On Swap-Distance Geometry of Voting Rules

Svetlana Obraztsova
National Technical University of Athens, Greece
Steklov Institute of Mathematics
St. Petersburg, Russia
svetlana.obraztsova@gmail.com

Edith Elkind
School of Physical and Mathematical Sciences
Nanyang Technological University, Singapore
eelkind@ntu.edu.sg

Piotr Faliszewski
AGH University of Science and Technology
Krakow, Poland
faliszew@agh.edu.pl

Arkadii Slinko
University of Auckland
Auckland, New Zealand
a.slinko@auckland.ac.nz

ABSTRACT

Axioms that govern our choice of voting rules are usually defined by imposing constraints on the rule's behavior under various transformations of the preference profile. In this paper we adopt a different approach, and view a voting rule as a (multi-)coloring of the election graph—the graph whose vertices are elections over a given set of candidates, and two vertices are adjacent if they can be obtained from each other by swapping adjacent candidates in one of the votes. Given this perspective, a voting rule $\mathcal{F}$ is characterized by the shapes of its "monochromatic components", i.e., sets of elections that have the same winner under $\mathcal{F}$. In particular, it would be natural to expect each monochromatic component to be convex, or, at the very least, connected. We formalize the notions of connectivity and (weak) convexity for monochromatic components, and say that a voting rule is connected/(weakly) convex if each of its monochromatic components is connected/(weakly) convex. We then investigate which of the classic voting rules have these properties. It turns out that while all voting rules that we consider are connected, convexity and even weak convexity are much more demanding properties. Our study of connectivity suggests a new notion of monotonicity, which may be of independent interest.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Multiagent systems*

Keywords

elections; swap distance; convexity; connectivity

1. INTRODUCTION

In many multi-agent systems and human collectives, agents have to work together and make group decisions, despite having different preferences over possible outcomes. Such decisions are traditionally made by means of voting, and therefore a thorough understanding of voting rules and their properties is important for the design and analysis of multiagent systems. Many such studies have been undertaken in the past, with the focus on either axiomatic or computational properties of voting rules (see, e.g., [19, 1] for overviews of classic social choice topics and [7, 16] for a discussion of computational aspects of voting). The former strand of research identifies desirable properties of voting rules; usually such properties specify how a voting rule should behave under various transformations of the preference profile (such us, e.g., merging two preference profiles that are identical or that have the same winner, shifting the winner upwards in some votes, etc.). The latter strand focuses, e.g., on efficient winner determination or on modifications of the profile that lead to a desired change of the election outcome (this description encompasses several forms of malicious behavior, such as manipulation [2, 9], control [3, 17] or bribery [15], as well as more benign actions such as campaign management [12, 24]).

In this paper, in contrast to the previous work, we adopt a bird's-eye view of elections and voting rules. Our approach is geometric in nature and is motivated by recent work on distance-based interpretation of voting rules [18, 14, 13]. Specifically, we view the set of all elections (with a fixed set of candidates C and a fixed number of voters) as the vertex set of a giant graph where we connect two vertices by an edge if they can be obtained from each other by swapping two adjacent candidates in some vote. As a result, the shortest path distance in our graph is the so-called *swap distance*. Now, each voting rule $\mathcal{F}$ can be viewed as a multicoloring of this graph with elements of C: the set of colors assigned to a vertex E is the set of winners in the election E under the rule $\mathcal{F}$ (we allow rules with multiple winners, i.e., voting correspondences). One can then ask which properties we would expect this coloring to have for a voting rule to be considered "reasonable". For instance, we would expect each "blob" of color to have the same size; this property is implied by (but not equivalent to) voting rule neutrality. However, perhaps a more immediate requirement is for each "blob" to consist of elections that are grouped together in the graph. At a minimum, we would want each "blob" to be connected, that is, to induce a connected subgraph of our graph. More ambitiously, we might ask whether each "blob" is convex, i.e., whether it is the case that for every two vertices that have the same color there is a shortest path between them in the graph that only goes through vertices of that color.

Appears in: *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013), Ito, Jonker, Gini, and Shehory (eds.), May, 6–10, 2013, Saint Paul, Minnesota, USA.*

In this paper, we study connectivity and convexity (in the above-defined sense) of a variety of voting rules. We start by investigating the weaker of these two properties, namely, connectivity (Section 3). We first give a simple argument which establishes that every weakly monotone rule is connected. We then identify a more relaxed monotonicity criterion, which we call *feeble monotonicity*, which turns out to be sufficient for connectivity, and may be of independent interest. We illustrate the power of this criterion by showing that three well-known voting rules that are not weakly monotone—Single Transferable Vote (STV), Coombs, and Dodgson—are feebly monotone. In contrast, Baldwin and Nanson, two sequential elimination rules that are based on the Borda score, are not even feebly monotone. Nevertheless, even for these rules we manage to show connectivity, by using a somewhat more subtle argument.

We then move on to the study of convexity (Section 4). We establish that a natural definition of convexity, which parallels the definition of connectivity used in Section 3, is so stringent that even the most basic voting rules do not satisfy it. We then consider a somewhat more permissive definition of convexity and demonstrate that it is satisfied by the Plurality rule. But even this relaxed definition turns out to be too demanding: we show that it is failed by Borda (and a large class of other scoring rules), Maximin and Copeland. On the other hand, we show that several natural consensus classes (classes of elections with a single, undisputed winner; e.g., elections where each votes ranks the same candidate first) do satisfy convexity. We also provide another example of a convex set of elections, namely, elections that are single-peaked with respect to a given axis (see Section 4.3 for a definition of single-peaked elections). Further, we argue that for a large class of voting rules, which includes all common rules, convexity implies feeble monotonicity.

2. PRELIMINARIES

We write $[n]$ to denote the set $\{1 \dots, n\}$. An *election* is specified by a set of *candidates* $C = \{c_1, \dots, c_m\}$ and a *preference profile* $\mathcal{R} = (R_1, \dots, R_n)$, where each R_i, $i \in [n]$, is a linear order over C. We will sometimes write $\succ_i$ in place of R_i. We will refer to the entries of $\mathcal{R}$ as *votes*; i.e., R_i is the vote of voter i in election $E = (C, \mathcal{R})$. We will say that voter i ranks a candidate $a \in C$ above $b \in C$ (or, prefers a to b) if $a \succ_i b$. We denote the top candidate in vote R_i by $top(R_i)$, i.e., we write $a = top(R_i)$ if $a \succ_i b$ for all $b \in C \setminus \{a\}$. Given a set of candidates X and two candidates $a, b \notin X$, we write $a \succ X \succ b$ to denote a vote where a is ranked above all candidates in X, b is ranked below all candidates in X, and the candidates in X are ranked in lexicographic order; we use $a \succ \overleftarrow{X} \succ b$ to denote the same vote with the order of the candidates in X reversed.

