On Swap Convexity of Voting Rules

Svetlana Obraztsova
Nanyang Technological University
Singapore
svetlana.obraztsova@gmail.com

Edith Elkind
University of Oxford
Oxford, UK
elkind@cs.ox.ac.uk

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

Abstract

Obraztsova et al. (2013) have recently proposed an intriguing convexity axiom for voting rules. This axiom imposes conditions on the shape of the sets of elections with a given candidate as a winner. However, this new axiom is both too weak and too strong: it is too weak because it defines a set to be convex if for any two elements of the set some shortest path between them lies within the set, whereas the standard definition of convexity requires all shortest paths between two elements to lie within the set, and it is too strong because common voting rules do not satisfy this axiom. In this paper, we (1) propose several families of voting rules that are convex in the sense of Obraztsova et al.; (2) put forward a weaker notion of convexity that is satisfied by most common voting rules; (3) prove impossibility results for a variant of this definition that considers all, rather than some shortest paths.

1 Introduction

Voting and elections are traditional ways in which societies of agents make decisions (Moulin 1991; Arrow, Sen, and Suzumura 2002). Since these decisions often carry significant weight, it is important to identify suitable voting rules for each application scenario. The requirements that a voting rule should satisfy, also called *axioms*, are application-dependent: e.g., in some settings in may be desirable to treat all opinions in the same way (voting rules with this property are called *anonymous*), while in other settings it may be more important to avoid rules that suffer from the so-called no-show paradox, where voters have an incentive to withdraw from the election. The axioms themselves can be classified in several ways. For instance, some axioms constrain the behavior of a voting rule in individual elections: this is the case, e.g., for the unanimity axiom, which requires that if all voters rank some candidate first, the rule outputs this candidate. Other axioms relate the behavior of the rule in elections obtained from each other by some transformation: e.g., the anonymity axiom describes the behavior of the rule on profiles obtained from each other by permuting votes.

Recently, Obraztsova et al. (2013) proposed two new axioms, which recognize the fact that collective decision-making is a dynamic process, and aim to capture how voting rules should behave when voters' preferences evolve. Specifically, consider a society that moves between two states A and B such that candidate a is considered to be the best choice in both states. We assume that the voters' preferences change gradually, that is, at each step some voter reports a small change in her beliefs. One can then expect that it should be possible to transition from A to B so that a remains the collective choice along the way; more ambitiously, we may require that this remains the case when we pick the *shortest* route from A to B.

Formally, Obraztsova et al. (2013) construct a graph whose vertices are all n-voter m-candidate elections, and there is an edge between two elections if one can be obtained from the other by swapping two adjacent candidates in some voter's ranking. The shortest path distance in this graph is the well-known *swap distance* (Deza and Deza 2009). Obraztsova et al. (2013) are interested in 'shapes' of sets of vertices of this graph that have a common defining property: for instance, these can be elections where a given candidate is ranked first by all voters or by a strict majority of voters (such elections are known, respectively, as unanimity consensus elections and majority consensus elections), elections that are single-peaked with respect to a fixed axis, or elections where a given candidate is a winner under a specific voting rule.

Yet, the results of Obraztsova et al. (2013) turn out to be rather discouraging. While they do show that for most natural voting rules the sets of elections with a given winner are connected in $\mathcal{G}_{n,m}$, i.e., it is possible to get from one election with winner a to another such election so that a stays an election winner along the way, it may be impossible to do so if one insists on staying on the shortest path between the two elections. That is, such sets are not convex, even in the sense of the somewhat non-orthodox definition that they use[1] (we elaborate on the differences between their

[1] We remark that this is not the first attempt to define a notion of convexity for voting rules: in particular, Saari (1990; 2003) proposed another notion of convexity, but his definition is very different in spirit, and, in particular, imposes constraints on the behavior of voting rules on elections with a different number of voters.

definition and the more standard one in Section 3; briefly put, their definition is quite natural, but less demanding). In fact, Obraztsova et al. (2013) are unable to identify a single voting rule that satisfies their convexity axiom (the only possible exception is the Plurality rule, which satisfies a relaxed variant of their definition). On the other hand, they show that the sets consisting of unanimity/majority consensus elections with a given winner are convex, and this is also the case for sets of elections that are single-peaked with respect to a fixed axis. Their conclusion, then, is that convexity is an appropriate notion of 'niceness' for consensus notions (such as the unanimity consensus or the majority consensus), but not for voting rules.

In this paper, we revisit this conclusion. First, we focus on the notion of convexity proposed by Obraztsova et al. (2013), and identify a family of rules that are convex, anonymous, neutral and unanimous (see Section 2 for the statements of these axioms). While it turns out that it is not too difficult to satisfy these axioms simultaneously if one is willing to accept rules that are not very decisive, we explore ways to increase decisiveness while maintaining the other desiderata, and identify a voting rule with the required properties that returns a unique winner on many profiles. Nevertheless, the rule we propose is not a commonly used voting rule, such as Plurality, Borda, or STV, and, indeed, it is already known that these rules are not convex in the sense of Obraztsova et al. (2013). Thus, if we are interested in a notion of convexity that can be used to distinguish among common voting rules, we need a different definition.

