

Good Rationalizations of Voting Rules

Edith Elkind*

School of Physical and Mathematical
Sciences
Nanyang Technological University
Singapore

Piotr Faliszewski†

Department of Computer Science
AGH Univ. of Science and Technology
Kraków, Poland

Arkadii Slinko

Department of Mathematics
University of Auckland
Auckland, New Zealand

Abstract

We explore the relationship between two approaches to rationalizing voting rules: the maximum likelihood estimation (MLE) framework originally suggested by Condorcet and recently studied in (Conitzer and Sandholm 2005; Conitzer, Rognlie, and Xia 2009) and the distance rationalizability (DR) framework (Meskanen and Nurmi 2008; Elkind, Faliszewski, and Slinko 2009). The former views voting as an attempt to reconstruct the correct ordering of the candidates given noisy estimates (i.e., votes), while the latter explains voting as search for the nearest consensus outcome. We provide conditions under which an MLE interpretation of a voting rule coincides with its DR interpretation, and classify a number of classic voting rules, such as Kemeny, Plurality, Borda and Single Transferable Vote (STV), according to how well they fit each of these frameworks. The classification we obtain is more precise than the ones that result from using MLE or DR alone: indeed, we show that the MLE approach can be used to guide our search for a more refined notion of distance rationalizability and vice versa.

Introduction

Various aspects of voting, and, more generally, preference aggregation, are an active research topic in the artificial intelligence community. Indeed, voting can be used in a variety of applications that range from decision-making in multiagent planning (Ephrati and Rosenschein 1997) to ranking movies (Ghosh et al. 1999) to aggregating the outputs of web search engines (Dwork et al. 2001).

Voting has a rich history going back to ancient times, and, unsurprisingly, the human societies explored many different approaches to joint decision making, resulting in a number of *voting rules*, or algorithms for determining the best alternative or the optimal ordering of the alternatives. A natural question, then, is which voting rule is most appropriate for a given scenario. One can try to answer this question by choosing a rule that satisfies the *voting axioms* that are most

pertinent to the problem in hand, such as *monotonicity* or *Pareto-optimality*. However, if the voters are seen as cooperative entities that aim to aggregate information about the alternatives in the presence of errors—a viewpoint advocated by Condorcet and appropriate for many of the AI applications of voting—an attractive approach to choose a voting rule that has a natural interpretation within the Maximum Likelihood Estimation (MLE) framework, or its younger cousin, the distance rationalizability framework (DR).

The MLE framework is based on the idea that there is some objective ordering of the candidates from best to worst, and each vote is a noisy estimate of this ordering. Thus, the role of the voting rule is to find a ranking that is most likely to be the objectively correct one. Formally, in the MLE framework we assume that there exists a probability distribution of votes that is conditioned on the correct ranking, and that each voter draws her vote randomly and independently from this distribution. If a voting rule outputs a ranking that is most likely to be the correct one given the distribution, then we say that this rule is MLE.

The first voting rule that has been shown to fit the MLE framework is the Kemeny rule (Young and Levenglick 1978; Young 1988; 1995). Subsequently, Conitzer and Sandholm (2005) have proved that all scoring rules (a large class of voting rules that includes Plurality, Borda and Veto, among others) are MLE, while many other voting rules, such as Bucklin, Copeland or Maximin are not (the case of Single Transferable Vote (STV) is more complicated and has been fully analyzed in (Conitzer, Rognlie, and Xia 2009)).

A related, but different approach is to interpret voting as search for a consensus. This is the main idea behind the distance rationalizability framework introduced in (Meskanen and Nurmi 2008; Elkind, Faliszewski, and Slinko 2009). In this framework, we define a distance between elections, and seek the closest consensus election (an election with a single, clear winner) to the given election. A voting rule is said to be distance-rationalizable if it elects the winner of the nearest consensus, for some distance and some class of consensus elections (such as, e.g., elections where all voters agree on the ranking of candidates, or elections where all voters rank the same candidate first). Meskanen and Nurmi (2008), and Elkind, Faliszewski and Slinko (2009) show that almost all voting rules are distance-rationalizable, including some rules that are provably not MLE.

*Supported by NRF Research Fellowship (NRF-RF2009-08).

†Supported by AGH University of Technology Grant no. 11.11.120.865, by Polish Ministry of Science and Higher Education grant N-N206-378637, and by Foundation for Polish Science's program Homing/Powroty.
Copyright © 2010, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

However, not all existing MLE and distance-rationalizability results are equally appealing. For example, the noise model in the MLE interpretation of the Kemeny rule is extremely natural (a voter provides a correct ranking of each pair of candidates with a fixed probability $p > 0.5$), while for the scoring rules the noise model of Conitzer and Sandholm (2005) does not seem to have a simple interpretation. Nevertheless, the MLE framework alone does not provide us with a principled way of formalizing this intuition. For distance rationalizability, the need to refine the original framework of (Meskanen and Nurmi 2008; Elkind, Faliszewski, and Slinko 2009) is even more striking. Indeed, Elkind, Faliszewski, and Slinko (2010) show that unless we place additional restrictions on the type of distances used, essentially any voting rule can be distance-rationalized, rendering distance-rationalizability results meaningless. To remedy this, they propose to focus on rationalizing voting rules via so-called *votewise* distances, which first measure by how much each voter has to modify its vote to reach a consensus, and then aggregate these measurements in a uniform manner. Some, but not all distance rationalizability results can be shown to hold with respect to such distances.

