Distance Rationalization of Voting Rules[1]

Edith Elkind, Piotr Faliszewski, and Arkadii Slinko

Abstract

The concept of *distance rationalizability* allows one to define new voting rules or "rationalize" existing ones via a consensus class of elections and a distance. A consensus class consists of elections in which there is a consensus in the society who should win. A distance measures the deviation of the actual election from consensus elections. Together, a consensus class and a distance define a voting rule: a candidate is declared an election winner if she is the consensus candidate in one of the nearest consensus elections. It is known that many classic voting rules are defined in this way or can be represented via a consensus class and a distance, i.e., *distance-rationalized*. In this paper, we focus on the power and the limits of the distance rationalizability approach. We first show that if we do not place any restrictions on the class of possible distances then essentially all voting rules are distance-rationalizable. Thus, to make the concept of distance ratioanalizability meaningful, we have to restrict the class of distances involved. To this end, we present a very natural class of distances, which we call *votewise* distances. We investigate which voting rules can be rationalized via votewise distances and study the properties of such 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, e.g., to select a unique alternative from a space of options available to them, even though the agents may have different opinions about the relative merits of these alternatives. 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 reflects the individual preferences of voters as well as possible, there is no universal agreement on how to reach this goal. As a consequence, there is a multitude of voting rules, and these rules are remarkably diverse (see, e.g., [4]).

Why cannot we settle on a single voting rule, which will aggregate the preferences optimally? One answer to this question is provided by the long list of impossibility theorems—starting with the famous Arrow's impossibility theorem [1]—which state that there is no voting rule (or a social welfare function) that simultaneously satisfies several natural desiderata. Thus in each real-life scenario we have to decide which of desired conditions we are willing to sacrifice.

An earlier view, initiated by Marquis de Condorcet, is that a voting rule must be a method for aggregating information. Voters have different opinions because they make errors of judgment; absent these errors, they would all agree on the best choice. The goal is to design a voting rule that identifies the best choice with highest probability. This approach is called *maximum likelihood estimation* and it has been actively pursued by Young who showed [22] that consistent application of Condorcet's ideas leads to the Kemeny rule [14]. It has been shown since then that several other voting rules can be obtained as maximum likelihood estimators for different models of errors (see Conitzer, Rognlie, and Xia [6] and

[1] This paper combines three earlier papers by the same authors: "On Distance Rationalizability of Some Voting Rules" (presented at TARK-2009), "On the Role of Distances in Defining Voting Rules" (presented at AAMAS-2010), and "Good Rationalizations of Voting Rules" (presented at AAAI-2010).

Conitzer and Sandholm [7]).

The third approach that has emerged recently in a number of papers (see, e.g., Baigent [2] and Meskanen and Nurmi [19]) can be called consensus-based. The result of each election is viewed as an imperfect approximation to some kind of electoral consensus. Under this view, the winner of a given election, 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 several ways of defining closeness. The heart of this approach is the decision which situations should be viewed as "electoral consensuses", be it the existence of Condorcet winner, universal agreement on which candidate is best, or something else. The concept of closeness should also be agreed upon. This approach is ideologically close to bargaining.

In this paper we concentrate on the third approach. To date, the most complete list of distance-rationalizable rules is provided by Meskanen and Nurmi [19] (but see also [2, 16, 15]). There, the authors show how to distance-rationalize many voting rules, including, among others, Plurality, Borda, Veto, Copeland, Dodgson, Kemeny, Slater, and STV. However, in Section 3 we show that the usefulness of these results is limited, as essentially every reasonable voting rule can be distance rationalized with respect to some distance and some notion of consensus. This indicates that the notion of distance rationalizability used in the early work is too broad to be meaningful. Hence, we have to determine what are the "reasonable" consensus classes and the "reasonable" distances and to reexamine all existing results.

In Section 4 we suggest a family of "good" distances (which we call *votewise* distances) and study voting rules that are distance rationalizable with respect to such distances. In particular, in Section 4.2 we show that many of the rules considered in [19], as well as all scoring rules and a variant of the Bucklin rule, can be rationalized via distances from this family. In contrast, we demonstrate that STV, which was shown to be distance-rationalizable in [19], is not distance-rationalizable via votewise distances, i.e., the restricted notion of distance rationalizability is indeed meaningful.

Now, the distance rationalizability framework can be viewed as a general method for specifying and analyzing voting rules. As such, it may be useful for proving results for entire families of voting rules, rather than isolated rules. For instance, a lot of recent research in computational social choice has focused on the complexity of determining (possible) election winners (see, e.g., [11, 17]), and the complexity of various types of attacks on elections (e.g., manipulation [8], bribery [9], and control [18, 10]).[2] However, most of the results in this line of work are specific to particular voting rules. We believe that the ability to describe multiple voting rules in a unified way (e.g., via the distance rationalizability framework) will lead to more general results. To provide an argument in favor of this belief, in Sections 4.1 and Section 4.3 we present initial results of this type, relating the type of distance and consensus used to rationalize a voting rule with the complexity of winner determination under this rule as well as the rule's axiomatic properties (such as anonymity, neutrality and consistency).

Due to space restrictions, all proofs are omitted. However, the reader may find many of them in the conference papers on which this paper is based (see the title footnote).

2 Preliminaries

2.1 Elections. An *election* is a pair $E = (C, V)$ where $C = \{c_1, \dots, c_m\}$ is the set of *candidates* and $V = (v_1, \dots, v_n)$ is an ordered list of *voters*. Each voter is represented by her *vote*, i.e., a strict, linear order over the set of candidates (also called a *preference order*).