A *voting rule*[1] is a mapping $\mathcal{F}$ that given an election $E = (C, \mathcal{R})$ outputs a non-empty set of candidates $W = \mathcal{F}(E) \subseteq C$; the candidates in W are called the *winners* of the election E under the voting rule $\mathcal{F}$. If $\mathcal{F}(E) = \{c\}$ for some $c \in C$ then candidate c is called the *unique winner* of the election E. If $|\mathcal{F}(E)| > 1$ then the elements of $\mathcal{F}(E)$ are called the *co-winners* of E.

Given a set of candidates C, $|C| = m$, and a positive integer n, we construct a graph $\mathcal{G}_{n,m} = (\mathcal{V}, \mathcal{E})$ whose vertices are all n-voter elections over C and where we have an edge between two elections $(C, \mathcal{R})$ and $(C, \mathcal{R}')$ if $\mathcal{R}'$ can be obtained from $\mathcal{R}$ by swapping two adjacent candidates in one of the votes. Note that $|\mathcal{V}| = (m!)^n$ and the degree of each vertex is $(m-1)n$. In what follows, we omit the indices n and m and write $\mathcal{G}$ in place of $\mathcal{G}_{n,m}$ where this does not lead to confusion.

One can think of a voting rule $\mathcal{F}$ as a (multi-)coloring of this graph: the set of colors is C and each vertex E is assigned a set of colors $\mathcal{F}(E)$. Given a candidate $a \in C$, we denote by $col_a(\mathcal{G}, \mathcal{F})$ the set of all vertices that are colored with a under $\mathcal{F}$: $col_a(\mathcal{G}, \mathcal{F}) = \{E \mid a \in \mathcal{F}(E)\}$. We will refer to the set $col_a(\mathcal{G}, \mathcal{F})$, $a \in C$, as a *monochromatic component* of $\mathcal{G}$ with respect to $\mathcal{F}$. We will be interested in how the choice of a voting rule affects the "shapes" of the monochromatic components. To formalize the notion of shape, we introduce the following definitions.

Definition 2.1. *A subset $\mathcal{V}'$ of vertices of the graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is said to be* connected *if for every $E, E' \in \mathcal{V}'$ there exists a path from E to E' that lies entirely within $\mathcal{V}'$. A voting rule $\mathcal{F}$ is said to be* connected *if for each $a \in C$ the set $col_a(\mathcal{G}, \mathcal{F})$ is connected.*

Definition 2.2. *A subset $\mathcal{V}'$ of vertices of the graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is said to be* convex *if for every $E, E' \in \mathcal{V}'$ there exists a path from E to E' that lies entirely within $\mathcal{V}'$ and has length d, where d is the length of a shortest path between E and E' in $\mathcal{G}$. A voting rule $\mathcal{F}$ is said to be* convex *if for each $a \in C$ the set $col_a(\mathcal{G}, \mathcal{F})$ is convex.*

We remark that given two elections, E and E', there may be many shortest paths between them. However, each such shortest path can be obtained by executing (in some order) exactly all the swaps of the form "swap candidates c_i and c_j in vote k", where the k-th vote in E ranks c_i ahead of c_j and the k-th vote in E' ranks c_j ahead of c_i (see, e.g., [11]).

While the notion of convexity defined above is satisfied by several *consensus classes* (see Section 4), it turns out to be too stringent for common voting rules: we consider several well-studied rules, and show that none of them is convex in the above sense. Therefore, we will now introduce a more relaxed notion, which we call *weak convexity*. Intuitively, a voting rule is weakly convex if the convexity condition is satisfied for pairs of elections with the same unique winner.

Definition 2.3. *A voting rule $\mathcal{F}$ is said to be* weakly convex *if for every $a \in C$ and every pair of elections E_1, E_2 such that $\mathcal{F}(E_1) = \mathcal{F}(E_2) = \{a\}$ there exists a path between E_1 and E_2 that lies entirely in the set $col_a(\mathcal{G}, \mathcal{F})$ and has length d, where d is the length of a shortest path between E and E' in $\mathcal{G}$.*

We mention that our definition of convexity, and—in general—the nature of our geometric approach, is very different from that used by Saari [22, 23]. The two main differences, as compared to the notion of convexity in [22], are that (a) Saari uses a very different space of elections (in particular, in his space elections with different number of voters are naturally allowed whereas this is not the case in our model), and (b) Saari is interested in the process of merging two elections (and, in effect, in changing the number of voters) whereas we transform the votes, without changing the number of voters. Further, two elections that are close

[1] Technically, the appropriate term would be a *voting correspondence*, but we use "rule" for brevity.

in terms of the swap distance may be quite distant in the election space used by Saari.

Voting rules We now describe the voting rules considered in this paper (see [5] for more details on voting rules). Unless explicitly specified otherwise, the winners are the candidates with the highest score. Given a profile $\mathcal{R}$ and a candidate c, we write $s(\mathcal{R}, c)$ to denote the score of c in $\mathcal{R}$ under the rule being defined.

Scoring rules Let m be the number of candidates. Given a vector $\alpha = (\alpha_1, \ldots, \alpha_m) \in \mathbb{R}^m$ with $\alpha_1 \geq \cdots \geq \alpha_m$, we define a *scoring rule* $\mathcal{F}_\alpha$ as follows. Each voter grants α_i points to the candidate she ranks in the i-th position; the score of a candidate is the sum of the scores he receives from all voters. In this work, we assume without loss of generality that the coordinates of α are nonnegative integers given in binary. Some well-known voting rules can be represented as families of scoring rules. For example, *Plurality* is given by the family of scoring rules with vectors of the form $(1, 0, \ldots, 0)$, *Veto* corresponds to a family of scoring rules with vectors $(1, \ldots, 1, 0)$, and *Borda* corresponds to a family of scoring rules with vectors $(m-1, \ldots, 1, 0)$.

Copeland A candidate a is said to win a *pairwise election* against candidate b if more than half of the voters prefer a to b. If exactly half of the voters prefer a to b then a is said to *tie* his pairwise election against b. Under the Copeland rule, each candidate gets 1 point for each pairwise election he wins and 0.5 points for each pairwise election he ties.