An interesting property of the votewise distances is that their definition is syntactically very similar to that of simple ranking scoring functions (SRSFs) introduced by Conitzer, Rognlie and Xia (2009) in the context of the MLE framework. Specifically, Conitzer, Rognlie and Xia (2009) show that SRSFs are, in fact, equivalent to MLE rules. Thus, our first goal in this paper is to better understand the connection between SRSFs and distance rationalizability. It turns out that, while in general these notions are incomparable, we can identify additional constraints under which they become almost identical. Interestingly, the SRSF for the Kemeny rule satisfies these constraints, which means that Kemeny can be shown to be both MLE and DR via the *same* underlying function. On the other hand, for scoring rules, this is not the case: while they have an interpretation within both frameworks, this interpretation is substantially different. In other words, if we simply ask which rules are both MLE and DR, scoring rules are indistinguishable from Kemeny, but if we ask which rules can be represented as MLE and DR *in a consistent manner*, the Kemeny rule emerges as a better choice. Thus, the DR framework can be used to refine the MLE framework, i.e., to explain why some rules are better maximum likelihood estimators than others. The converse is also true: the very connection between SRSFs and DR is an additional argument in favor of votewise distances, as only such distances can be interpreted as SRSFs.

To illustrate this idea, in the second part of the paper, we consider four rules—Kemeny, Plurality, Borda, and STV—and rank them according to how well they can be represented as DR, MLE, or both. Our approach places all these rules in different categories, with Kemeny being the best, Plurality a close second, and STV failing the test completely. This allows us to conclude that combining MLE and DR leads to a better understanding of voting rules than either of these approaches on its own.

Preliminaries

Elections. An *election* is modelled as a pair $E = (C, V)$, where $C = \{c_1, \dots, c_m\}$ is a set of *candidates* and $V = (v_1, \dots, v_n)$ is a list of *voters*. Each voter v_i is described by a linear order \succ_i over C , called her *preference order*. A collection of preference orders is called a *preference profile*. When the set of candidates is fixed, we will sometimes identify E with V . We interpret \succ_i as the ranking of candidates according to the i -th voter. Thus, $a \succ_1 b \succ_1 c$ means that the first voter prefers a to b to c . For brevity, we will often write abc in place of $a \succ_i b \succ_i c$. For a set of candidates C , we denote by $\mathcal{L}(C)$ the set of all possible preference orders over C . Given a linear order v over a candidate set C and a permutation $\pi : C \rightarrow C$, let $\pi(v)$ denote the order obtained from v by replacing each candidate $c \in C$ with $\pi(c)$. We say that a function ϕ defined on $\mathcal{L}(C)^k$ for $k \geq 1$ is *neutral* if for any $\pi : C \rightarrow C$ and any $u_1, \dots, u_k \in \mathcal{L}(C)$ we have $\phi(u_1, \dots, u_k) = \phi(\pi(u_1), \dots, \pi(u_k))$. Neutrality is a very natural requirement in the context of voting, so from now on we assume that all functions on $\mathcal{L}(C)^k$ that we consider (voting rules, distances, noise models, etc.) are neutral.

Preference functions and voting rules. In this paper, we distinguish voting rules, i.e., mappings from elections to subsets of candidates, and preference functions, i.e., mappings from elections to subsets or candidate rankings. Formally, a *voting rule* is a mapping F that given an election $E = (C, V)$ outputs a set $W \subseteq C$ of *election winners*, and a *preference function* is a mapping f that given an election $E = (C, V)$ outputs a set $R = \{r_1, \dots, r_t\}$ of rankings from $\mathcal{L}(C)$. The interpretation here is that f views each of r_1, \dots, r_t as an equally good ranking of the candidates from C given votes V . Given a preference function f , we can construct a voting rule $F_f : E \rightarrow 2^C$ by setting $F_f(E) = \{c \in C \mid c \text{ is ranked first in some } r \in f(E)\}$.

A number of prominent preference functions are defined via families of so-called scoring protocols. A *scoring protocol* for m candidates can be identified with a vector $(\alpha_1, \dots, \alpha_m)$ that satisfies $\alpha_1 \geq \dots \geq \alpha_m \geq 0$. Under this protocol, a candidate c receives α_j points from each voter that puts c in the j -th position in her ranking. Given an election (C, V) with $|C| = m$, the corresponding preference function f_α outputs a set of linear orders that rank members of C in the order of decreasing number of points (there may be many rankings satisfying this condition if some candidates have the same number of points). We will focus on two most prominent families of scoring protocols: Plurality and Borda. Plurality is defined via scoring protocols of the form $(1, 0, \dots, 0)$, i.e., under Plurality candidates get points for being ranked first only. For elections with m candidates, Borda is defined via vector $(m-1, m-2, \dots, 0)$.