[2] These references are only examples; an overview of literature is far beyond the scope of this paper.

We will refer to the list V as a *preference profile*, and we denote the number of voters in V by $|V|$. The number of alternatives will be denoted by $|C|$.

A *voting rule* $\mathcal{R}$ is a function that given an election $E = (C, V)$ returns a set of *election winners* $\mathcal{R}(E) \subseteq C$. Note that it is legal for the set of winners to contain more that one candidate. To simplify notation, we will sometimes write $\mathcal{R}(V)$ instead of $\mathcal{R}(E)$. We sometimes consider voting rules defined for a particular number of candidates (or even a particular set of candidates) only.

Below we define several prominent voting rules.

Scoring rules. For any sequence of non-negative real numbers $(\alpha_1, \ldots, \alpha_m)$, we can define a *scoring rule* $\boldsymbol{\mathcal{R}_{(\alpha_1,\ldots,\alpha_m)}}$ for elections with m candidates as follows: each candidate receives α_j points for each vote that ranks her in the jth position. The winner(s) are the candidate(s) with the highest score. Note that a scoring rule is defined for a fixed number of candidates. However, many standard voting rules can be defined via families of scoring rules. For example, *Plurality* is defined via the family of vectors $(1, 0, \ldots, 0)$, *veto* is defined via the family of vectors $(1, \ldots, 1, 0)$, and *Borda* is defined via the family of vectors $(m-1, m-2, \ldots, 0)$; k*-approval* is the scoring rule with $\alpha_i = 1$ for $i \leq k$, $\alpha_i = 0$ for $i > k$.

Bucklin and Simplified Bucklin. Given a positive integer k, $1 \leq k \leq |C|$, we say that a candidate c is a k*-majority winner* if more than $\frac{|V|}{2}$ voters rank c among the top k candidates. Let k' be the smallest positive integer such that there is at least one k'-majority winner for E. The *Bucklin score* of a candidate c is the number of voters that rank her in top k' positions. The *Bucklin winners* are the candidates with the highest Bucklin score; clearly, all of them are k'-majority winners. The *simplified Bucklin winners* are all k'-majority winners.

Single Transferable Vote (STV). In STV the winner is chosen as follows. We find a candidate with the lowest Plurality score (i.e., one that is ranked first the least number of times) and remove him from the votes. We repeat the process until a single candidate remains; this candidate is declared to be the winner. For STV the issue of handling ties—that is, the issue of the order in which candidates with lowest Plurality scores are deleted—is quite important, and is discussed in detail by Conitzer, Rognlie and Xia [6]. However, the results in our paper are independent of the tie-breaking rule.

Dodgson. Dodgson voting is based on measuring closeness to becoming a Condorcet winner. A *Condorcet winner* is a candidate that is preferred to any other candidate by a majority of voters. The Dodgson score of a candidate c is the smallest number of swaps of adjacent candidates that have to be performed on the votes to make c a Condorcet winner. The winner(s) are the candidate(s) with the lowest score.

Kemeny. Let $\succ$ and $\succ'$ be two preference orders over C. The number of *disagreements* between $\succ$ and $\succ'$, denoted $t(\succ, \succ')$, is the number of pairs of candidates c_i, c_j such that either $c_i \succ c_j$ and $c_j \succ' c_i$ or $c_j \succ c_i$ and $c_i \succ' c_j$. A candidate c_i is a Kemeny winner if there exists a preference order $\succ$ such that c_i is ranked first in $\succ$ and $\succ$ minimizes the sum $\sum_{i=1}^{n} t(\succ, \succ_i)$. We note that usually the Kemeny rule is defined to return the ranking $\succ$ that minimizes $\sum_{i=1}^{n} t(\succ, \succ_i)$, or a set of such rankings in case of a tie; however, here we focus on rules that return sets of winners and not rankings.

2.2 Distances. Let X be a set. A function $d\colon X \to \mathbb{R} \cup \{\infty\}$ is a *distance* (or, a *metric*) if for each $x, y, z \in X$ it satisfies the following four conditions: (a) $d(x, y) \geq 0$ (nonnegativity), (b) $d(x, y) = 0$ if and only if $x = y$ (identity of indiscernibles), (c) $d(x, y) = d(y, x)$ (symmetry), and (d) $d(x, y) \leq d(x, z) + d(z, y)$ (triangle inequality). If d satisfies all of the above conditions except the second one (identity of indiscernibles) then d is called a *pseudodistance*.

In the context of elections, it is useful to consider both distances over votes and over entire elections (that is, distances where the set X is the set of all linear orders over some given candidate set, and distances where X is the set of all possible elections); we remark that the former can be extended to the latter in a natural way (see the paragraph below and Section 4).

Two particularly useful distances over votes are the *discrete distance* and the *swap distance*.[3] Let C be a set of candidates and let u and v be two votes over C. The *discrete distance* $d_{\mathrm{discr}}(u, v)$ is defined to be 0 if $u = v$ and to be 1 otherwise. The *swap distance* $d_{\mathrm{swap}}(u, v)$ is the least number of swaps of adjacent candidates that transform vote u into vote v. Any distance d over votes can be extended in several ways to the distance over the profiles. For example, for any two elections, $E' = (C', V')$ and $E'' = (C'', V'')$, where $C' = C''$ and $V' = (v'_1, \dots, v'_n)$, $V'' = (v''_1, \dots, v''_n)$, we may define $\widehat{d}(E', E'') = \sum_{i=1}^{n} d(v'_i, v''_i)$ (and we set $\widehat{d}(E', E'') = \infty$ if the candidate sets are different or the profiles have different number of voters).