To achieve this goal, we make use of the notion of star convexity that is firmly rooted in the topology literature. Intuitively, a set is star convex if it has a 'central' point such that all other points in the set can be reached from this point along shortest paths that stay within the set. This notion turns out to be closely related to the notion of feeble monotonicity (which Obraztsova et al. used to show their connectivity results), and for many (but not all!) voting rules we can show that the set of elections that have a given candidate as a winner is star convex.

Emboldened by these positive results, we revisit the definition of convexity used by Obraztsova et al. (2013). Their definition is somewhat non-orthodox in that, instead of requiring that all shortest paths between two elements of a set lie within this set, they only require that this holds for *some* shortest path. Now, of course, the 'all-paths' definition is even more demanding than the original one, but given that we have identified a voting rule that is convex according to the original definition, perhaps there are reasonable rules that satisfy this more stringent definition as well? We show that this is not the case, by proving an impossibility result.

Our results indicate that convexity, as defined by Obraztsova et al. (2013), is a meaningful, if stringent condition not just for consensus classes, but also for voting rules. In contrast, star convexity turns out to be closely related to existing notions of monotonicity, while building on a very different intuition; this axiom provides additional justification for rules such as STV that are not monotone in the standard sense, but are nevertheless considered fairly well-behaved.

The remainder of this paper is structured as follows. First, in Section 2 we introduce the necessary background regarding voting rules. Then, in Section 3 we formulate the notions of convexity that we are going to consider, starting with the original definition of Obraztsova et al. (2013). Next, we present several families of voting rules that are convex with respect to the original definition (Section 4). Section 5 presents our results for star convexity. We conclude the paper by presenting the impossibility results regarding 'all-paths' convexity (Section 6). We omit some proofs due to space constraints.

2 Preliminaries

We write $[n]$ to denote the set $\{1, \dots, n\}$. Given a finite set of *candidates* $C = \{c_1, \dots, c_m\}$, a *vote* over C is a linear order over C; a *profile* over C is an ordered list $\mathcal{R} = (R_1, \dots, R_n)$ of n votes over C. Given a profile $\mathcal{R} = (R_1, \dots, R_n)$, we will refer to the elements of $[n]$ as *voters*; i.e., R_i is the vote of voter i in profile $\mathcal{R}$. We will sometimes write $\succ_i$ in place of R_i. 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 $b \in C \setminus \{a\}$. We say that two voters *disagree* on a pair of candidates $(a, b) \in C \times C$ if one of the voters ranks a above b, while the other voter ranks b above a; otherwise, we say that these voters *agree* on (a, b).

A *voting rule* is a mapping $\mathcal{F}$ that given a profile $\mathcal{R}$ over a set of candidates C outputs a non-empty set of candidates $W = \mathcal{F}(\mathcal{R}) \subseteq C$; the candidates in W are called the *winners* in $\mathcal{R}$ under the voting rule $\mathcal{F}$. A canonical example of a voting rule is the Plurality rule, which, given a profile $\mathcal{R} = (R_1, \dots, R_n)$ over a set of candidates C, outputs the set of all candidates $a \in C$ such that $|\{i \in [n] : top(R_i) = a\}| \geq |\{i \in [n] : top(R_i) = b\}|$ for all $b \in C$.

A candidate $a \in C$ is a *Condorcet winner* in a profile $\mathcal{R}$ over C if for each $c \in C \setminus \{a\}$ more than half of the votes in $\mathcal{R}$ rank a above c. We say that a profile $\mathcal{R}$ is a *unanimity consensus with winner* a if all votes in $\mathcal{R}$ rank a first; $\mathcal{R}$ is a *majority consensus with winner* a if more than half of the votes in $\mathcal{R}$ rank a first; and $\mathcal{R}$ is a *Condorcet consensus with winner* a if a is a Condorcet winner in $\mathcal{R}$.

Given a profile $\mathcal{R} = (R_1, \dots, R_n)$ over a set of candidates C, we say that a candidate $a \in C$ is *nominated* in $\mathcal{R}$ if $top(R_i) = a$ for some $i \in [n]$; we say that a is *popular* in $\mathcal{R}$ if $|\{i \in [n] : top_i(R_i) = a\}| \geq \frac{n}{|C|}$. We denote the set of all nominated candidates in $\mathcal{R}$ by $Nom(\mathcal{R})$, and the set of all popular candidates by $Pop(\mathcal{R})$. For any profile $\mathcal{R}$ we have $Pop(\mathcal{R}) \subseteq Nom(\mathcal{R})$ and, by the pigeonhole principle, $Pop(\mathcal{R}) \neq \emptyset$; thus, Nom and Pop are voting rules.

A vote R over a set of candidates C is said to be *single-peaked with respect to an ordering $<$ of C* if for every pair of candidates $a, b \in C \setminus \{top(R)\}$ such that $b < a < top(R)$ or $top(R) < a < b$ it holds that R ranks a above b; the order $<$ is often called an *axis*. A profile $(R_1, \dots, R_n)$ is said to be *single-peaked with respect to an axis $<$* if for every $i \in [n]$ the vote R_i is single-peaked with respect to $<$.