In Single Transferable Vote (STV) preference function the rankings are created as follows: We find a candidate with the lowest Plurality score, remove him from the votes, place him on the last available position in the output ranking, and repeat the process with the modified votes until all candidates are processed. For STV the issue of handling ties—that is, the issue of the order in which candidates with equal Plurality scores are handled—is quite important, and is discussed

in detail by Conitzer, Rognlie and Xia (2009); however, our results are independent of the choice of a tie-breaking rule.

Finally, given an election $E = (C, V)$ with $|V| = n$, Kemeny's preference function outputs the rankings \succ that minimize the expression $\sum_{i=1}^n d_s(\succ_i, \succ)$, where $d_s(\succ_i, \succ)$ is the number of swaps of adjacent candidates needed to transform \succ into \succ_i . (Equivalently, $d_s(\succ_i, \succ)$ is the number of inversions between \succ and \succ_i .)

MLEs, SRSFs and DR

We will now formally define the two approaches to thinking about voting rules that are discussed in this paper, i.e., the maximum likelihood estimation (MLE) framework and the distance rationalizability (DR) framework, as well as the notion of simple ranking scoring functions (SRSFs) that provides a bridge between them. Let us fix a candidate set C throughout this section.

MLEs and SRSFs. The following overview is based on (Conitzer, Rognlie, and Xia 2009). For each $v, r \in \mathcal{L}(C)$, a *noise model* ν specifies a conditional probability $P_\nu(v|r)$, that is, the probability that a voter submits a ranking v given that the "correct" ranking is r . We say that a preference function f can be interpreted as a *maximum likelihood estimator (MLE)* if there exists a noise model ν such that for each preference profile $V = (v_1, \dots, v_n)$ over C , it holds that $f(C, V) = \arg \max_{r \in \mathcal{L}(C)} \prod_{i=1}^n P_\nu(v_i|r)$. Intuitively, this definition assumes that the votes are distributed according to ν in an i.i.d. fashion.

A preference function f is a *simple ranking scoring function* if there exists a function $s_f : \mathcal{L}(C) \times \mathcal{L}(C) \rightarrow \mathbb{R}^+ \cup \{0\}$ such that for each collection $V = (v_1, \dots, v_n)$ of voters over C we have $f(C, V) = \arg \min_{r \in \mathcal{L}(C)} \sum_{i=1}^n s_f(v_i, r)$.¹ Via a slight abuse of notation, we will often refer to the function s_f itself as the simple ranking scoring function. Each preference function f that can be interpreted as an MLE is an SRSF: if f is MLE via a noise model ν , we let $s_f(v, u) = -\ln(P_\nu(v|u))$ for any $u, v \in \mathcal{L}(C)$; the converse is also true.

Preference functions that can be interpreted as MLE include the Kemeny rule and all scoring rules.

Distance rationalizability. The definition of distance rationalizability has two main ingredients: a notion of distance and a notion of consensus. A *distance* d (or, metric \hat{d}) over some set X is a function $d : X \times X \rightarrow \mathbb{R}$ such that for each $x, y, z \in X$ it holds that (a) $d(x, y) \geq 0$, (b) $d(x, y) = 0$ if and only if $x = y$, (c) $d(x, y) = d(y, x)$, and (d) $d(x, y) + d(y, z) \geq d(x, z)$. The last condition is called the *triangle inequality*. If d satisfies all conditions except (b), then d is called a *pseudodistance*.

A *consensus* is a set of elections with a clear winner. The three most standard consensus classes are the *strong unanimity consensus* \mathcal{S} , which consists of all elections in which all voters rank the candidates in the same way, the *weak unanimity consensus* \mathcal{U} , which consists of all elections in which

all voters rank the same candidate first, and the *Condorcet consensus* \mathcal{C} , which consists of all elections that have a *Condorcet winner*, i.e., a candidate that would beat any other candidate in a pairwise election. Additionally, the *majority consensus* \mathcal{M} consists of all elections in which a majority of voters ranks the same candidate first.

The following definition is adapted from (Elkind, Faliszewski, and Slinko 2009). A voting rule F is said to be *distance-rationalizable (DR)* with respect to a consensus class $\mathcal{X} \in \{\mathcal{S}, \mathcal{U}, \mathcal{M}, \mathcal{C}\}$ if for any $n \geq 1$ there is a distance d over $\mathcal{L}^n(C)$ such that for each collection $V = (v_1, \dots, v_n)$ of voters, a candidate is a winner in (C, V) under F if and only if he is a winner in a nearest (with respect to d) election in \mathcal{X} . The definition above differs from the one given in (Elkind, Faliszewski, and Slinko 2009) in that it requires the distance to be defined on profiles of the same length; for our purposes, this distinction is irrelevant.