2.3 Consensus classes. Intuitively, we say that an election $E = (C, V)$ is a consensus if it has an undisputed winner. Formally, a *consensus class* is a pair $(\mathcal{E}, \mathcal{W})$ where $\mathcal{E}$ is a set of elections and $\mathcal{W}\colon \mathcal{E} \to C$ is a mapping which for each election $E \in \mathcal{E}$ assigns a unique alternative, which is called the *consensus alternative (winner).* We consider the following four natural classes that can be accepted by societies as consensus:

Strong unanimity. Denoted $\mathcal{S}$, this class contains elections $E = (C, V)$ where all voters report the same preference order. The consensus alternative is the candidate ranked first by all the voters.

Unanimity. Denoted $\mathcal{U}$, this class contains all elections $E = (C, V)$ where all voters rank some candidate c first. The consensus alternative is c.

Majority. Denoted $\mathcal{M}$, this class contains all elections $E = (C, V)$ where more than half of the voters rank some candidate c first. The consensus alternative is c.

Condorcet. Denoted $\mathcal{C}$, this class contains all elections $E = (C, V)$ with a Condorcet winner (defined above). The Condorcet winner is the consensus alternative.

2.4 Distance rationalizability. We now define the concept of distance rationalizability of a voting rule which has been used in the previous work.

Definition 2.1. *Let d be a distance over elections and let $\mathcal{K} = (\mathcal{E}, \mathcal{W})$ be a consensus class. We define the $(\mathcal{K}, d)$-*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 $c_i = \mathcal{W}(E')$. The set of $(\mathcal{K}, d)$-winners of an election $E = (C, V)$ consists of those candidates in C whose $(\mathcal{K}, d)$-score is the smallest.*

Definition 2.2. *A voting rule $\mathcal{R}$ is* distance-rationalizable *via a consensus class $\mathcal{K} = (\mathcal{E}, \mathcal{W})$ and a distance d over elections ($(\mathcal{K}, d)$-rationalizable), if for each election E, a candidate c is an $\mathcal{R}$-winner of E if and only if she is a $(\mathcal{K}, d)$-winner of E.*

Meskanen and Nurmi [19] show that many of the common voting rules are distance-rationalizable in a very natural way. For example, Kemeny is $(\mathcal{S}, \widehat{d}_{\mathrm{swap}})$-rationalizable, Borda is $(\mathcal{U}, \widehat{d}_{\mathrm{swap}})$-rationalizable, and Dodgson is $(\mathcal{C}, \widehat{d}_{\mathrm{swap}})$-rationalizable. It is quite remarkable that these three major voting rules are rationalized by the same distance. It is also easy to see that Plurality is $(\mathcal{U}, \widehat{d}_{\mathrm{discr}})$-rationalizable.

We remark that the notion of distance rationalizability introduced in Definition 2.2 allows for arbitrary consensus classes and distances; as we will see in the next section, this lack of constraints results in a definition that is too broad to be practically applicable.

[3] Swap distance is also called Kendall tau distance, Dodgson distance and bubble-sort distance.

3 Unrestricted Distance-Rationalizability: an Impasse

We say that a voting rule $\mathcal{R}$ over a set of candidates C satisfies *nonimposition* if for every $c \in C$ there exists an election with the set of candidates C in which c is the unique winner under $\mathcal{R}$. Clearly, nonimposition is a very weak condition that is satisfied by all common voting rules. Nevertheless, it turns out to be sufficient for unrestricted distance-rationalizability.

Theorem 3.1. *For any voting rule $\mathcal{R}$ over a set of candidates C that satisfies nonimposition, there is a consensus class $(\mathcal{K}, \mathcal{W})$ and a distance d such that $\mathcal{R}$ is $(\mathcal{K}, d)$-rationalizable.*

The consensus class used in the proof of Theorem 3.1 is somewhat artificial. However, the following theorem shows that a similar result holds for our natural consensus notions, too.

Definition 3.2. *Let $\mathcal{R}$ be a voting rule and let $(\mathcal{E}, \mathcal{W})$ be a consensus class. We say that $\mathcal{R}$ is* compatible *with $(\mathcal{E}, \mathcal{W})$, or $(\mathcal{E}, \mathcal{W})$-*compatible *if for each election $E = (C, V)$ in $\mathcal{E}$ it holds that $\mathcal{R}(E) = \{\mathcal{W}(E)\}$.*

Theorem 3.3. *For any consensus class $\mathcal{K} \in \{\mathcal{S}, \mathcal{U}, \mathcal{M}, \mathcal{C}\}$, a voting rule $\mathcal{R}$ is $(\mathcal{K}, d^{\mathcal{K}})$-rationalizable for some distance $d^{\mathcal{K}}$ if and only if $\mathcal{R}$ is $\mathcal{K}$-compatible.*

The proof of Theorem 3.3 is fairly simple: we construct the distance so that any given election is at distance 1 from all consensus elections with appropriate winners and at distance 2 from any other election.

Effectively, Theorem 3.3 shows that any interesting voting rule is distance-rationalizable with respect to the strong unanimity consensus. Thus, knowing that a rule is distance rationalizable—even with respect to a standard notion of consensus—provides no further insight into the properties of the rule. Moreover, the dichotomy between distance-rationalizable and non-distance-rationalizable rules becomes essentially meaningless.

However, the distances employed in the proof of Theorem 3.3 are very unnatural. In particular, the following proposition holds.