Axioms Given a permutation $\pi : C \to C$ and a vote R over C, let $\pi(R)$ denote the vote such that for every pair of candidates $(a, b) \in C \times C$ it holds that a is ranked above b in R if and only if $\pi(a)$ is ranked above $\pi(b)$ in $\pi(R)$. We extend this notation to profiles: given a permutation $\pi : C \to C$ and a profile $\mathcal{R} = (R_1, \ldots, R_n)$ over C, we denote by $\pi(\mathcal{R})$ the profile $(\pi(R_1), \ldots, \pi(R_n))$. Given a pair of candidates $(a, b) \in C \times C$ and a profile $\mathcal{R}$ over C, we denote by $\mathcal{R}_{a \leftrightarrow b}$ the profile $\pi(\mathcal{R})$, where π is given by $\pi(a) = b$, $\pi(b) = a$, $\pi(c) = c$ for $c \in C \setminus \{a, b\}$. Also, given a permutation $\sigma : [n] \to [n]$ and a profile $\mathcal{R} = (R_1, \ldots, R_n)$, we denote by $\sigma(\mathcal{R})$ the profile $(R_{\sigma(1)}, \ldots, R_{\sigma(n)})$.

Below we state the axioms that we consider in this paper.

Definition 2.1. *A voting rule $\mathcal{F}$ is* neutral *if for every set of candidates C, every profile $\mathcal{R}$ over C, every permutation $\pi : C \to C$, and every candidate $a \in C$ it holds that $a \in \mathcal{F}(\mathcal{R})$ if and only if $\pi(a) \in \mathcal{F}(\pi(\mathcal{R}))$.*

Definition 2.2. *A voting rule $\mathcal{F}$ is* anonymous *if for every profile $\mathcal{R} = (R_1, \ldots, R_n)$ and every permutation $\sigma : [n] \to [n]$ it holds that $\mathcal{F}(\mathcal{R}) = \mathcal{F}(\sigma(\mathcal{R}))$.*

Definition 2.3. *A voting rule $\mathcal{F}$ is* unanimous *if for every set of candidates C and every profile $\mathcal{R}$ over C in which each voter ranks the same candidate $a \in C$ in the top position in their vote it holds that $\mathcal{F}(\mathcal{R}) = \{a\}$.*

Definition 2.4. *A voting rule $\mathcal{F}$ is* weakly monotone *if for every set of candidates C, every profile $\mathcal{R}$ over C, and every candidate $a \in \mathcal{F}(\mathcal{R})$ it holds that every profile $\mathcal{R}'$ obtained from $\mathcal{R}$ by picking a vote R_i in $\mathcal{R}$ with $top(R_i) \neq a$ and swapping a with the candidate ranked right above her in R_i satisfies $a \in \mathcal{F}(\mathcal{R}')$.*

We also consider the following weaker notion of monotonicity proposed by Obraztsova et al. (2013).

Definition 2.5. *A voting rule $\mathcal{F}$ is* feebly monotone *if for every set of candidates C, every profile $\mathcal{R}$ over C, and every candidate $a \in \mathcal{F}(\mathcal{R})$ 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 profile $\mathcal{R}'$ obtained from $\mathcal{R}$ by swapping a with the candidate ranked right above her in R_i satisfies $a \in \mathcal{F}(\mathcal{R}')$.*

Swap distance and the graph $\mathcal{G}_{n,m}$ Given two votes R_i and R_j over a set of candidates C, the *swap distance* between R_i and R_j, denoted by $d_{swap}(R_i, R_j)$, is the number of pairs of candidates $(a, b) \in C \times C$ such that $a \succ_i b$, but $b \succ_j a$. This notion can be extended to profiles: given two n-voter profiles $\mathcal{R} = (R_1, \ldots, R_n)$ and $\mathcal{R}' = (R'_1, \ldots, R'_n)$ over a set of candidates C, the *swap distance* $d_{swap}(\mathcal{R}, \mathcal{R}')$ is computed as $\sum_{i \in [n]} d_{swap}(R_i, R'_i)$.

Given a set of candidates C, $|C| = m$, and a positive integer n, consider the graph $\mathcal{G}_{n,m} = (\mathcal{V}, \mathcal{E})$ whose vertices are all n-voter profiles over C and where we have an edge between two profiles $\mathcal{R}$ and $\mathcal{R}'$ if and only if $d_{swap}(\mathcal{R}, \mathcal{R}') = 1$; that is, we have an edge between $\mathcal{R} = (R_1, \ldots R_n)$ and $\mathcal{R}' = (R'_1, \ldots, R'_n)$ if and only if there is a voter $j \in [n]$ such that $d_{swap}(R_j, R'_j) = 1$ and for each $i \in [n] \setminus \{j\}$ we have $R_i = R'_i$. We write $\mathcal{G}$ instead of $\mathcal{G}_{n,m}$ when the values of n and m are clear from the context. Note that the length of a shortest path between two vertices of this graph is exactly the swap distance between them. A voting rule $\mathcal{F}$ can be viewed as a coloring of this graph (albeit one where each vertex may have several different colors): we identify each candidate with a color, so that the set of colors assigned to a profile $\mathcal{R}$ is exactly $\mathcal{F}(\mathcal{R})$. Given a candidate $a \in C$, we denote by $col_a(\mathcal{G}_{n,m}, \mathcal{F})$ the set of all vertices that are colored with a under $\mathcal{F}$: $col_a(\mathcal{G}_{n,m}, \mathcal{F}) = \{\mathcal{R} : a \in \mathcal{F}(\mathcal{R})\}$.