Recall that a *norm* on a vector space S over \mathbb{R} is a mapping $N : S \rightarrow \mathbb{R}$ that satisfies (a) $N(\alpha u) = |\alpha|N(u)$ for any $\alpha \in \mathbb{R}$, $u \in S$, (b) $N(u + v) \leq N(u) + N(v)$ for any $u, v \in S$, and (c) $N(u) = 0$ if and only if u is the zero vector. The class of votewise distances introduced in (Elkind, Faliszewski, and Slinko 2010) consists of all product metrics obtained by composing a distance $d : \mathcal{L}(C) \times \mathcal{L}(C) \rightarrow \mathbb{R}^+ \cup \{0\}$ over individual votes and a norm N on \mathbb{R}^n . Formally, a distance $\hat{d} : \mathcal{L}^n(C) \times \mathcal{L}^n(C) \rightarrow \mathbb{R}^+ \cup \{0\}$ is said to be *votewise* if there exist a distance $d : \mathcal{L}(C) \times \mathcal{L}(C) \rightarrow \mathbb{R}^+ \cup \{0\}$ and a norm $N : \mathbb{R}^n \rightarrow \mathbb{R}^+ \cup \{0\}$ such that for any $\mathbf{u} = (u_1, \dots, u_n)$, $\mathbf{v} = (v_1, \dots, v_n)$ we have $\hat{d}(\mathbf{u}, \mathbf{v}) = N(d(u_1, v_1), \dots, d(u_n, v_n))$. It is said to be *additively votewise* if N is the ℓ_1 -norm, i.e., $\hat{d}(\mathbf{u}, \mathbf{v}) = d(u_1, v_1) + \dots + d(u_n, v_n)$. Note that we can define \hat{d} in the same manner for an arbitrary function $d : \mathcal{L}(C) \times \mathcal{L}(C) \rightarrow \mathbb{R}^+ \cup \{0\}$, i.e., d need not to be a metric.

Elkind, Faliszewski and Slinko (2010) show that essentially any voting rule is DR with respect to \mathcal{S} . However, the distance used in their construction is not votewise. The rules that are known to be DR via votewise distances include a variant of the Bucklin rule, the Dodgson rule, the Kemeny rule, Plurality, and all "good" scoring rules, i.e., those with $\alpha_i \neq \alpha_j$ for any $i \neq j$; for all of these rules except for Bucklin, the corresponding distance is additively votewise. These DR results make use of all four consensus classes listed above: Bucklin is DR with respect to \mathcal{M} , Dodgson is DR with respect to \mathcal{C} , Kemeny is DR with respect to \mathcal{S} , and Plurality and the "good" scoring rules are DR with respect to \mathcal{U} .

SRSF vs DR. There is a remarkable similarity between the definition of a simple ranking scoring function and that of a distance-rationalizable voting rule. However, in general the two notions are incomparable.

First, in the definition of SRSF, the score of a profile is obtained as a sum of individual scores, while the definition of distance-rationalizability allows arbitrary distances. Thus, for the purposes of the comparison, we need to focus on rules that are DR via additively votewise distances. Note that Elkind, Faliszewski and Slinko (2010) argue that we should restrict ourselves to votewise distances when proving DR

¹Note that Conitzer, Rognlie and Xia (2009) use $\arg \max$ instead of $\arg \min$ in their definition of SRSFs; our definition is equivalent and more natural in our setting.

results; the comparison with SRSFs provides another argument in favor of that position. Second, in the definition of an SRSF, we try to minimize the sum of scores with respect to a single ranking, while in the definition of a DR rule we compute the distance to a consensus, i.e., a *profile* of rankings. This leads us to another refinement of the notion of distance-rationalizability: namely, rationalizability with respect to the consensus class \mathcal{S} . Indeed, a strong unanimity consensus can be represented by a single vote, so finding a vote u that minimizes the expression $\sum_{i=1}^n s_f(v_i, u)$ is equivalent to finding a profile $\mathbf{u} \in \mathcal{S}$ that minimizes $\widehat{s}_f((v_1, \dots, v_n), \mathbf{u})$, where \widehat{s}_f is the additively votewise function that corresponds to s_f . With these constraints in place, given a voting rule that is rationalized with respect to \mathcal{S} via some distance d on votes, we can form a noise model so that for each $u, v \in \mathcal{L}(C)$, $P_d(u|v) = ce^{-d(u,v)}$, where c is a normalization constant. This noise model almost proves that f is an MLE: It leads to rankings with correct candidate ranked first, but makes no guarantees as to how further candidates are ranked (recall that the definition of DR applies to voting rules, not to preference functions).

Conversely, consider a preference function F_f that corresponds to an SRSF function s_f . For the voting rule F_f to be distance-rationalizable via \widehat{s}_f , the function s_f needs to be a metric. However, the definition of an SRSF imposes no restrictions on s_f : in particular, it can be asymmetric, or fail the triangle inequality.