Proposition 3.4. *Let $\mathcal{R}$ be a voting rule that is $(\mathcal{K}, d^{\mathcal{K}})$-rationalizable via a consensus class $\mathcal{K} \in \{\mathcal{S}, \mathcal{U}, \mathcal{M}, \mathcal{C}\}$ and the distance $d^{\mathcal{K}}$ constructed in the proof of Theorem 3.3. If $d^{\mathcal{K}}$ is polynomial-time computable then the winner determination problem for $\mathcal{R}$ is in* P.

For example, this implies that, if P $\neq$ NP, the distance produced in the proof of Theorem 3.3 for the rationalization of Kemeny rule with respect to $\mathcal{S}$ is not polynomial-time computable. On the other hand, we know that Kemeny does have a very natural rationalization with respect to $\mathcal{S}$ via distance $\widehat{d}_{\mathrm{swap}}$. The requirement that the distance should be polynomial-time computable is essential for the distance rationalizability framework to be interesting, in addition to further, structural, restrictions on the distances that we will introduce in the next section.

4 Rationalizability via Votewise Distances

The results of the previous section make it clear that we need to restrict the set of distances that we consider. To identify an appropriate restriction, consider rationalizations of Borda and Plurality via distances $\widehat{d}_{\mathrm{swap}}$ and $\widehat{d}_{\mathrm{discr}}$, respectively (see the end of Section 2). To build either of these distances, we first defined a distance over votes and then extended it to a distance over elections (with the same candidate sets and equal-cardinality voter lists) via summing the distances between respective votes. This technique can be interpreted as taking the direct product of the metric spaces that correspond to individual votes, and defining the distance on the resulting space via the ℓ_1-norm. It turns out that distances obtained in this manner (possibly using norms other than ℓ_1), which we will call *votewise* distances, are very versatile and expressive. They are also attractive from the social choice point of view, as they exhibit continuous and monotone dependence on the voters' opinions.

In this section we will define votewise distances and attempt to answer the following three questions regarding voting rules that can be rationalized via them:

(a) What properties do such rules have?
(b) Which rules can be rationalized with respect to votewise distances?
(c) What is the complexity of winner determination for such rules?

Definition 4.1. *Given a vector space S over $\mathbb{R}$, a* norm *on S is a mapping N from S to $\mathbb{R}$ that satisfies the following properties:*

(i) positive scalability: $N(\alpha u) = |\alpha| N(u)$ for all $u \in S$ and all $\alpha \in \mathbb{R}$;
(ii) positive semidefiniteness: $N(u) \geq 0$ for all $u \in S$, and $N(u) = 0$ if and only if $u = 0$;
(iii) triangle inequality: $N(u + v) \leq N(u) + N(v)$ for all $u, v \in S$.

A well-known class of norms on $\mathbb{R}^n$ are the *p-norms* ℓ_p given by $\ell_p(x_1, \ldots, x_n) = (\sum_{i=1}^{n}(|x_i|^p))^{\frac{1}{p}}$, with the convention that $\ell_\infty(x_1, \ldots, x_n) = \max\{x_1, \ldots, x_n\}$. A norm N on $\mathbb{R}^n$ is said to be *symmetric* if it satisfies $N(x_1, \ldots, x_n) = N(x_{\sigma(1)}, \ldots, x_{\sigma(n)})$ for any permutation $\sigma : [1, n] \to [1, n]$; clearly, all p-norms are symmetric. We can now define our family of votewise distances.

Definition 4.2. *We say that a function d on pairs of preference profiles is* votewise *if the following conditions hold:*

1. $d(E, E') = +\infty$ if E and E' have a different set of candidates or a different number of voters.
2. For any set of candidates C, there exists a distance $d_C(\cdot, \cdot)$ defined on votes over C;
3. For any $n \in \mathbb{N}$, there exists a norm N_n on $\mathbb{R}^n$ such that for any two preference profiles $E = (C, U)$, $E' = (C, V)$ with $U = (u_1, \ldots, u_n)$ and $V = (v_1, \ldots, v_n)$ we have $d(E, E') = N_n(d_C(u_1, v_1), \ldots, d_C(u_n, v_n))$.

It is well known that any function defined in this manner is a metric. Thus, in what follows, we refer to votewise functions as *votewise distances*; we will also use the term "N-votewise distance" to refer to a votewise distance defined via a norm N, and denote a votewise distance that is based on a distance d over votes by $\widehat{d}$. Similarly, we will use the term *N-votewise rules* to refer to voting rules that can be distance-rationalized via one of our four consensus classes and an N-votewise distance.

An important special case of our framework is when N_n is the ℓ_1-norm, i.e., $N_n(x_1, \ldots, x_n) = x_1 + \cdots + x_n$; we will call any such distance an *additively votewise* distance, or, in line with the notation introduced above, an *ℓ_1-votewise* distance. So far, ℓ_1-votewise distances were the only votewise distances used in distance rationalizability constructions:[4] Meskanen and Nurmi [19] use them to distance-rationalize the Kemeny rule, Dodgson, Plurality and Borda, and we will show that the construction for Borda can be generalized to all scoring rules (also using an ℓ_1-votewise distance). However, N-votewise distances with $N \neq \ell_1$ are almost as easy to work with as ℓ_1-votewise distances and may be useful for rationalizing natural voting rules. In fact, later on we will see that simplified Bucklin is an ℓ_∞-votewise rule.

[4] However, see [23, Footnote 7].