3 Definitions of Convexity

Obraztsova et al. (2013) defined convexity as follows.

Definition 3.1. *A subset $\mathcal{V}'$ of the vertices of the graph $\mathcal{G}_{n,m} = (\mathcal{V}, \mathcal{E})$ is said to be* convex *if for every $\mathcal{R}, \mathcal{R}' \in \mathcal{V}'$ there exists a path from $\mathcal{R}$ to $\mathcal{R}'$ that lies entirely within $\mathcal{V}'$ and has length d, where d is the length of a shortest path between $\mathcal{R}$ and $\mathcal{R}'$ in $\mathcal{G}$. A voting rule $\mathcal{F}$ is said to be* convex *if for every set of candidates C, every $n \geq 1$, and every $a \in C$ the set $col_a(\mathcal{G}_{n,|C|}, \mathcal{F})$ is convex.*

$\forall$-Convexity Definition 3.1 is different from the usual one, used in topology, in that instead of requiring *all* shortest paths to stay within the set, it requires this for only *one* of them. In our context, the standard definition of convexity takes the following form.

Definition 3.2. *A subset $\mathcal{V}'$ of the vertices of the graph $\mathcal{G}_{n,m} = (\mathcal{V}, \mathcal{E})$ is* $\forall$-convex *if for every $\mathcal{R}, \mathcal{R}' \in \mathcal{V}'$ and every shortest path ρ between $\mathcal{R}$ and $\mathcal{R}'$ in $\mathcal{G}$ it holds that every profile $\mathcal{R}''$ on ρ belongs to $\mathcal{V}'$. A voting rule $\mathcal{F}$ is* $\forall$-convex *if for every set of candidates C, every $n \geq 1$, and every $a \in C$ the set $col_a(\mathcal{G}_{n,|C|}, \mathcal{F})$ is $\forall$-convex.*

To avoid confusion, from now on we will refer to the notion of convexity from Definition 3.1 as $\exists$-convexity. Clearly, if a set of profiles is $\forall$-convex then it is $\exists$-convex, but the converse is not true, as shown by the following example.

Example 3.3. Fix $n \geq 1$, and consider the space of all n-voter profiles over the set of candidates $\{a, b\}$.

Theorem 4.8 in the work of Obraztsova et al. implies that the set of profiles where a strict majority of voters prefer a to b is $\exists$-convex. However, this set if not $\forall$-convex for all $n \geq 3$. Indeed, let $k = \lfloor \frac{n}{2} \rfloor + 1$, and consider a pair of profiles $\mathcal{R}$, $\mathcal{R}'$, where in $\mathcal{R}$ the first k voters prefer a to b and the remaining voters prefer b to a, whereas in $\mathcal{R}'$ the last k voters prefer a to b and the remaining voters prefer b to a. In both profiles a is a strict majority winner, but there is a path from $\mathcal{R}$ to $\mathcal{R}'$ that goes through a profile where $2(n - k) \geq n/2$ voters rank b first: on this path, we first swap a and b in the preferences of the first $n - k$ voters, and then swap them in the preferences of the last $n - k$ voters.

We discuss $\forall$-convexity in more detail in Section 6.

Star Convexity In geometry, a subset S of $\mathbb{R}^N$ is said to be *star convex*, or *star-shaped*, if there exists a point $c \in S$ such that for every $a \in S$ the shortest path between a and c lies entirely in S; the set of all points c that have this property is called the *kernel* of S (Preparata and Shamos 1985). To extend this definition to subsets of the graph $\mathcal{G}_{n,m}$, we have to decide whether we are interested in *all* shortest paths

between c and a or *some* shortest path between these points; we refer to the corresponding notions as *star* $\forall$*-convexity* and *star* $\exists$*-convexity*. Formally, we have the following definition.

Definition 3.4. *A subset $\mathcal{V}'$ of vertices of the graph $\mathcal{G}_{n,m} = (\mathcal{V}, \mathcal{E})$ is* star $\forall$-convex *if there exists a profile $\mathcal{R} \in \mathcal{V}'$ such that for every $\mathcal{R}' \in \mathcal{V}'$ all shortest paths between $\mathcal{R}$ and $\mathcal{R}'$ lie entirely within $\mathcal{V}'$; the set of all such $\mathcal{R}$ is called the* $\forall$-kernel *of $\mathcal{V}'$. Further, $\mathcal{V}'$ is* star $\exists$-convex *if there exists a profile $\mathcal{R} \in \mathcal{V}'$ such that for every $\mathcal{R}' \in \mathcal{V}'$ there exists a path of length d between $\mathcal{R}$ and $\mathcal{R}'$ that lies entirely within $\mathcal{V}'$, where d is the length of a shortest path between $\mathcal{R}$ and $\mathcal{R}'$ in $\mathcal{G}_{n,m}$; the set of all such $\mathcal{R}$ is called the* $\exists$-kernel *of $\mathcal{V}'$. A voting rule $\mathcal{F}$ is* star $\forall$-convex *(respectively,* star $\exists$-convex*) if for every set of candidates C, every $n \geq 1$, and every $a \in C$ the set $col_a(\mathcal{G}_{n,|C|}, \mathcal{F})$ is star $\forall$-convex (respectively, star $\exists$-convex).*

We explore these two notions of star convexity in Section 5.