Nevertheless, there exists a voting rule for which we can show that it is both SRSF and additively votewise DR by using the same function, namely, the Kemeny rule! Indeed, it is not hard to see that the function d_s in the definition of Kemeny rule is a metric. Thus, the Kemeny rule can be consistently explained in both frameworks. It is interesting to ask if other voting rules also have this property; arguably, such rules provide the most principled approach to preference aggregation. Thus, in the next section, we subject three classic voting rules—Plurality, Borda, and STV—to this test.

Main Results

In this section we implement the program outlined in the introduction and in the preceding section. Specifically, we show that (a) Plurality is additively votewise DR with respect to \mathcal{S} , but not via any of its SRSFs, (b) Borda is not additively votewise DR with respect to \mathcal{S} , and (c) STV is not votewise DR with respect to \mathcal{S}, \mathcal{U} or \mathcal{C} , even if we allow a very general class of norms instead of ℓ_1 .

Our proofs proceed by constructing counterexamples for the special case of three candidates. Unless stated otherwise, we assume that the candidate set is $C = \{a, b, c\}$.

Consider a distance d over C . By symmetry and neutrality, d is completely described by its values on the pairs (abc, abc) , (abc, acb) , (abc, bac) , (abc, bca) , (abc, cab) , and (abc, cba) . Further, we have $d(abc, abc) = 0$, and by neutrality and symmetry we have $d(abc, bca) = d(abc, cab)$ (to see this, note that the permutation π given by $\pi(a) = c$, $\pi(b) = a$, $\pi(c) = b$ transforms abc into cab and bca into abc). Set $d(abc, acb) = T$, $d(abc, bac) = B$, $d(abc, cba) = C$, $d(abc, bca) = d(abc, cab) = S$. Table 1 gives the values

	abc	acb	bac	bca	cab	cba
abc	0	T	B	S	S	C
acb	T	0	S	C	B	S
bac	B	S	0	T	C	S
bca	S	C	T	0	S	B
cab	S	B	C	S	0	T
cba	C	S	S	B	T	0

Table 1: The values of d for each pair of votes over C .

of d for each pair of preference orders.² Since d is a distance, we have $T, C, B, S > 0$. For a collection $V = (v_1, \dots, v_n)$ of voters and a ranking r , by $\widehat{d}(V, r)$ we mean $\sum_{i=1}^n d(v_i, r)$.

Plurality. Meskanen and Nurmi (2008) show that Plurality is DR with respect to \mathcal{U} . The distance employed in their construction is additively votewise. Further, Plurality is rationalizable with respect to \mathcal{S} via an additively votewise pseudodistance: for any two votes $u, v \in \mathcal{L}(C)$, $u \neq v$, we can set $d(u, v) = 0$ if u and v rank the same candidate first and $d(u, v) = 1$ otherwise. We can strengthen the latter result to additively votewise distance-rationalizability.

Theorem 1. *Plurality is additively votewise DR with respect to \mathcal{S} .*

Proof. We start with the construction for three candidates. Let d be the distance given by $S = T = 1$, $B = C = 2$. We claim that \widehat{d} rationalizes Plurality with respect to \mathcal{S} .

To see this, consider an election $E = (C, V)$ with $C = \{a, b, c\}$ that has a_1 voters with preferences abc , a_2 voters with preferences acb , b_1 voters with preferences bca , b_2 voters with preferences bac , c_1 voters with preferences cab , and c_2 voters with preferences cba .

We have $\widehat{d}(V, abc) = a_2 + b_1 + c_1 + 2b_2 + 2c_2$, $\widehat{d}(V, acb) = a_1 + b_2 + c_2 + 2b_1 + 2c_1$. Thus, the distance from V to the nearest profile in \mathcal{S} with winner a is $\min\{a_2 + b_2 + c_2, a_1 + b_1 + c_1\} + (b_1 + c_1 + b_2 + c_2)$. Symmetrically, the distance from V to the nearest profile in \mathcal{S} with winner b is $\min\{a_2 + b_2 + c_2, a_1 + b_1 + c_1\} + (a_1 + c_1 + a_2 + c_2)$, and the distance from V to the nearest profile in \mathcal{S} with winner c is $\min\{a_2 + b_2 + c_2, a_1 + b_1 + c_1\} + (a_1 + b_1 + a_2 + b_2)$. We observe that the first component of these expressions is identical, and the second component counts the number of voters that do not rank a (respectively, b , c) first. Thus, the set of Plurality winners coincides with the set of winners in the nearest strong consensus profiles with respect to \widehat{d} .