4.1 Properties of Votewise Rules

In this section we consider three basic properties of voting rules. Specifically, given a consensus class $\mathcal{K}$ and a votewise distance $\widehat{d}$, we ask under which circumstances the voting rule that is distance-rationalizable via $(\mathcal{K}, \widehat{d})$ is anonymous, neutral, or consistent. To start, we recall the formal definitions of these properties.

Let $E = (C, V)$ be an election with $V = (v_1, \dots, v_n)$, and let σ and π be permutations of V and C, respectively. For any $C' \subseteq C$, set $\pi(C') = \{\pi(c) \mid c \in C'\}$. Let $\tilde{\pi}(v)$ be the vote obtained from v by replacing each occurrence of a candidate $c \in C$ by an occurrence of $\pi(c)$; we can extend this definition to preference profiles by setting $\tilde{\pi}(v_1, \dots, v_n) = (\tilde{\pi}(v_1), \dots, \tilde{\pi}(v_n))$.

Anonymity. A voting rule is *anonymous* if its result depends only on the number of voters reporting each preference order. Formally, a voting rule $\mathcal{R}$ is anonymous if for each election $E = (C, V)$ with $V = (v_1, \dots, v_n)$ and each permutation σ of V, the election $E' = (C, \sigma(V))$ satisfies $\mathcal{R}(E) = \mathcal{R}(E')$.

Neutrality. A voting rule is *neutral* if its result does not depend on the candidates' names. Formally, a voting rule $\mathcal{R}$ is neutral if for each election $E = (C, V)$, where $C = \{c_1, \dots, c_m\}$ and each permutation π of C, the election $E' = (C, \tilde{\pi}(V))$ satisfies $\mathcal{R}(E) = \pi^{-1}(\mathcal{R}(E'))$.

Consistency. A voting rule $\mathcal{R}$ is *consistent* if for any two elections $E_1 = (C, V_1)$ and $E_2 = (C, V_2)$ such that $\mathcal{R}(E_1) \cap \mathcal{R}(E_2) \neq \emptyset$, the election $E = (C, V_1 + V_2)$ (i.e., the election where the collections of voters from E_1 and E_2 are concatenated) satisfies $\mathcal{R}(E) = \mathcal{R}(E_1) \cap \mathcal{R}(E_2)$. This property was introduced by Young [21] and is also known as *reinforcement* [5].

For votewise distance-rationalizable rules, a symmetric norm produces an anonymous rule.

Proposition 4.3. *Suppose that a voting rule $\mathcal{R}$ is $(\mathcal{K}, \widehat{d})$-rationalizable, where $\mathcal{K} \in \{\mathcal{S}, \mathcal{U}, \mathcal{M}, \mathcal{C}\}$ and $\widehat{d}$ is an N-votewise distance, where N is a symmetric norm. Then $\mathcal{R}$ is anonymous.*

In contrast, neutrality is inherited from the underlying distance over votes.

Definition 4.4. *Let C be a set of candidates and let d be a distance on votes over C. We say that d is* neutral *if for each permutations π over C and any two votes u and v over C it holds that $d(u, v) = d(\tilde{\pi}(u), \tilde{\pi}(v))$. Further, we say that a votewise distance $\widehat{d}$ that corresponds to a distance d on votes is* neutral *if d is.*

Proposition 4.5. *Suppose that a voting rule $\mathcal{R}$ is $(\mathcal{K}, \widehat{d})$-rationalizable, where $\mathcal{K} \in \{\mathcal{S}, \mathcal{U}, \mathcal{M}, \mathcal{C}\}$ and $\widehat{d}$ is a neutral votewise distance. Then $\mathcal{R}$ is neutral.*

It is natural to ask if the converse of Proposition 4.5 is also true, i.e., if every neutral votewise rule can be rationalized via a neutral distance. Indeed, paper [6] provides a positive answer to a similar question in the context of representing voting rules as maximum likelihood estimators. However, the natural extension of the approach of [6] is not necessarily applicable in our setting. Nevertheless, all votewise distances that have so far arisen in the study of distance rationalizability of natural voting rules are neutral.

Our results for anonymity and neutrality are applicable to all consensus classes considered in this paper. In contrast, when discussing consistency, we need to limit ourselves to the unanimity consensus, and to ℓ_p-votewise rules.

Theorem 4.6. *Suppose that a voting rule $\mathcal{R}$ is $(\mathcal{U}, \widehat{d})$-rationalizable, where $\widehat{d}$ is an ℓ_p-votewise distance. Then $\mathcal{R}$ is consistent.*

While Theorem 4.6 may hold for some norms other than ℓ_p, we cannot hope to prove it for all votewise distances: fundamentally, consistency is a constraint on the relationship among N_s, N_t and N_{s+t} (i.e., the norms used for s voters, t voters, and $s+t$ voters), and our definition of votewise distances allows us to select norms N_n for different values of n independently of each other. Further, for our proof to work, the consensus class should be closed with respect to "splitting" and "merging" of the consensus profiles, and neither of the classes $\mathcal{S}$, $\mathcal{C}$, and $\mathcal{M}$ satisfies both of these conditions. Indeed, for $\mathcal{S}$ and $\mathcal{C}$ the conclusion of the theorem itself is not true: the counterexamples are provided by the Kemeny rule and the Dodgson rule, respectively (both are not consistent, yet rationalizable via $\widehat{d}_{\mathrm{swap}}$).