4 $\exists$-Convexity

Obraztsova et al. (2013) did not identify any $\exists$-convex voting rules in their work. To close this gap, we describe several families of rules that are anonymous, neutral, unanimous and $\exists$-convex. We start with the least decisive rule in this class, and then explore various ways to refine it.

Let $\mathcal{F}_{\mathcal{U}}$ be a voting rule defined as follows: given a profile $\mathcal{R} = (R_1, \dots, R_n)$, if there exists an $a \in C$ such that $top(R_i) = a$ for all $i \in [n]$, output $\{a\}$; else, output C.

Theorem 4.1. *$\mathcal{F}_{\mathcal{U}}$ is anonymous, neutral and unanimous. It is $\exists$-convex whenever $n \geq 3$.*

Proof. It is immediate that $\mathcal{F}_{\mathcal{U}}$ is anonymous, neutral and unanimous. Now, suppose that $n \geq 3$, and consider two profiles $\mathcal{R} = (R_1, \dots, R_n)$ and $\mathcal{R}' = (R'_1, \dots, R'_n)$ over a candidate set C such that for some candidate $a \in C$ we have $a \in \mathcal{F}_{\mathcal{U}}(\mathcal{R}) \cap \mathcal{F}_{\mathcal{U}}(\mathcal{R}')$, or, equivalently, for each candidate $b \in C \setminus \{a\}$ neither $\mathcal{R}$ nor $\mathcal{R}'$ is a unanimity consensus with winner b. We will now describe a shortest path from $\mathcal{R}$ to $\mathcal{R}'$ that avoids all profiles where all voters rank some candidate $b \in C \setminus \{a\}$ first.

First, we use bubble sort to rearrange the candidates in the bottom $m-1$ positions in each vote R_i so that they are ordered according to R'_i. Formally, we consider a maximal sequence of profiles starting with $\mathcal{R}$ where each profile is obtained from its predecessor by picking a voter $i \in [n]$ and two candidates $b, c \in C \setminus top(R_i)$ that are adjacent in i's preference order and such that the current vote of voter i disagrees with R'_i on (b, c), and ordering b and c according to R'_i. Let $\mathcal{R}^* = (R^*_1, \dots, R^*_n)$ be the last profile in this sequence. This sequence of profiles forms a prefix of a shortest path from $\mathcal{R}$ to $\mathcal{R}'$, and, since we do not touch the top candidate in any vote, we have $\mathcal{F}_{\mathcal{U}}(\mathcal{R}^*) = \mathcal{F}_{\mathcal{U}}(\mathcal{R})$.

Now, if $top(R_i) \neq top(R'_i)$ for some $i \in [n]$, then $top(R'_i)$ is ranked in the second position in R^*_i. Thus, to transform $\mathcal{R}^*$ into $\mathcal{R}'$ it remains to swap the first and the second candidate in R^*_i for each $i \in [n]$ such that $top(R_i) \neq top(R'_i)$.

Suppose first that $top(R'_i) = a$ for some $i \in [n]$. If $top(R_i) = a$ as well, we can perform the remaining swaps in any order: at any point in time, at least one voter ranks a first, and this ensures that a is among the winners under $\mathcal{F}_{\mathcal{U}}$. If $top(R_i) \neq a$, we can first swap a and $top(R_i)$ in the preferences of the first voter: from then on, there exists a voter than ranks a first.

On the other hand, suppose that $top(R'_i) \neq a$ for all $i \in [n]$. Since $a \in \mathcal{F}_{\mathcal{U}}(\mathcal{R}')$, there exist two distinct candidates $b, c \in C$ and two voters $i, j \in [n]$ such that $top(R'_i) = b$, $top(R'_j) = c$. As we assume that $n \geq 3$, there is also a voter $k \in [n] \setminus \{i, j\}$; without loss of generality, we can assume that $top(R_k) \neq b$. We can then first swap b and $top(R_i)$ in the preferences of the i-th voter (and do nothing if $top(R_i) = b$), and then swap c and $top(R_j)$ in the preferences of the j-th voter. We then perform the remaining swaps. After the first step, the i-th voter and the k-th voter rank different candidates in the first position, and from the second step onward the i-th and the j-th voter rank different candidates in the first position, so our rule outputs C on all these profiles. □

The rule $\mathcal{F}_{\mathcal{U}}$ is the minimally decisive rule that satisfies anonymity, neutrality and unanimity: it outputs a single winner if and only if this is required by the unanimity axiom, and otherwise it outputs the entire set C. There are two ways to obtain a more decisive rule based on similar intuition: first, we can try to broaden the range of scenarios where our rule identifies a clear winner (e.g., by replacing the unanimity consensus with Condorcet consensus), and second, even when there is no clear winner, we can try to eliminate the least popular candidates and output $Nom(\mathcal{R})$ or, even more decisively, $Pop(\mathcal{R})$ instead of C. It turns out that these ideas can be combined.