We will now extend this construction to any number of candidates. Fix $C = \{c_1, \dots, c_m\}$. For two votes $u, v \in \mathcal{L}(C)$, we say that v can be obtained from u by a *cyclic shift* if there exists an $i \in [m]$ and a permuta-

²If a vote u is obtained from a vote v by permuting the second and the third candidate (and leaving the top candidate in place), we have $d(u, v) = T$, if u is obtained from v by permuting the first and the second candidate (and leaving the bottom candidate in place), we have $d(u, v) = B$, if u is obtained from v by permuting the first and the third candidate (and leaving the center in place), we have $d(u, v) = C$, and if u is obtained from v by a cyclic shift, we have $d(u, v) = S$.

tion $\pi : C \rightarrow C$ such that $v = \pi(c_1) \succ \dots \succ \pi(c_m)$, $u = \pi(c_i) \succ \dots \succ \pi(c_m) \succ \pi(c_1) \succ \dots \succ \pi(c_{i-1})$.

Partition $\mathcal{L}(C)$ into m groups L_1, \dots, L_m , where the voters in L_i rank c_i on top. Set $s = (m-1)!$ and, for each $i \in [m]$, number the votes in L_i as v_i^1, \dots, v_i^s so that for any $i, j \in [m]$ the vote v_j^t can be obtained from the vote v_i^t by a cyclic shift. This is possible, since for each $u_i^t, i \in [m], t \in [s]$ and each $j \in [m]$, there is exactly one vote in L_j that can be obtained from v_i^t by a cyclic shift. Now, set $d(v_i^t, v_j^r) = 1$ if either $i = j$ or $t = r$, but $(i, t) \neq (j, r)$, and set $d(v_i^t, v_j^r) = 2$ if $i \neq j, t \neq r$. Observe that since $d(u, v) \in \{1, 2\}$ for $u \neq v$, the mapping d satisfies the triangle inequality; it is also symmetric and neutral.

Consider a preference profile V . For any $i \in [m], t \in [s]$, let a_i^t denote the number of voters in V with preferences v_i^t .

We have $\hat{d}(V, v_i^t) = \sum_{r \in [s] \setminus \{t\}} a_i^r + \sum_{j \in [m] \setminus \{i\}} a_j^t + 2 \sum_{j \in [m] \setminus \{i\}} \sum_{r \in [s] \setminus \{t\}} a_j^r = \sum_{j \in [m] \setminus \{i\}} \sum_{r \in [s]} a_j^r + \sum_{j \in [m]} \sum_{r \in [s] \setminus \{t\}} a_j^r$. Consequently, the distance from V to the nearest profile in \mathcal{S} with winner c_i is given by $\sum_{j \in [m] \setminus \{i\}} \sum_{r \in [s]} a_j^r + \min_{t \in [s]} \sum_{j \in [m]} \sum_{r \in [s] \setminus \{t\}} a_j^r$. The second component of this expression does not depend on i , while its first component counts the number of voters that do not rank c_i first. Thus, the nearest strong unanimity consensus to V has c_i as a winner if and only if i minimizes the sum $\sum_{j \in [m] \setminus \{i\}} \sum_{r \in [s]} a_j^r$ over all $i \in [m]$, i.e., c_i has the largest number of first-place votes. Thus, Plurality is distance-rationalizable with respect to \mathcal{S} via \hat{d} . \square

Yet, it is impossible to set the values B, C, S, T so that d rationalizes Plurality with respect to \mathcal{S} in such a way that the nearest \mathcal{S} -consensus orders the candidates by their Plurality scores: If d is a distance that additively rationalizes Plurality with respect to \mathcal{S} , then d is not an SRSF for Plurality.

Indeed, let $k \geq 2$ be an integer, and consider a collection V of $2k-1$ voters where k voters have preference order acb and $k-1$ voters have preference order bca . The Plurality scores of a, b , and c are, respectively, $k, k-1$, and 0 . Thus, the nearest ranking should be abc . However, it is impossible to set B, C, S , and T to ensure this, while keeping d a metric. To see this, note that $d(V, abc) = kT + (k-1)S$ and $d(V, acb) = (k-1)C$. We would have to have $T + (k-1)T + (k-1)S < (k-1)C$. Yet, this is impossible, because by triangle inequality we have $d(acb, abc) + d(abc, bca) \geq d(acb, bca)$, that is, $T + S \geq C$.

Borda. Like all scoring rules, the Borda rule is an SRSF. Further, it is known to be DR with respect to unanimity consensus \mathcal{U} via an additively votewise distance (Meskanen and Nurmi 2008). In fact, the distance used in this construction is just the distance d_s that rationalizes the Kemeny rule with respect to \mathcal{S} , i.e., Borda and Kemeny are rationalized via the same distance, but with respect to different consensus classes. However, the SRSF for Borda is not a distance. Our next result explains why this is the case.