Maximin The *Maximin score* of a candidate $c \in C$ is given by the number of votes c gets in his worst pairwise election, i.e., $\min_{d \in C \setminus \{c\}} |\{i \mid c \succ_i d\}|$.

Dodgson A candidate $a \in C$ is said to be the *Condorcet winner* of an election $E = (C, \mathcal{R})$ if he beats every other candidate in their pairwise election. Given an election E and a candidate $a \in C$, the *Dodgson score* of a in E is the distance in $\mathcal{G}$ from E to the nearest election that has a as its Condorcet winner. The *Dodgson winner(s)* are the candidates with the smallest Dodgson score.

Successive elimination rules For any voting rule $\mathcal{F}$ that is defined by assigning scores to candidates (so that the candidate with the highest score wins), we can define a *successive elimination variant* of this rule, denoted by SE-$\mathcal{F}$, as follows. Under SE-$\mathcal{F}$, the election proceeds in rounds. During each round, the candidate with the lowest $\mathcal{F}$-score is eliminated, and the candidates' $\mathcal{F}$-scores are recomputed. The winner is the candidate that survives until the end. If several candidates have the lowest $\mathcal{F}$-score, we assume that the candidate to be eliminated is chosen according to the lexicographic order over the candidates: if S is the set of candidates that have the lowest $\mathcal{F}$ score in some round, we eliminate the candidate c_j such that $j \geq i$ for all $c_i \in S$. (Other tie-breaking rules are also popular; see, e.g., [8] for parallel-universes tie-breaking.) Several classic voting rules can be described as sequential elimination rules: *STV* is SE-Plurality, *Coombs* is SE-Veto, and *Baldwin* is SE-Borda. *Nanson* is another rule based on the idea of successive elimination. Just as Baldwin, it works with Borda scores, but instead of eliminating a single candidate with the lowest Borda score, it eliminates all candidates whose Borda scores are strictly below the average. We mention that successive elimination rule have recently received some attention in the context of their algorithmic properties [20, 10].

3. CONNECTIVITY

We start by presenting a simple, yet powerful, argument showing that many voting rules are connected. This argument makes use of a classic notion of monotonicity of voting rules, known as *weak monotonicity.*[2]

Definition 3.1. *A voting rule $\mathcal{F}$ is said to be* weakly monotone *if for every election $E = (C, \mathcal{R})$ and every candidate $a \in \mathcal{F}(E)$ it holds that every election $E' = (C, \mathcal{R}')$ obtained from E by picking a vote R_i with $top(R_i) \neq a$ and swapping a with the candidate ranked right above her in R_i satisfies $a \in \mathcal{F}(E')$.*

It is not hard to verify that all scoring rules, Copeland, and Maximin are weakly monotone. In contrast, neither Dodgson nor any of the successive elimination rules (including Nanson) has this property (for the case of Dodgson, see [6] for a survey of its deficiencies; for the case of Baldwin and Nanson—and elimination rules in general—see, e.g., [21] and the references therein, in particular, the work of Smith [25]).

We will say that a voting rule $\mathcal{F}$ is *unanimity-consistent* if in every election E where some candidate a is ranked first by all voters, we have $a \in \mathcal{F}(E)$. Clearly, all voting rules defined in Section 2 are unanimity-consistent.

Proposition 3.2. *Every weakly monotone and unanimity-consistent voting rule is connected.*

Proof. Fix a weakly monotone and unanimity-consistent voting rule $\mathcal{F}$, a candidate a and two elections $E^1 = (C, \mathcal{R})$, $E^2 = (C, \mathcal{R}^2)$ in $col_a(\mathcal{G}, \mathcal{F})$. We will argue that there is a path from E^1 to E^2 that is contained in $col_a(\mathcal{G}, \mathcal{F})$.

Let $\widehat{E^1}$ (respectively, $\widehat{E^2}$) be the election obtained by moving a to the top of each vote in $\mathcal{R}^1$ (respectively, $\mathcal{R}^2$), while leaving the relative ordering of all other candidates unchanged. Election $\widehat{E^1}$ can be reached from E^1 by repeatedly swapping a with a candidate ranked right above her in some vote. Thus, by weak monotonicity of $\mathcal{F}$ there is a path $\mathcal{P}^1$ from E^1 to $\widehat{E^1}$ in $col_a(\mathcal{G}, \mathcal{F})$. Similarly, there is a path $\mathcal{P}^2$ from E^2 to $\widehat{E^2}$ in $col_a(\mathcal{G}, \mathcal{F})$. Now, both in $\widehat{E^1}$ and in $\widehat{E^2}$ candidate a is ranked first in all votes. Thus, we can transform $\widehat{E^1}$ into $\widehat{E^2}$ by repeatedly swapping adjacent candidates without touching a. Let $\mathcal{P}^{12}$ be the path in $\mathcal{G}$ that corresponds to this transformation. Since $\mathcal{F}$ is unanimity-consistent, this path is contained in $col_a(\mathcal{G}, \mathcal{F})$. Thus, we can obtain a path from E^1 to E^2 in $col_a(\mathcal{G}, \mathcal{F})$ by gluing together $\mathcal{P}^1$, $\mathcal{P}^{12}$, and (the reverse of) $\mathcal{P}^2$. □

Corollary 3.3. *Copeland, Maximin, and all scoring rules are connected.*

Note that the proof of Proposition 3.2 does not use the full power of weak monotonicity: for the proof to go through, it suffices to be able to move the current winner a up in *some* vote without making it lose, whereas weak monotonicity criterion ensures that we can move a up in *every* vote where she is not ranked first. This motivates the following definition, which we believe may be of independent interest.

[2]The qualifier "weak" is used to distinguish it from another classic notion of monotonicity, *strong* monotonicity (see, e.g., [19]). We do not define strong monotonicity here as it is not relevant to our work.

Definition 3.4. *A voting rule* $\mathcal{F}$ *is said to be* feebly monotone *if for every election* $E = (C, \mathcal{R})$ *and every candidate* $a \in \mathcal{F}(E)$ *it holds that either (1)* $top(R_i) = a$ *for all* $i \in [n]$*, or (2) there exists a vote* R_i *with* $top(R_i) \neq a$ *such that the election* $E' = (C, \mathcal{R}')$ *obtained from* $\mathcal{R}$ *by swapping* a *with the candidate ranked right above her in* R_i *satisfies* $a \in \mathcal{F}(E')$.

As argued above, the proof of Proposition 3.2 actually shows the following stronger claim.