4.2 ℓ_p-Votewise Rules

Now that we know that ℓ_p-votewise rules have some desirable properties, let us see which voting rules are in fact ℓ_p-votewise distance rationalizable. We will generally focus on additively votewise rules, but we will look at ℓ_∞ as well. Naturally, we expect the answer to this question to strongly depend on the consensus notion used. Thus, let us consider unanimity, strong unanimity, majority, and Condorcet consensuses one by one.

We start with the unanimity consensus. By combining Propositions 4.3, 4.5 and Theorem 4.6, we conclude that any rule that is $(\mathcal{U}, \widehat{d})$-rationalizable, where $\widehat{d}$ is a neutral ℓ_1-votewise distance, is neutral, anonymous and consistent; it is not hard to check that the conclusion still holds if $\widehat{d}$ is a pseudodistance rather than a distance. In contrast, Young's famous characterization result [21] says that every voting rule that has all three of these properties is either a scoring rule or a composition of scoring rules (see [21] for an exact definition of composition of voting rules). It turns out that our framework allows us to refine Young's result by characterizing exactly the scoring rules themselves rather than their compositions. Moreover, we can actually "extract" the scoring rule from the corresponding distance, albeit not efficiently (see Section 4.3 for a discussion of the related complexity issues).

Theorem 4.7. *Let $\mathcal{R}$ be a voting rule. There exists a neutral ℓ_1-votewise pseudodistance $\widehat{d}$ such that $\mathcal{R}$ is $(\mathcal{U}, \widehat{d})$-rationalizable if and only if $\mathcal{R}$ can be defined via a family of scoring rules.*[5]

That is, the above theorem gives a complete characterization of voting rules rationalizable via neutral ℓ_1-votewise distances with respect to the unanimity consensus. However, the situation with respect to other consensus notions is more difficult.

Let us consider strong unanimity next. Intuitively, strong unanimity is quite challenging to work with as it provides very little flexibility. Meskanen and Nurmi [19] have shown that Kemeny is ℓ_1-votewise with respect to $\mathcal{S}$, but, at least at first, it seems that no other natural rule is. Interestingly, and very counterintuitively, Plurality is also ℓ_1-votewise with respect to strong unanimity.

Theorem 4.8. *There exists an ℓ_1-votewise distance $\widehat{d}$ such that Plurality rule is $(\mathcal{S}, \widehat{d})$-rationalizable.*

Naturally, this result suggests that, perhaps, all scoring rules are votewise distance-rationalizable with respect to $\mathcal{S}$. However, this turns out to be false.

Theorem 4.9. *There is no ℓ_1-votewise distance $\widehat{d}$ such that Borda rule is $(\mathcal{S}, \widehat{d})$-rationalizable.*

[5]Note that in this paper, following Young [21], we do not require $(\alpha_1, \ldots, \alpha_m)$ to be nondecreasing or integer. Indeed, the distance rationalizability framework does not impose any ordering over different positions in a vote, so it works equally well for a scoring rule with, e.g., $\alpha_1 < \alpha_2$.

Thus, the class of rules ℓ_1-votewise rationalizable with respect to $\mathcal{S}$ is rather enigmatic. On the one hand, it does contain Kemeny, a very complex rule, and Plurality, a very simple rule, yet it does not contain other natural scoring rules such as Borda. We believe that characterizing this class exactly is a very interesting research problem, particularly so since the rules in this class can be shown to be related to MLERIV rules of [7] and and [6] (we omit a description of this connection here due to space constraints).

Our understanding of rules that are votewise rationalizable with respect to $\mathcal{C}$ and $\mathcal{M}$ is even more limited. For example, Meskanen and Nurmi [19] have shown that Dodgson is ℓ_1-votewise rationalizable with respect to $\mathcal{C}$, and it is easy to see that no scoring rule is distance-rationalizable with respect to $\mathcal{C}$ because scoring rules are not Condorcet-consistent [20]. It is very interesting if, e.g., Young's rule is votewise with respect to $\mathcal{C}$ (however, see Section 5 for some comments). For the case of $\mathcal{M}$, we can show that simplified Bucklin is ℓ_∞-votewise with respect to $\mathcal{M}$; note that this result provides an argument for considering votewise distances that use a norm other that ℓ_1.

Theorem 4.10. *Simplified Bucklin is ℓ_∞-votewise with respect to consensus $\mathcal{M}$.*

The regular Bucklin rule is also rationalizable via a distance very similar to the one for simplified Bucklin but, nonetheless, not votewise. Finding further natural voting rules that are votewise rationalizable with respect to either $\mathcal{C}$ or $\mathcal{M}$ is an open question.

We conclude this section with a quick look at the STV rule. Conitzer, Rognlie, and Xia [6] have shown that STV is not MLERIV. It can be shown that this implies that STV is not distance-rationalizable via an ℓ_1-votewise distance with respect to $\mathcal{S}$. It turns out that this result can be extended to (almost) any votewise distance as well as two other consensus classes, namely, $\mathcal{U}$ and $\mathcal{C}$.

Definition 4.11 ([3])**.** *A norm N in $\mathbb{R}^n$ is* monotonic in the positive orthant*, or* $\mathbb{R}^n_+$-monotonic*, if for any two vectors $(x_1, \ldots, x_n), (y_1, \ldots, y_n) \in \mathbb{R}^n_+$ such that $x_i \leq y_i$ for all $i = 1, \ldots, n$ we have $N(x_1, \ldots, x_n) \leq N(y_1, \ldots, y_n)$.*

We say that a votewise distance is *monotonic* if the respective norm is monotonic in the positive orthant. We remark that monotonicity is a very weak constraint that is satisfied by any reasonable norm.