Theorem 2. *For three candidates, Borda is not DR with respect to \mathcal{S} via a neutral additively votewise distance.*

Proof. Suppose that Borda is additively votewise DR with

respect to \mathcal{S} via a distance d on votes given by T, C, B , and S , and consider two families of preference profiles, $V_1(k)$ and $V_2(k)$, where $k > 0$. $V_1(k)$ and $V_2(k)$ both contain k voters with preference order acb , k voters with preference order bca , and one extra voter. In the case of $V_1(k)$ this extra voter has preference order cab , and in the case of $V_2(k)$, the extra voter has preference acb . We have

$$\begin{aligned} d_1 &= \hat{d}(V_1(k), cab) = kB + kS, \\ d_2 &= \hat{d}(V_1(k), cba) = kS + kB + T, \\ d_3 &= \hat{d}(V_1(k), acb) = kC + B, \\ d_4 &= \hat{d}(V_1(k), bca) = kC + S. \end{aligned}$$

Naturally, $d_1 < d_2$. Further, for each $k > 0$, the unique winner of $V_1(k)$ is c , and the unique winner of $V_2(k)$ is a . Therefore, for any $k > 0$ it holds that $d_1 < d_3$, and $d_1 < d_4$, that is, in particular, for each k we have $kB + kS - B < kC$. On the other hand, for $V_2(k)$ we have:

$$\begin{aligned} d'_1 &= \hat{d}(V_2(k), cab) = (k+1)B + kS, \\ d'_2 &= \hat{d}(V_2(k), cba) = (k+1)S + kB, \\ d'_3 &= \hat{d}(V_2(k), acb) = kC, \\ d'_4 &= \hat{d}(V_2(k), abc) = (k+1)T + kS. \end{aligned}$$

By triangle inequality, $T + S \geq C$, so $d'_3 \leq d'_4$. Since a wins in $V_2(k)$, we have $d'_3 < d'_1$ and $d'_3 < d'_2$. In particular, for each $k > 0$ it holds that $kC < kB + kS + B$. By combining this inequality with the previous one, we get that for each $k > 0$ it holds that $B + S - \frac{B}{k} < C < B + S + \frac{B}{k}$. Since k can be arbitrarily large, we have $C = B + S$.

Now, consider a preference profile $V_3 = (abc, acb)$. Clearly, a is the unique Borda winner for V_3 . We have $d''_1 = \hat{d}(V_3, abc) = T$, $d''_2 = \hat{d}(V_3, acb) = T$, and $d''_3 = \hat{d}(V_3, cab) = S + B$. Since a is the unique winner, it holds that $T < S + B$. In particular, $T < B$. However, by triangle inequality we know that $T + S \geq C$. We also know that $C = B + S$, so we have $T + S \geq B + S$, that is, $T \geq B$. This is a contradiction. \square

Single Transferable Vote. Conitzer, Rognlie, and Xia (2009) have shown that STV is not MLE. By our observations regarding the relationship between MLE and DR, this implies that STV cannot be rationalized with respect to \mathcal{S} via an additively votewise distance. We will now strengthen this result to (almost) arbitrary votewise distances. We need the following definition.

Definition 1 (Bauer, Stoer, and Witzgall 1961). *A norm N in \mathbb{R}^n is monotonic in the positive orthant, or \mathbb{R}_+^n -monotonic, if for any two vectors $(x_1, \dots, x_n), (y_1, \dots, y_n) \in \mathbb{R}_+^n$ such that $x_i \leq y_i$ for all $i = 1, \dots, n$ we have $N(x_1, \dots, x_n) \leq N(y_1, \dots, y_n)$.*

Bauer, Stoer, and Witzgall (1961) provide a discussion of norms that are monotonic in the positive orthant. We remark that this is a fairly weak notion of monotonicity: the class of \mathbb{R}_+^n -monotonic norms strictly contains the class of all *monotonic* norms (as defined in (Bauer, Stoer, and Witzgall 1961)). The requirement of \mathbb{R}_+^n -monotonicity is very

natural when the norm in question is to be used to construct a product metric, as in our case.

We say that a votewise distance is *monotonic* if the respective norm is monotonic in the positive orthant.

Theorem 3. *For three candidates, STV (together with any intermediate tie-breaking rule) is not distance-rationalizable with respect to the strong unanimity and any neutral anonymous monotonic votewise distance.*

Proof. For the sake of contradiction, suppose that STV can be rationalized with respect to \mathcal{S} via a neutral anonymous monotonic votewise distance \widehat{d} , and let N denote the corresponding norm. Consider a profile V with $2k + 1$ voters, $k \geq 2$, where the voters' preferences are as follows: k voters have preferences given by abc , k voters have preferences given by bca , one voter has preferences given by cab . We have

$$\begin{aligned} d_1 &= \widehat{d}(V, abc) = N(0, \dots, 0, S, \dots, S, S), \\ d_2 &= \widehat{d}(V, acb) = N(T, \dots, T, C, \dots, C, B), \\ d_3 &= \widehat{d}(V, bca) = N(S, \dots, S, 0, \dots, 0, S), \\ d_4 &= \widehat{d}(V, bac) = N(B, \dots, B, T, \dots, T, C). \end{aligned}$$

Clearly, under STV candidate a is the unique winner in V . Thus, it must be the case that $\min\{d_1, d_2\} < \min\{d_3, d_4\}$. By symmetry we have $d_1 = d_3$, and hence $d_2 < d_4$. Also, by symmetry we get $d_4 = N(T, \dots, T, C, B, \dots, B)$. Hence, by monotonicity $C < B$.