Corollary 3.5. *Every feebly monotone and unanimity-consistent voting rule is connected.*

Of course, the reader may wonder if the class of feebly monotone voting rules is different from the class of weakly monotone voting rules. We will now show that this is indeed the case, by proving that STV, Coombs, and Dodgson are feebly monotone (as we have mentioned, all of these rules are known not to be weakly monotone).

Theorem 3.6. *The STV rule is feebly monotone.*

Proof. Consider an election $E = (C, \mathcal{R})$ and let a be the STV winner of E (note that under our definition of STV, every election has a unique STV winner). The case where a is the unanimous winner of E is trivial, so we assume that there is a vote in $\mathcal{R}$ where a is not ranked first.

We rename the candidates so that $C = \{c_1, \ldots, c_m\}$, where for each j, $j = 1, \ldots, m-1$, c_j is the candidate eliminated in the j-th round, and where c_m is a. For each vote R_i in $\mathcal{R}$ where a is not ranked on top, we define $above_i(a)$ to be the candidate that R_i ranks right above a. Further, for each such R_i we define ℓ_i to be the first elimination round such that before round ℓ_i starts, all candidates ranked in R_i ahead of the candidate $above_i(a)$ are eliminated. Observe that until round ℓ_i, voter i does not contribute to the Plurality scores of either $above_i(a)$ or a. Thus, if $above_i(a)$ is eliminated before round ℓ_i, we can safely swap $above_i(a)$ and a in R_i: the Plurality scores of all candidates in all rounds would not be affected, and therefore a would win in the modified election as well.

It remains to consider the case where for each vote R_i with $top(R_i) \neq a$ the candidate ranked right above a in R_i is not eliminated before round ℓ_i. Let $\ell = \max\{\ell_i \mid a \neq top(R_i)\}$, let R_k be some vote with $\ell_k = \ell$, let $c = above_k(a)$, and let r be an integer such that $c = c_r$. We will now argue that if we swap a and c in R_k then a will still be the STV winner.

Let E' be the election obtained by swapping a and c in R_k. If STV operates in the same way on E and E', i.e., if STV eliminates candidates in the same order in both elections then, naturally, a is the unique winner of E'. Now, suppose there is a round where STV eliminates different candidates in E and E', and let ℓ' be the first such round. The reader can verify that we have $\ell' \geq \ell$ and, moreover, it has to be the case that in round ℓ' STV eliminates c from E' (whereas in E candidate c is eliminated in round $r > \ell'$).

We claim that at the start of round ℓ' all voters in E who rank c first (and, hence, all voters in E' who rank c first) rank a right after c. That is, after round ℓ' in E' all votes for c transfer to a. Indeed, suppose there exists some voter i in E' who at the beginning of round ℓ' ranks c first but does not rank a second. Let $b = above_i(a)$, $b \neq c$, be the candidate ranked right above a in the original ordering R_i (i.e., before we started eliminating candidates). Note that by our assumption candidate b is not eliminated until round ℓ_i. Since at the beginning of round ℓ' candidate c still appears above a in vote i (both in E and in E'), it has to be the case that $\ell_i > \ell' \geq \ell$, a contradiction with our choice of ℓ.

Thus, after round ℓ', all c's votes in E' transfer to a. Observe that after round ℓ' the Plurality score of each candidate $c' \in C \setminus \{c, a\}$ in E' is the same as his Plurality score in E after round $\ell' - 1$, whereas the Plurality score of a in E' after round ℓ' may be higher than her score in E after round ℓ'. Note also that in E candidate a is not eliminated until the very end (and recall that the tie-breaking rule is lexicographic). Therefore, the behavior of STV on E' in rounds $\ell' + 1, \ldots, r$ is identical to its behavior on E in rounds $\ell', \ldots, r - 1$, i.e., it eliminates the same candidates in the same order (note that all votes that transfer to c in E transfer to a in E', and thus do not affect the relative scores of other candidates). Then, in round r STV eliminates $c = c_r$ from E (as argued above, all these votes transfer to a), and from round $r + 1$ onwards STV operates identically on both E and E'. In particular, this means that a is the STV winner of E', which is what we wanted to prove. □

A similar argument, omitted due to space constraints, shows that the Coombs rule is feebly monotone as well.

Theorem 3.7. *The Coombs rule is feebly monotone.*

For the Dodgson rule, the argument is somewhat different.

Theorem 3.8. *The Dodgson rule is feebly monotone.*

Proof. Consider an election $E = (C, \mathcal{R})$ and let a be a Dodgson winner of E. Suppose first that a is the Condorcet winner of E. Then a remains the Condorcet winner (and hence the unique Dodgson winner) after it is shifted upwards in any of the votes, so we are done.

Now, suppose that a's Dodgson score is $s > 0$ (and hence the Dodgson score of any other candidate is at least s), and consider a sequence of swaps of length s that makes a the Condorcet winner. It can be assumed without loss of generality that each of these swaps involves a and moves her upwards in the vote. Suppose that the first of these swaps takes place in the i-th vote, i.e., it swaps a with candidate c ranked right above her in R_i. Consider the election E' obtained from E by swapping a and c in R_i. In E', the Dodgson score of a is $s - 1$: it is at most $s - 1$, because we can perform the rest of the swaps in the original sequence that makes a the Condorcet winner, and it is at least $s - 1$, because otherwise the Dodgson score of a in E would be less than s. On the other hand, the Dodgson score of any other candidate in E' is at least $s - 1$: any sequence of swaps that makes a candidate $b \neq a$ the Condorcet winner in E' combined, if necessary, with swapping b and a in vote i, would make b the Condorcet winner in E. Thus, a is a Dodgson winner in E'. □

While the set of feebly monotone voting rules is considerably broader than the set of weakly monotone voting rules, there exist voting rules that are not even feebly monotone. In particular,this is the case for the Nanson rule and the Baldwin rule.