Theorem 4.12. *STV (together with any intermediate tie-breaking rule) is not distance-rationalizable with respect to either of $\mathcal{S}$, $\mathcal{U}$, or $\mathcal{C}$ and any neutral anonymous monotonic votewise distance.*

Note that Meskanen and Nurmi [19] show that STV can be distance-rationalized with respect to $\mathcal{U}$, but their distance is not votewise, and it is not immediately clear whether it is polynomial-time computable.

4.3 Winner Determination for Votewise Rules

Now that we have some understanding of the nature of votewise rules, we are ready to study the complexity of determining winners under them.[6] Clearly, to prove upper bounds on the complexity of this problem, we need to impose restrictions on the complexity of the distance itself. Thus, in what follows, we focus on distances that take values in $\mathbb{Z} \cup \{\infty\}$ and are polynomial-time computable; we will call a distance *normal* if it has both of these properties. We remark that restricting ourselves to distances with values in $\mathbb{Z} \cup \{\infty\}$ may prevent us from using ℓ_p-distances for values of p other than 1 and ∞. For example, taking

[6] We assume the reader is familiar with standard notions of complexity theory and fixed-parameter complexity. Due to space limits we cannot provide appropriate background in the paper.

the p-th root of an integer may yield a non-integer value. However, it is easy to see that for winner-determination, instead of using an ℓ_p-distance d, we can use function d^p, despite the fact that it is not a distance. This is so, because for winner-determination we only need to compare distances between elections.

The winner determination problem can be formally stated as follows.

Definition 4.13. *Let $\mathcal{R}$ be a voting rule. In the* $\mathcal{R}$-winner problem *we are given an election $E = (C, V)$ and a candidate $c \in C$ and we ask whether $c \in \mathcal{R}(E)$.*

This problem can be hard even for ℓ_1-votewise rules: for Dodgson and Kemeny it is known to be Θ_2^p-complete [11, 12]. On the positive side, for both of these rules the winner determination problem can be solved in polynomial time if the number of candidates is fixed. In fact, a stronger statement is true: the winner determination problem for both Dodgson and Kemeny is fixed parameter tractable with respect to the number of candidates.

We will now show that from the complexity perspective, Dodgson and Kemeny exhibit some of the worst possible behavior.

Theorem 4.14. *Suppose that a voting rule $\mathcal{R}$ is $(\mathcal{K}, d)$-rationalizable, where $\mathcal{K} \in \{\mathcal{S}, \mathcal{U}, \mathcal{M}, \mathcal{C}\}$, and d is a normal distance that satisfies $d((C_1, V_1), (C_2, V_2)) = +\infty$ whenever $C_1 \neq C_2$ or $|V_1| \neq |V_2|$. Then the $\mathcal{R}$-winner problem is in* P^{NP}. *Moreover, if, in addition, for any two elections $E_1 = (C, V_1)$, $E_2 = (C, V_2)$, the distance $d(E_1, E_2)$ is either $+\infty$ or at most polynomial in $|C| + |V_1| + |V_2|$, then the $\mathcal{R}$-winner problem is in Θ_2^p.*

Note that the distance used to rationalize Dodgson and Kemeny is polynomially bounded. On the other hand, there are natural distances that are not polynomially bounded; this includes distances that appear in our distance rationalizabiliy constructions for scoring rules with "large" coefficients.

If, in addition to being normal, the distance in question is an ℓ_1-votewise distance, the winner determination problem is fixed-parameter tractable with respect to the number of candidates.

Theorem 4.15. *Suppose that a voting rule $\mathcal{R}$ is $(\mathcal{K}, d)$-rationalizable, where $\mathcal{K} \in \{\mathcal{S}, \mathcal{U}, \mathcal{M}, \mathcal{C}\}$. and d is a normal ℓ_1-votewise distance. Then the $\mathcal{R}$-winner problem is FPT with respect to the number of candidates.*

In the previous section we have seen that neutral ℓ_1-votewise rules that use unanimity consensus correspond to families of scoring rules. Thus, one would expect their winner problems to be in P. Note, however, that in our setting we are given the distance, but not the scoring vector and computing the latter from the former might be hard. Nevertheless, it turns out that in this setting we can easily determine the winner if we are allowed to use polynomial-size *advice*.

Theorem 4.16. *Suppose that a voting rule $\mathcal{R}$ is distance-rationalizable via a normal neutral ℓ_1-votewise distance and unanimity consensus. Then $\mathcal{R}$-winner is in* P/poly.

P/poly is a complexity class that captures the power of polynomial computation "with advice." Karp–Lipton theorem [13] says that if there is an NP-hard problem in P/poly then the Polynomial Hierarchy collapses. Thus, for voting rules that are distance-rationalizable via a normal neutral ℓ_1-votewise distance and the consensus class $\mathcal{U}$ the winner determination problem is unlikely to be NP-hard. In contrast, this problem is hard for both Dodgson and Kemeny, even though they are both rationalizable via a normal neutral ℓ_1-votewise distance (and consensus classes $\mathcal{C}$ and $\mathcal{S}$, respectively). Thus, from computational perspective, the unanimity consensus appears to be easier to work with than the strong consensus and the Condorcet consensus. Indeed, both $\mathcal{S}$ and $\mathcal{C}$ impose "global" constraints on the closest consensus and $\mathcal{U}$ only imposes "local" ones.

