 [cs.GT], arXiv.org, May 2009.
- [Faliszewski *et al.*, 2006] P. Faliszewski, E. Hemaspaandra, and L. Hemaspaandra. The complexity of bribery in elections. In *Proceedings of the 21st National Conference on Artificial Intelligence*, pages 641–646. AAAI Press, July 2006.
- [Klamler, 2005a] C. Klamler. Borda and condorcet: Some distance results. *Theory and Decision*, 59(2):97–109, 2005.
- [Klamler, 2005b] C. Klamler. The copeland rule and condorcet’s principle. *Economic Theory*, 25(3):745–749, 2005.
- [Meskanen and Nurmi, 2008] T. Meskanen and H. Nurmi. Closeness counts in social choice. In M. Braham and F. Steffen, editors, *Power, Freedom, and Voting*. Springer-Verlag, 2008.
- [Moulin, 1991] H. Moulin. *Axioms of Cooperative Decision Making*. Cambridge University Press, 1991.
- [Rothe *et al.*, 2003] J. Rothe, H. Spakowski, and J. Vogel. Exact complexity of the winner problem for Young elections. *Theory of Computing Systems*, 36(4):375–386, 2003.
- [Xia and Conitzer, 2008] L. Xia and V. Conitzer. Generalized scoring rules and the frequency of coalitional manipulability. In *Proceedings of the 9th ACM Conference on Electronic Commerce*, pages 109–118. ACM Press, July 2008.