Theorem 3.9. *The Baldwin rule is not feebly monotone.*

PROOF. We construct an election over the set of candidates $C = \{x_0, \dots, x_4, a, b\}$ as follows. Let $\mathcal{R}^0$ be a 5-voter preference profile over $\{x_0, \dots, x_4\}$ where voter i, $i = 1, \dots, 5$, ranks the candidates as:

$$x_{i \bmod 5} \succ x_{i+1 \bmod 5} \succ \dots \succ x_{i-1 \bmod 5}.$$

Let $\mathcal{R}^1$ be a 5-voter profile over C obtained by inserting a and b into, respectively, the top position and the second-to-last position in each vote in $\mathcal{R}^0$; for instance, the first voter in $\mathcal{R}^1$ ranks the candidates as $a \succ x_1 \succ x_2 \succ x_3 \succ x_4 \succ b \succ x_0$. Similarly, let $\mathcal{R}^2$ be a 5-voter profile over C obtained by inserting a and b into, respectively, the second-to-last and the second position in each vote in $\mathcal{R}^0$; for instance, the first voter in $\mathcal{R}^2$ ranks the candidates as $x_1 \succ b \succ x_2 \succ x_3 \succ x_4 \succ a \succ x_0$. Finally, let $\mathcal{R}^3$ be a 5-voter profile over C obtained by inserting a and b into, respectively, the first and the last position in each vote in $\mathcal{R}^0$; for instance, the first voter in $\mathcal{R}^3$ ranks the candidates as $a \succ x_1 \succ x_2 \succ x_3 \succ x_4 \succ x_0 \succ b$.

Now, consider the election obtained by taking 4 copies of $\mathcal{R}^1$, 5 copies of $\mathcal{R}^2$, and 1 copy of $\mathcal{R}^3$. Assume that the tie-breaking rule favors b over all other candidates.

In this election the Borda score of a is 175, the Borda score of b is 145, and the Borda score of each x_i, $i = 1, \dots, 4$, is 146. Thus, in the first round b is eliminated. In the next five rounds, candidates $x_1, \dots, x_0$ are eliminated, and finally a wins.

Now, suppose that we move a one position up in some vote in $\mathcal{R}^2$. This lowers the score of some x_i, $i = 0, \dots, 4$, to 145. Hence, because of the tie-breaking rule, in the modified election x_i is the first candidate to be eliminated. Then all other candidates in $C \setminus \{a, b\}$ are eliminated one by one, starting with $x_{i-1 \bmod 5}$ and ending with $x_{i+1 \bmod 5}$. In the last round the election is a tie between a and b, so b wins because of the tie-breaking rule. This is the case no matter in which vote a was moved upwards. Hence, the theorem is proved. □

THEOREM 3.10. *The Nanson rule is not feebly monotone.*

This theorem is witnessed by the following 12-voter 7-candidate profile (preference orders are formed by the matrix columns):

$$\begin{pmatrix}
a_7 & a_4 & a_3 & a_2 & a_1 & a_6 & a_7 & a_4 & a_5 & a_1 & a_2 & a_6 \\
a_5 & a_3 & a_2 & a_7 & a_6 & a_5 & a_3 & a_5 & a_1 & a_7 & a_6 & a_3 \\
a_4 & a_2 & a_7 & a_1 & a_5 & a_4 & a_4 & a_1 & a_7 & a_2 & a_3 & a_4 \\
a_3 & a_7 & a_1 & a_6 & a_4 & a_3 & a_5 & a_7 & a_2 & a_6 & a_4 & a_5 \\
a_2 & a_1 & a_6 & a_5 & a_3 & a_2 & a_1 & a_2 & a_6 & a_3 & a_5 & a_1 \\
a_6 & a_6 & a_5 & a_4 & a_2 & a_1 & a_6 & a_6 & a_3 & a_4 & a_1 & a_2 \\
a_1 & a_5 & a_4 & a_3 & a_7 & a_7 & a_2 & a_3 & a_4 & a_5 & a_7 & a_7
\end{pmatrix}$$

Borda scores of all candidates in this profile equal 36 and thus Nanson rule declares all candidates to be the winners of the election. In particular, a_6 is among the winners. However, in every vote where a_6 is not ranked first, either a_1 or a_2 is ranked right above a_6. A straightforward calculation shows that if we move a_6 up, a_6 ceases to be a winner.

THEOREM 3.11. *Both the Baldwin rule and the Nanson rule are connected.*

We omit the full proof of Theorem 3.11 due to space constraints. Briefly, for both rules, we follow the same line of reasoning as in the proof of Proposition 3.2. Specifically, given two profiles with the same winner a, in each profile we identify the candidate(s) that are first to be eliminated, and shift them to the bottom in all votes. We do so repeatedly until we reach a pair of profiles such that a is ranked first by all voters in both profiles. We then use the fact that both Baldwin and Nanson are unanimity-consistent to complete the proof.

4. CONVEXITY

Clearly, convexity is a much more demanding property than connectivity, but also a much more desirable one. Consider a society where the voters' preferences evolve in some direction over time (we use the terms "evolve" and "direction" informally here). For example, suppose that initially the voters prefer less expensive energy sources that result in higher pollution, but over time—as some sort of awareness campaign is conducted—their preferences move towards more expensive, but clean energy sources. The fact that preferences "evolve in some direction" could be modeled by the fact that these preferences move over some shortest path between two elections. A voting rule's convexity would ensure that if the society switches from some aggregated top-preference a to b, then over the course of further evolution it never switches back to a (unless the direction of the evolution changes).

On a more technical ground, for unanimity-consistent voting rules convexity can be viewed as a strengthening of feeble monotonicity (and, if one were to change the definition to speak of *all* shortest paths rather than of *at least one* shortest path, then it would have been a strengthening of weak monotonicity). Indeed, consider a unanimity-consistent voting rule $\mathcal{F}$ and an election E where no candidate is ranked first by all voters, and let a be a candidate in $\mathcal{F}(E)$. Let E' be the election obtained from E by shifting a to the top position in each vote. By unanimity consistency we have $a \in \mathcal{F}(E')$. Now, if $\mathcal{F}$ is convex, then there is some shortest path from E to E' such that a is a co-winner of every election on this path. Further, by properties of the swap distance, the first election after E on this path is identical to E except that in some vote a is swapped with the candidate originally preceding her. Since our choice of E was, in essence, arbitrary, this reasoning shows that if $\mathcal{F}$ is convex then it is feebly monotone (and if the definition of convexity spoke of all shortest paths, then the same reasoning would imply weak monotonicity of $\mathcal{F}$).

Convexity is a very intriguing property. Our first set of results shows that it is extremely difficult for a voting rule to satisfy convexity, and even the Plurality rule fails to be convex in the sense of Definition 2.2 (but it is weakly convex in the sense of Definition 2.3). On the other hand, our second set of results shows that convexity is satisfied by many natural consensus classes (i.e., mappings that are very similar to voting rules, but for each election either return a unique winner or indicate that there is no winner). This difference between voting rules and consensus classes is quite striking, and suggests that perhaps convexity should be viewed as a normative requirement that should be imposed on consensus classes, but not necessarily on voting rules built around them (the idea of building a voting rule around a consensus class is the basis of the work on distance rationalization of voting rules [18, 14, 13]).