Given a profile $\mathcal{R} = (R_1, \dots, R_n)$ over a set of candidates C and $q > \frac{n}{2}$, we say that a candidate $a \in C$ is a *q-Condorcet winner* if for every $b \in C \setminus \{a\}$ at least q voters prefer a to b. Let $\mathcal{F}_{q,Cond}$ be the rule that, given a profile $\mathcal{R}$ over C, outputs a candidate $a \in C$ if she is a q-Condorcet winner in $\mathcal{R}$; if there is no such candidate, $\mathcal{F}_{q,Cond}$ outputs C. We will also consider variants of this rule that, given a profile $\mathcal{R}$ that does not have a q-Condorcet winner, output all candidates in $Nom(\mathcal{R})$, or, alternatively, in $Pop(\mathcal{R})$; we denote these rules by $\mathcal{F}^{Nom}_{q,Cond}$ and $\mathcal{F}^{Pop}_{q,Cond}$ It turns out that all these rules are $\exists$-convex whenever $q \geq \lfloor \frac{n}{2} \rfloor + 2$.

Theorem 4.2. *For every integer q such that $\lfloor \frac{n}{2} \rfloor + 2 \leq q \leq n$ the rules $\mathcal{F}_{q,Cond}$, $\mathcal{F}^{Nom}_{q,Cond}$ and $\mathcal{F}^{Pop}_{q,Cond}$ are anonymous, neutral, and unanimous. They are $\exists$-convex whenever $n \geq 3$.*

The rules considered in Theorem 4.2 are considerably more decisive that $\mathcal{F}_{\mathcal{U}}$. This is particularly true for $\mathcal{F}^{Pop}_{q,Cond}$ with $q = \lfloor \frac{n}{2} \rfloor + 2$, which is the most decisive rule among the ones we consider. While variants of these rules with $q = \lfloor \frac{n}{2} \rfloor + 1$ are even more decisive, we can show that they are not $\exists$-convex, so Theorem 4.2 is tight in this regard.

The rules discussed in this section are similar in spirit to the classic rule of Black, which outputs the Condorcet winner when it exists and outputs all Borda winners otherwise (under the Borda rule, each candidate in an m-candidate election gets j points from each voter who ranks her in posi-

tion $m-j$, and the winners are the candidates with the highest total number of points) (Black 1958). Our rules, too, output the 'obvious' winner when it exists and otherwise they output all 'acceptable' candidates.

5 Star Convexity

In this section we provide an in-depth exploration of the two notions of star convexity proposed in Section 3. We start by observing that both the majority consensus and the Condorcet consensus are star $\forall$-convex (and hence star $\exists$-convex).

Proposition 5.1. *For any candidate set C and a candidate $a \in C$, both the set of all majority consensuses with winner a and the set of all Condorcet consensuses with winner a are star $\forall$-convex. Moreover, the unanimity consensus with winner a lies in the $\forall$-kernel of both of these sets.*

Proof. Let $\mathcal{R}$ be a unanimity consensus with winner a, and let $\mathcal{R}'$ be a majority consensus with winner a. At least $\lfloor n/2 \rfloor + 1$ voters in $\mathcal{R}'$ rank a first. In every profile on every shortest path from $\mathcal{R}'$ to $\mathcal{R}$ these voters still rank a first, so all these profiles are majority consensuses with winner a.

Similarly, let $\mathcal{R}''$ be a Condorcet consensus with winner a. Consider a candidate $c \in C \setminus \{a\}$. At least $\lfloor n/2 \rfloor + 1$ voters in $\mathcal{R}''$ rank a above c. In every profile on every shortest path from $\mathcal{R}''$ to $\mathcal{R}$ these voters still rank a above c. As this holds for every $c \in C \setminus \{a\}$, all these profiles are Condorcet consensuses with winner a. □

5.1 Star $\forall$-Convexity vs $\exists$-Convexity

It is instructive to compare $\exists$-convexity and star $\forall$-convexity. Intuitively, both are obtained from $\forall$-convexity by replacing a $\forall$ quantifier with an $\exists$ quantifier: the former relaxes $\forall$-convexity by considering, for each pair of profiles, *some* shortest path between these profiles (rather than *all* shortest paths), while the latter fixes *some* profile and only considers pairs of profiles that involve this profile, but then imposes constraints on *all* shortest paths between the profiles in each pair. Obviously, star $\exists$-convexity is a relaxation of both of these notions. We will now demonstrate that star $\forall$-convexity and $\exists$-convexity are incomparable: there exists a set that is $\exists$-convex, but not star $\forall$-convex and vice versa. For the first direction, we show that the set of profiles that are single-peaked with respect to a given axis is not star $\forall$-convex. Since, as shown by Obraztsova et al. (2013), this set is $\exists$-convex, we know that it also is star $\exists$-convex.

Proposition 5.2. *For every candidate set C with $|C| \geq 4$, every linear order $<$ over C and every $n > 0$, the set of n-voter profiles that are single-peaked with respect to $<$ is not star $\forall$-convex.*

Proof. Fix a linear order $<$ over C, and let $\mathcal{P}$ be the set of all n-voter profiles that are single-peaked with respect to $<$. Assume without loss of generality that $<$ orders the candidates as $c_1 < \cdots < c_m$.