Now, consider the profile W obtained by replacing the last voter in V by a voter whose preferences are cba . We have

$$\begin{aligned} d'_1 &= \widehat{d}(W, abc) = N(0, \dots, 0, S, \dots, S, C), \\ d'_2 &= \widehat{d}(W, acb) = N(T, \dots, T, C, \dots, C, S), \\ d'_3 &= \widehat{d}(W, bca) = N(S, \dots, S, 0, \dots, 0, B), \\ d'_4 &= \widehat{d}(W, bac) = N(B, \dots, B, T, \dots, T, S). \end{aligned}$$

In W , the STV-winner is b , so we have $\min\{d'_1, d'_2\} > \min\{d'_3, d'_4\}$. Furthermore, by symmetry, we have $d'_3 = N(0, \dots, 0, S, \dots, S, B)$. As $C < B$, by monotonicity we conclude that $d'_1 \leq d'_3$. This implies that $d'_2 > d'_4$. However, by symmetry we have $d'_4 = N(T, \dots, T, B, \dots, B, S)$, so by monotonicity $d'_2 \leq d'_4$, a contradiction. \square

We can use similar ideas to show that STV is not distance-rationalizable with respect to \mathcal{U} and a neutral anonymous monotonic votewise distance.

Theorem 4. *For three candidates, STV (together with any intermediate tie-breaking rule) is not distance-rationalizable with respect to weak unanimity and any neutral anonymous monotonic votewise distance.*

Finally, we remark that STV is not DR with respect to \mathcal{C} since it is not Condorcet-consistent. Note that Meskanen and Nurmi (2008) show that STV can be distance-rationalized with respect to \mathcal{U} . Their distance is neutral, but not votewise (note that the notions of anonymity and monotonicity are only defined for votewise distances, so they are not applicable here). Further, it is not immediately clear if this distance is polynomial-time computable. Thus, of all rules we have considered, STV is DR in the weakest possible sense.

Conclusions

Maximum likelihood estimation and distance rationalizability provide two related, but distinct approaches to understanding voting rules. In this paper, we have classified several prominent voting rules according to how well they can be explained in each of these frameworks. According to this criterion, the best voting rule is Kemeny rule, which can be shown to fit both the MLE framework and the DR framework via the same underlying function. Plurality is a close second: it is both an SRSF and DR via a “good” distance with respect to our strongest consensus class \mathcal{S} (as well as a weaker consensus \mathcal{U}); however, we show that one cannot use the same function to prove both of these results. Borda rule, too, is both an SRSF and additively votewise DR, but only with respect to a weaker consensus class \mathcal{U} . Finally, STV fails this test completely: we show that it cannot be rationalized via a votewise distance with respect to any of the standard consensus classes. Thus, our approach provides a more refined classification of voting rules than either MLE or DR alone.

References

- Bauer, F.; Stoer, J.; and Witzgall, C. 1961. Absolute and monotonic norms. *Numerische Mathematik* 3:257–264.
- Conitzer, V., and Sandholm, T. 2005. Common voting rules as maximum likelihood estimators. In *Proc. of UAI-05*, 145–152.
- Conitzer, V.; Rognlie, M.; and Xia, L. 2009. Preference functions that score rankings and maximum likelihood estimation. In *Proc. of IJCAI-09*, 109–115.
- Dwork, C.; Kumar, R.; Naor, M.; and Sivakumar, D. 2001. Rank aggregation methods for the web. In *Proc. of WWW-01*, 613–622.
- Elkind, E.; Faliszewski, P.; and Slinko, A. 2009. On distance rationalizability of some voting rules. In *Proc. of TARK-09*, 108–117.
- Elkind, E.; Faliszewski, P.; and Slinko, A. 2010. On the role of distances in defining voting rules. In *Proc. of AAMAS-10*. To appear.
- Ephrati, E., and Rosenschein, J. 1997. A heuristic technique for multi-agent planning. *Annals of Mathematics and Artificial Intelligence* 20(1–4):13–67.
- Ghosh, S.; Mundhe, M.; Hernandez, K.; and Sen, S. 1999. Voting for movies: The anatomy of recommender systems. In *Proc. of Agents-99*, 434–435.
- Meskanen, T., and Nurmi, H. 2008. Closeness counts in social choice. In Braham, M., and Steffen, F., eds., *Power, Freedom, and Voting*. Springer-Verlag.
- Young, H., and Levenglick, A. 1978. A consistent extension of Condorcet's election principle. *SIAM Journal on Applied Mathematics* 35(2):285–300.
- Young, H. 1988. Condorcet's theory of voting. *American Political Science Review* 82(2):1231–1244.
- Young, H. 1995. Optimal voting rules. *Journal of Economic Perspectives* 9(1):51–64.