We conclude the discussion of convexity by showing that the set of elections that are single-peaked with respect to a given axis is convex. This result can be interpreted as an argument that single-peakedness is essentially a form of consensus.

4.1 Convexity of Voting Rules

We start by arguing that even the Plurality rule is not convex in the sense of Definition 2.2. This argument can be adapted to apply to many other voting rules.

Proposition 4.1. *The Plurality rule is not convex.*

Proof. Let $C = \{a, b, c\}$. Let $E^1 = (C, \mathcal{R}^1)$ be a 3-voter election with $\mathcal{R}^1 = (R^1_1, R^1_2, R^1_3)$, where $R^1_1 = a \succ b \succ c$, $R^1_2 = b \succ c \succ a$, and $R^1_3 = c \succ b \succ a$. Further, let $E^2 = (C, \mathcal{R}^2)$ be a 3-voter election with $\mathcal{R}^2 = (R^2_1, R^2_2, R^2_3)$, where $R^2_1 = R^1_1$, $R^2_2 = R^1_3$, and $R^2_3 = R^1_2$, i.e., E^2 is obtained from E^1 by swapping the last two voters. Clearly, the shortest path from E^1 to E^2 involves swapping b and c in each of the last two votes. However, after any such swap the candidate who was swapped upwards becomes the unique winner. Thus, even though a is a co-winner in both E^1 and E^2, it is not a co-winner in some election on every shortest path from E^1 to E^2. □

However, Plurality can be shown to be weakly convex.

Theorem 4.2. *The Plurality rule is weakly convex.*

Proof. Consider two elections, $E^1 = (C, \mathcal{R}^1)$ and $E^2 = (C, \mathcal{R}^2)$, that both have some candidate $a \in C$ as their unique Plurality winner. For each $c \in C$ and $i = 1, 2$, let $s^i(c)$ denote the Plurality score of candidate c in E^i.

We will construct a shortest path from E^1 to E^2 on which a is the election co-winner as follows. We start with $E = E^1$. Then, for every index $i \in [n]$ such that $a \notin top(R^1_i)$ but $a \in top(R^2_i)$, we move a to the top of the i-th vote, and then optimally rearrange the rest of the candidates so that the i-th vote coincides with R^2_i. Clearly, each of these swaps is needed to get from E^1 to E^2, so the sequence of elections that we go through during this process is a prefix of some shortest path from E^1 to E^2. Let s be the Plurality score of a at this point; clearly, we have $s \geq \max\{s^1(a), s^2(a)\}$. Since a is the unique Plurality winner in both E^1 and E^2, the score of any other candidate at this point is at most $s-1$.

Let $S = \{i \mid top(R^1_i) \neq a, top(R^2_i) \neq a, top(R^1_i) \neq top(R^2_i)\}$. We will now process the votes in S one by one. Specifically, we pick an arbitrary voter $i \in S$, set $c = top(R^2_i)$, move c to the top, and then optimally rearrange the rest of the vote so that it coincides with R^2_i. During each step of this transformation, the top-ranked candidate in the i-th vote is either $top(R^1_i)$ or c; hence, the score of each candidate other than c remains at most $s-1$. Now, if also the score of c remains at most $s-1$, we simply pick another arbitrary voter $j \in S$ and repeat the process. On the other hand, if the score of c becomes $s > s^2(c)$, there must exist another voter $j \in S$ with $top(R^1_j) = c$, $top(R^2_j) \neq c$. Then we can apply the same procedure to the j-th voter; as a result, the score of c goes back to $s-1$ (though the score of $top(R^2_j)$ may now go up to s). We repeat this step until we have handled all voters in S, picking the next voter to process either arbitrarily or according to the candidate promoted to the top at the previous step. An easy inductive argument shows that if at step k candidate d was promoted to the top, the scores of all candidates in $C \setminus \{a, d\}$ after step k are at most $s-1$, and d's score is at most s.

Finally, we transform each vote i with $top(R^1_i) = a$, $top(R^2_i) \neq a$ into R^2_i by first moving $top(R^2_i)$ to the top and then optimally rearranging the rest of the vote. At each step of this process a's score is at least $s^2(a)$ and the score of each other candidate d is at most $s^2(d) < s^2(a)$. Since at each step of the transformation we perform the necessary swaps only, we have found a shortest path from E^1 to E^2. □

Other common voting rules, such as, e.g., the Borda rule, are not even weakly convex.

Theorem 4.3. *The Borda rule is not weakly convex.*

Proof. Let $C = \{a, b, c, d, e\}$. We construct a 4-voter election E over C as follows. The first two votes are $a \succ b \succ c \succ d \succ e$ and $e \succ c \succ b \succ a \succ d$ and the next two votes are $e \succ c \succ d \succ a \succ b$ and $a \succ e \succ c \succ b \succ d$. We define election E' to be identical to E except we replace the first two voters with, respectively, $c \succ d \succ a \succ b \succ e$ and $e \succ a \succ d \succ c \succ b$.

It is easy to verify that in both elections e has 11 points, a and c have 10 points each, and b and d have at most 10 points each. Thus, e is the unique winner of both E and E'. However, we claim that on any shortest path from E to E' (a) the Borda score of either a or c goes up to at least 12, and (b) the Borda score of e remains unchanged. Together, (a) and (b) imply that every shortest path from E to E' goes through an election E'' where e is not a Borda winner and hence the Borda rule is not weakly convex.

To see why (a) holds, note that on each shortest path from E to E', the second swap of c upwards is with a (in the first vote), and the second swap of a upwards is with c (in the second vote). Now, one of these swaps, call it S, has to happen first. However, by construction of E and E', this means that the candidate being swapped upwards in S could not yet have been swapped downwards in the other vote. Thus, after S this candidate has at least 12 points. On the other hand, (b) holds because in the respective votes of E and E' candidate e is ranked either first or last. □

It is easy to see that the construction used in the proof of Theorem 4.3 works for many other scoring rules. By tweaking this construction slightly, we can extend Theorem 4.3 as follows.