Suppose for the sake of contradiction that $\mathcal{P}$ is star $\forall$-convex, and let $\mathcal{R}$ be some profile in the $\forall$-kernel of $\mathcal{P}$. We can assume without loss of generality that the first voter in $\mathcal{R}$ ranks some candidate c_j with $j \geq \lfloor m/2 \rfloor + 1$ first. Since $m \geq 4$, this means that he ranks c_3 above c_2. Now, consider the profile $\mathcal{R}'$ where all voters rank the candidates as $c_1 \succ \cdots \succ c_m$; clearly, this profile is single-peaked with respect to $<$. Let $\mathcal{R}''$ be the profile obtained from $\mathcal{R}'$ by swapping c_2 and c_3 in the first vote. Obviously, $\mathcal{R}''$ lies on a shortest path from $\mathcal{R}'$ to $\mathcal{R}$. However, it is not single-peaked with respect to $<$: the only vote that ranks c_1 first and is single-peaked with respect to $<$ is $c_1 \succ \cdots \succ c_m$. □

An example of a subset of $\mathcal{G}_{n,m}$ that is star $\forall$-convex, but not $\exists$-convex is provided by a ball of radius d, i.e., the set $\mathcal{B}(E, d)$ of all profiles that are at distance at most d from a given profile E.

Proposition 5.3. *For any profile $\mathcal{R}$ and any $d \geq 1$ the set $\mathcal{B}(\mathcal{R}, d)$ is star $\forall$-convex. However, if $d = m - 1$, then $\mathcal{B}(\mathcal{R}, d)$ is not $\exists$-convex.*

Proof. To prove our first claim, we observe that $\mathcal{R}$ lies in the $\forall$-kernel of $\mathcal{B}(\mathcal{R}, d)$: for every profile $\mathcal{R}'$ at distance $d' \leq d$ from $\mathcal{R}$ it holds that every profile on every shortest path from $\mathcal{R}'$ to $\mathcal{R}$ is at distance at most d' from $\mathcal{R}$.

To prove the second claim, fix a profile $\mathcal{R} = (R_1, \ldots, R_n)$ over $\{c_1, \ldots, c_m\}$. Assume without loss of generality that R_1 is given by $c_1 \succ \cdots \succ c_m$. Let $\mathcal{R}' = (R'_1, \ldots, R'_n)$ be the profile obtained from $\mathcal{R}$ by pushing c_1 to the bottom of the first vote, i.e., R'_1 is given by $c_2 \succ \cdots \succ c_m \succ c_1$, and $R'_i = R_i$ for $i = 2, \ldots, n$. Similarly, let $\mathcal{R}'' = (R''_1, \ldots, R''_n)$ be the profile obtained from $\mathcal{R}$ by pushing c_m to the top of the first vote, i.e., R''_1 is given by $c_m \succ c_1 \cdots \succ c_{m-1}$, and $R''_i = R_i$ for $i = 2, \ldots, n$. Note that both $\mathcal{R}'$ and $\mathcal{R}''$ are at swap distance $m - 1$ from $\mathcal{R}$.

Now, let $\mathcal{R}^* = (R^*_1, \ldots, R^*_n)$ be the first profile on some shortest path from $\mathcal{R}'$ to $\mathcal{R}''$. Clearly, $R^*_i = R_i$ for each $i = 2, \ldots, m$. Further, the only pairs of candidates that R'_1 and R''_1 disagree on are $(c_1, c_2), \ldots, (c_1, c_{m-1})$ and $(c_m, c_2), \ldots, (c_m, c_{m-1})$. Among these pairs, the only one that is adjacent in R'_1 (and therefore available for swapping) is (c_m, c_{m-1}), so it has to be the case that the first vote in $\mathcal{R}^*$ is $c_2 \succ \cdots \succ c_m \succ c_{m-1} \succ c_1$. But then $\mathcal{R}^*$ is at swap distance m from $\mathcal{R}$, i.e., does not belong to $\mathcal{B}(\mathcal{R}, m-1)$. □

5.2 Star Convexity of Voting Rules

We now move on to the study of voting rules. Our results for star $\forall$-convexity are similar to known results for $\exists$-convexity: while we were able to get positive results for consensus classes, most voting rules turn out not to be star $\forall$-convex. We provide a proof for Plurality, but a similar argument works for many other common voting rules.

Proposition 5.4. *If $m \geq 3$ and $n \geq 9$, the Plurality rule is not star $\forall$-convex.*

Proof. Fix a candidate set C and a candidate $a \in C$, and suppose for the sake of contradiction that the set of all n-voter profiles over C that have a as their Plurality winner has a non-empty $\forall$-kernel. Let $\mathcal{R}$ be some profile in this $\forall$-kernel. Pick two distinct candidates $b, c \in C \setminus \{a\}$. We can assume without loss of generality that the first $\lceil n/2 \rceil$ voters in $\mathcal{R}$ prefer b to c. Now, consider a profile $\mathcal{R}'$ where the first

$\lfloor n/3 \rfloor$ voters rank c first and b second, the next $\lfloor n/3 \rfloor$ voters rank b first, and the remaining voters rank a first. Clearly, a is a Plurality winner at this profile. Let $\mathcal{R}''$ be the profile obtained from $\mathcal{R}'$ by swapping b and c in the first $\lfloor n/3 \rfloor$ votes. Clearly, this profile is on a shortest path from $\mathcal{R}'$ to $\mathcal{R}$. However, in this profile b is ranked first $2\lfloor n/3 \rfloor$ times, while a is ranked first $n - 2\lfloor n/3 \rfloor$ times, and for $n \geq 9$ we have $2\lfloor n/3 \rfloor > n - 2\lfloor n/3 \rfloor$, so a is not a Plurality winner at this profile. □

However, if we further relax our notion of convexity and consider star $\exists$-convexity, we obtain a positive result, by establishing a link between this notion of convexity and feeble monotonicity.