5 Conclusions and Open Problems

In this paper we have presented general results regarding the recently introduced distance rationalizability framework. Our paper has two main contributions. First, we have shown that without any restrictions, essentially every reasonable voting rule is distance-rationalizable and further refinement of this framework is needed. Second, we have put forward a natural class of distances to consider—votewise distances—and proved that the rules which can be distance-rationalized using such distances have several desirable properties. We have identified a number of votewise rules, as well as showed that some rules are not votewise rationalizable with respect to standard consensus classes, and established complexity results for winner determination under votewise rules.

Are votewise distances the only natural distances that one should consider? Such distances are based on the assumption that, given an election $E = (C, V)$, if a voter changes her opinion in a minor way, then the resulting election $E' = (C, V')$ must not deviate from E too far. However some rules have discontinuous nature by definition, especially Young's rule which picks the winner of a largest Condorcet-consistent subelection. It is unlikely that such rules can be distance-rationalized via a votewise distance. Indeed, it can be shown that Young's rule and Maximin can be rationalized with respect to $\mathcal{C}$ via fairly intuitive distances that operate on profiles with different numbers of voters: in the case of Maximin we are, essentially, adding voters, and in the case of Young, we are deleting voters. (We omit the definitions of these rules and the construction due to space constraints). However, neither of these rules is known to be votewise rationalizable. Thus, it would be desirable to extend the class of "acceptable" distances to include some non-votewise distances; how to do this is an interesting research direction.

We mention that our work is closely related to a sequence of papers of Conitzer, Rognlie, Sandholm, and Xia [7, 6] on interpreting voting rules as maximum likelihood estimators. There are some very interesting connections (and differences) between the two approaches, but, unfortunately, due to space constraints, we cannot elaborate on them here.

Acknowledgements: This research was partially supported by NRF (Singapore NRF Fellowship RF-2009-08) and AGH University of Technology (Grant no. 11.11.120.865).

References

[1] K. Arrow. *Social Choice and Individual Values.* John Wiley and Sons, 1951 (revised editon, 1963).

[2] N. Baigent. Metric rationalisation of social choice functions according to principles of social choice. *Mathematical Social Sciences*, 13(1):59–65, 1987.

[3] F. Bauer, J. Stoer, and C. Witzgall. Absolute and monotonic norms. *Numerische Matematic*, 3:257–264, 1961.

[4] 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. Elsevier, 2002.

[5] P. Yu. Chebotarev and E. Shamis. Characterizations of scoring methods for preference aggregation. *Annals of Operations Research*, 80:299 332, 1998.

[6] 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.

[7] 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.

[8] V. Conitzer, T. Sandholm, and J. Lang. When are elections with few candidates hard to manipulate? *Journal of the ACM*, 54(3):Article 14, 2007.

[9] P. Faliszewski, E. Hemaspaandra, and L. Hemaspaandra. How hard is bribery in elections? *Journal of Artificial Intelligence Research*, 35:485–532, 2009.

[10] P. Faliszewski, E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Llull and Copeland voting computationally resist bribery and constructive control. *Journal of Artificial Intelligence Research*, 35:275–341, 2009.

[11] E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Exact analysis of Dodgson elections: Lewis Carroll's 1876 voting system is complete for parallel access to NP. *Journal of the ACM*, 44(6):806–825, 1997.

[12] E. Hemaspaandra, H. Spakowski, and J. Vogel. The complexity of Kemeny elections. *Theoretical Computer Science*, 349(3):382–391, 2005.

[13] R. Karp and R. Lipton. Some connections between nonuniform and uniform complexity classes. In *Proceedings of the 12th ACM Symposium on Theory of Computing*, pages 302–309. ACM Press, April 1980. An extended version has also appeared as: Turing machines that take advice, *L'Enseignement Mathématique*, 2nd series, 28:191–209, 1982.

[14] J. Kemeny. Mathematics without numbers. *Daedalus*, 88:575–591, 1959.

[15] C. Klamler. Borda and Condorcet: Some distance results. *Theory and Decision*, 59(2):97–109, 2005.

[16] C. Klamler. The Copeland rule and Condorcet's principle. *Economic Theory*, 25(3):745–749, 2005.

[17] K. Konczak and J. Lang. Voting procedures with incomplete preferences. In *Proceedins of the Multidisciplinary IJCAI-05 Worshop on Advances in Preference Handling*, pages 124–129, July/August 2005.

[18] R. Meir, A. Procaccia, J. Rosenschein, and A. Zohar. The complexity of strategic behavior in multi-winner elections. *Journal of Artificial Intelligence Research*, 33:149–178, 2008.

[19] T. Meskanen and H. Nurmi. Closeness counts in social choice. In M. Braham and F. Steffen, editors, *Power, Freedom, and Voting*. Springer-Verlag, 2008.

[20] H. Moulin. *Axioms of Cooperative Decision Making*. Cambridge University Press, 1991.

[21] H. Young. Social choice scoring functions. *SIAM Journal on Applied Mathematics*, 28(4):824–838, 1975.

[22] H. Young. Optimal voting rules. *Journal of Economic Perspectives*, 9(1):51–64, 1995.

[23] W. Zwicker. Consistency without neutrality in voting rules: When is a vote an average? *Mathematical and Computer Modelling*, 48(9–10):1357–1373, 2008.

Edith Elkind
Division of Mathematical Sciences
School of Physical and Mathematical Sciences
Nanyang Technological University
Singapore
Email: `eelkind@gmail.com`

Piotr Faliszewski
Department of Computer Science
AGH University of Science and Technology
Kraków, Poland
Email: `faliszew@agh.edu.pl`

Arkadii Slinko
Department of Mathematics
University of Auckland
Auckland, New Zealand
Email: `slinko@math.auckland.ac.nz`