Corollary 4.4. *Suppose that a scoring rule $\mathcal{F}_\alpha$ satisfies one of the following conditions: (1) there exists an $i \in \{1, \ldots, m-2\}$ such that $\alpha_i - \alpha_{i+1} = 2$, $\alpha_{i+1} - \alpha_{i+2} = 1$; (2) there exists an $i \in \{1, \ldots, m-2\}$, a $j \in \{1, \ldots, i-1\}$ and a $k \in \{1, \ldots, j-1\}$ such that $\alpha_i - \alpha_{i+1} = 1$, $\alpha_{i+1} - \alpha_{i+2} = 1$, $\alpha_j - \alpha_{j+1} \geq 1$, and $\alpha_k - \alpha_{k+1} \geq 1$. Then $\mathcal{F}_\alpha$ is not weakly convex.*

A somewhat similar argument shows that neither Copeland nor Maximin is weakly convex.

Theorem 4.5. *Neither the Copeland rule nor the Maximin rule is weakly convex.*

Given that Borda, Copeland, and Maximin are not weakly convex, how difficult is it to tell whether two elections constitute a witness to their nonconvexity?

Definition 4.6. *Let $\mathcal{F}$ be a voting rule. An instance of the $\mathcal{F}$-ShortestPath problem is given by a set of candidates C and two preference profiles $\mathcal{R}$ and $\mathcal{R}'$ over C with*

$\mathcal{F}(C, \mathcal{R}) = \mathcal{F}(C, \mathcal{R}') = \{a\}$ *for some* $a \in C$. *This is a "yes" instance if there is a shortest path between* $(C, \mathcal{R})$ *and* $(C, \mathcal{R}')$ *in* $\mathcal{G}$ *that is contained in* $col_a(\mathcal{G}, \mathcal{F})$, *and a "no"-instance otherwise.*

This problem is NP-complete even for the Borda rule (as well as for Maximin and Copeland). We interpret this result as saying that the topology of $\mathcal{G}$ (colored by these rules) is quite intricate.

THEOREM 4.7. *For each* $\mathcal{F}$ *in* {*Borda, Maximin, Copeland*}, *the problem* $\mathcal{F}$-SHORTESTPATH *is* NP-*complete.*

4.2 Convexity of Consensus Classes

We now move on to the discussion of convexity of consensus classes. A *consensus class* $\mathcal{F}$ is a mapping that given an election $E = (C, \mathcal{R})$ either outputs a single (undisputed) consensus winner of this election or indicates that there is no consensus winner. The notion of a consensus class lies at the heart of distance-based interpretation of voting rules [18, 14, 13], where to determine the winners of an election E, we identify the nearest consensus election(s) (with respect to a given distance function, which does not need to be the swap distance), and output the consensus winners of these elections. The four most-often studied consensus classes are:

1. Strong unanimity ($\mathcal{S}$): A candidate a is a strong unanimity winner of an election E if all voters in E have the same preference order and a is the top candidate with respect to this order.
2. Weak unanimity ($\mathcal{U}$): A candidate a is a weak unanimity winner of an election E if a is ranked first by all voters in E.
3. Majority ($\mathcal{M}$): A candidate a is a majority winner of an election E if a is ranked first by more than half of the voters in E.
4. Condorcet ($\mathcal{C}$): A candidate a is a consensus winner in election E with respect to the Condorcet consensus if a is the Condorcet winner in this election.

The notion of convexity can be naturally extended from voting rules to consensus classes. Since each election either has a single consensus winner or has no winners at all, for consensus classes the notions of convexity and weak convexity are equivalent. We will now show that $\mathcal{U}$, $\mathcal{M}$, and $\mathcal{C}$ are all convex, but the strong unanimity consensus $\mathcal{S}$ is not.

THEOREM 4.8. *The consensus classes* $\mathcal{U}$, $\mathcal{M}$, *and* $\mathcal{C}$ *are convex.*

PROOF. Let $\mathcal{K}$ be one of the consensus classes $\mathcal{U}$, $\mathcal{M}$, and $\mathcal{C}$, and let $E^1 = (C, \mathcal{R}^1)$ and $E^2 = (C, \mathcal{R}^2)$ be two elections that have the same $\mathcal{K}$-consensus winner $a = \mathcal{K}(E^1) = \mathcal{K}(E^2)$. For each $\mathcal{K} \in \{\mathcal{U}, \mathcal{M}, \mathcal{C}\}$ we will show that there is a shortest path from E^1 to E^2 where a is the $\mathcal{K}$-consensus winner of every intermediate election.

We consider the majority consensus first. The proof is very similar to that of Proposition 3.2. Let $\widehat{E^1}$ (respectively, $\widehat{E^2}$) be an election obtained from E^1 (respectively, from E^2) by shifting a to the top of each vote in $\mathcal{R}^1$ (respectively, in $\mathcal{R}^2$) such that the corresponding vote in $\mathcal{R}^2$ (respectively, in $\mathcal{R}^1$) ranks a first. It is easy to see that there is a shortest path from E^1 to E^2 that first goes from E^1 to $\widehat{E^1}$, then to $\widehat{E^2}$, and finally to E^2, and such that a is the majority winner of every election on that path. The same proof works for weak unanimity consensus $\mathcal{U}$.

Let us now consider the Condorcet consensus. We proceed in two stages. For each stage, we describe the order in which swaps of adjacent candidates should be performed.

At the first stage, we consider the votes in $\mathcal{R}^1$ one by one. Pick a vote R_i^1 in $\mathcal{R}^1$. Candidate a partitions it into two blocks: the candidates ranked above a and those ranked below a. We sort the upper block according to R_i^2 by bubble sort, i.e., by performing the necessary swaps only. This transformation is on a shortest path from E^1 to E^2, and, since it does not change the order of a relative to the other candidates, a remains the Condorcet winner at each step. Now the upper block can be further subdivided into two subblocks (one or both can be empty): the candidates ranked above a in R_i^2 and those ranked below it in R_2^i. Note that the candidates in each subblock occur contiguously in the vote after the sorting step. We swap a with all candidates in the lower subblock of the upper block, i.e., the candidates that are ranked above a in R_i^1, but below it in R_i^2; clearly, all these swaps are necessary. Next, we bubble-sort all candidates ranked below a in the vote according to R_i^2. Again, all these swaps are necessary and do not change the order of a relative to the other candidates, so we stay on a shortest path and a remains the Condorcet winner. The first stage ends when we have done this for all votes.

At the beginning of the second stage, the lower block of the i-th vote can be subdivided into two contiguous subblocks: the candidates that appear above a in R_i^2 and those that appear below it in R_i^2, with the former subblock appearing above the latter. To get to E^2, it suffices to swap a with all candidates in the upper subblock of the lower block of each vote; this is what we do during the second stage. Since a is only shifted downwards during this procedure, and it is the Condorcet winner at the end, it is also the Condorcet winner at all the intermediate steps. □

Given that each of $\mathcal{U}$, $\mathcal{M}$, and $\mathcal{C}$ are convex (and, consequently, connected), it is natural to ask about the status of strong unanimity. Here the answer is somewhat disappointing: it is easy to see that strong unanimity is not convex, and, moreover, not even connected. Indeed, for every strong unanimity election E with $m > 1$ candidates and more than one voter, any single swap of adjacent candidates transforms E into an election without a strong unanimity consensus winner.

However, there is a different interpretation of the strong unanimity consensus under which it is both connected and convex (albeit for trivial reasons). While for weak unanimity, majority, and Condorcet consensus we focused on the consensus winner, for the case of strong unanimity one could argue that we have more: a consensus ranking. Then, if one colored the swap-distance graph not by (consensus) winners, but by consensus rankings, then the resulting coloring would of course be connected and convex (because there would be exactly one vertex of each color). We leave it to the readers to decide whether they find this interpretation convincing.

4.3 Convexity of Single-Peaked Elections

We conclude our discussion of convexity by considering single-peaked elections [4].

DEFINITION 4.9. *Let* C *be a set of candidates and let* L *be a linear order over* C. *We say that a preference order* $\succ$

is consistent with L *if for every triple of distinct candidates* $(a, b, c) \in C^3$ *it holds that* $((a\ L\ b\ L\ c) \vee (c\ L\ b\ L\ a)) \implies (a \succ b \implies a \succ c)$. *An election* $E = (C, \mathcal{R})$ *is* single-peaked with respect to L *if every vote in* E *is consistent with* L.

Single-peaked elections are very well studied and have a number of desirable properties: for instance, every single-peaked election has a (weak) Condorcet winner. It turns out that the set of all elections that are single-peaked with respect to a given axis is convex (we omit the proof).

THEOREM 4.10. *For every candidate set* C*, every linear order* L *over* C *and every* $n > 0$*, the set of* n*-voter elections that are single-peaked with respect to* L *is convex.*

5. CONCLUSIONS

Swap distance is one of the most natural measures of similarity between elections (with a fixed candidate set and a fixed number of voters). In this paper we have considered the shapes of sets formed by elections with a particular winner (according to a given voting rule or a given consensus class) in the graph of elections defined via the swap distance. To do so, we have adapted the standard geometric notions of connectedness and convexity. It turns out that while all the voting rules that we have considered are connected (that is, the sets of elections with a given winner are connected), none of them is convex, and only the Plurality rule is weakly convex. On the other hand, we have identified interesting examples of convex sets of elections, such as consensus elections with a given winner (for three standard notions of consensus) and elections that are single-peaked with respect to a given ordering of candidates. In the course of our discussion, we have defined a new notion of monotonicity (feeble monotonicity), and argued that for unanimity-consistent voting rules convexity is a strengthening of feeble monotonicity. It would be very interesting to explore further connections between connectivity, convexity, and other properties of voting rules.

Acknowledgements Svetlana Obraztsova was supported by the project ALGONOW of the research funding program THALIS (co-financed by the European Social Fund-ESF and Greek national funds) as well as by grants of Russian government (FCP and RFFI). Edith Elkind was supported by National Research Foundation (Singapore) under grant 2009-08. Piotr Faliszewski was supported in part by AGH Univ. grant 11.11.120.865, Homing/Powroty program, and by Poland's NCN's grant DEC-2011/03/B/ST6/01393.

6. REFERENCES

[1] K. Arrow, A. Sen, and K. Suzumura, editors. *Handbook of Social Choice and Welfare, Volume 1*. Elsevier, 2002.

[2] J. Bartholdi, C. Tovey, and M. Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6(3):227–241, 1989.

[3] 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.

[4] D. Black. *The Theory of Committees and Elections*. Cambridge University Press, 1958.

[5] 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.

[6] F. Brandt. Some remarks on Dodgson's voting rule. *Mathematical Logic Quarterly*, 55(4):460–463, 2009.

[7] F. Brandt, V. Conitzer, and U. Endriss. Computational social choice. In G. Weiß, editor, *Multiagent Systems*. MIT Press, 2012. Forthcoming.

[8] V. Conitzer, M. Rognlie, and L. Xia. Preference functions that score rankings and maximum likelihood estimation. In *Proceedings of IJCAI-09*, pages 109–115, 2009.

[9] 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.

[10] J. Davies, N. Narodytska, and T. Walsh. Eliminating the weakest link: Making manipulation intractable? In *Proceedings of AAAI-12*, pages 1333–1339, 2012.

[11] M. Deza and E. Deza. *Encyclopedia of Distances*. Springer-Verlag, 2009.

[12] E. Elkind and P. Faliszewski. Approximation algorithms for campaign management. In *Proceedings of WINE-10*, pages 473–482, 2010.

[13] E. Elkind, P. Faliszewski, and A. Slinko. On the role of distances in defining voting rules. In *Proceedings of AAMAS-10*, pages 375–382, 2010.

[14] E. Elkind, P. Faliszewski, and A. Slinko. Rationalizations of condorcet-consistent rules via distances of Hamming type. *Social Choice and Welfare*, 39(4):891–905, 2012.

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

[16] P. Faliszewski, E. Hemaspaandra, and L. Hemaspaandra. Using complexity to protect elections. *Commun. ACM*, 53(11):74–82, 2010.

[17] E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Anyone but him: The complexity of precluding an alternative. *Artificial Intelligence*, 171(5–6):255–285, 2007.

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

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

[20] N. Narodytska, T. Walsh, and L. Xia. Manipulation of Nanson's and Baldwin's rule. In *Proceedings of AAAI-11*, pages 713–718, 2011.

[21] P. Pattanaik. Positional rules of collective decision-making. In K. Arrow, A. Sen, and K. Suzumura, editors, *Handbook of Social Choice and Welfare, Volume 1*, pages 361–394. Elsevier, 2002.

[22] D. Saari. Consistency of decision processes. *Annals of Operations Research*, 23(1):103–137, 1990.

[23] D. Saari. *Basic Geometry of Voting*. Springer-Verlag, 2003.

[24] I. Schlotter, P. Faliszewski, and E. Elkind. Campaign management under approval-driven voting rules. In *Proceedings of AAAI-11*, pages 726–731, 2011.

[25] J. Smith. Aggregation of preferences with variable electorate. *Econometrica*, 41:1027–1041, 1973.