Theorem 5.5. *Every feebly monotone unanimous voting rule $\mathcal{F}$ is star $\exists$-convex. Moreover, for every candidate $a \in C$ every unanimity consensus with winner a is in the $\exists$-kernel of $col_a(\mathcal{G}, \mathcal{F})$.*

Proof. Fix a feebly monotone unanimous voting rule $\mathcal{F}$ and a candidate a. Consider two profiles $\mathcal{R} = (R_1, \ldots, R_n)$ and $\mathcal{R}' = (R'_1, \ldots, R'_n)$ over a candidate set C such that $\mathcal{R}$ is a unanimity consensus with winner $a \in \mathcal{F}(\mathcal{R})$. If all voters in $\mathcal{R}'$ rank a first, then we are done: in any profile $\mathcal{R}''$ that lies on a shortest path from $\mathcal{R}'$ to $\mathcal{R}$ candidate a is ranked first in all votes, so $a \in \mathcal{F}(\mathcal{R}'')$ by unanimity. Otherwise, by feeble monotonicity of $\mathcal{F}$ we can find a voter $i \in [n]$ with $top(R'_i) \neq a$ such that a remains a winner after being swapped with the candidate ranked right above her in R'_i. When we perform this swap, we obtain a profile that lies on a shortest path from $\mathcal{R}'$ to $\mathcal{R}$. We repeat this step until we arrive to a profile $\mathcal{R}^*$ where a is ranked first in every vote. Now we can transform this profile into $\mathcal{R}$ by only swapping the pairs of candidates that contribute to the swap distance between $\mathcal{R}^*$ and $\mathcal{R}$; by unanimity all profiles along the way have a as a winner. □

To appreciate the power of Theorem 5.5, recall that almost all common voting rules, including, e.g., Plurality, Borda, and STV, are feebly monotone (Obraztsova et al. 2013).

It is natural to ask whether the $\exists$-kernel of $col_a(\mathcal{G}, \mathcal{F})$ is equal to the respective unanimity consensus. In general, this is not the case: for many common voting rules the $\exists$-kernel of $col_a(\mathcal{G}, \mathcal{F})$ is considerably larger. In particular, for weakly monotone rules that are *majority-consistent*, i.e., elect the majority winner when one exists, the $\exists$-kernel contains the respective majority consensus.

Proposition 5.6. *Every weakly monotone majority-consistent rule $\mathcal{F}$ is star $\exists$-convex, and for every candidate $a \in C$ every majority consensus with winner a is in the $\exists$-kernel of $col_a(\mathcal{G}, \mathcal{F})$.*

It would be interesting to characterize the $\exists$-kernel of $col_a(\mathcal{G}, \mathcal{F})$ for common voting rules. E.g., it is not clear if for weakly monotone Condorcet-consistent rules the kernel of $col_a(\mathcal{G}, \mathcal{F})$ contains the respective Condorcet consensus.

The following result can be seen as a partial converse of Proposition 5.6: it shows that star $\forall$-convexity, together with other standard axioms, implies weak monotonicity.

Proposition 5.7. *Every star $\forall$-convex anonymous, neutral and unanimous voting rule $\mathcal{F}$ is weakly monotone. Moreover, the $\forall$-kernel of $col_a(\mathcal{G}, \mathcal{F})$ consists of unanimity consensus profiles with winner a.*

6 $\forall$-Convexity

In this section, we explore the possibilities offered by the notion of convexity that is most similar to the classic notion of convexity in geometry, namely, $\forall$-convexity.

6.1 Examples and Counterexamples

While Example 3.3 illustrates that even simple voting rules for two candidates may fail to be $\forall$-convex, the following proposition establishes that there are interesting sets of profiles that have this property.

Proposition 6.1. *Given a candidate set C and a candidate $a \in C$, let $\mathcal{U}_n(a)$ denote the set of all n-voter profiles over C where each voter ranks a first. Then $\mathcal{U}_n(a)$ is $\forall$-convex.*

Proof. Let $\mathcal{R}$ and $\mathcal{R}'$ be two n-voter profiles over C such that both in $\mathcal{R}$ and in $\mathcal{R}'$ all voters rank a first. Consider a voter $i \in [n]$. For every $c \in C \setminus \{a\}$, the votes R_i and R'_i agree on (a, c), so the pair (a, c) does not contribute to $d_{swap}(R_i, R'_i)$. Thus, on any shortest path from $\mathcal{R}$ to $\mathcal{R}'$ in $\mathcal{G}_{n,m}$ candidate a is not swapped with c in the i-th vote. Hence, in every profile on this path voter i ranks a first. As this is true for every voter i, the proof is complete. □

Proposition 6.1 can be viewed as a special case of a more general construction, which is suggested by Nehring and Puppe (2007). Given a candidate set C, $|C| = m$, let $S = \{(c_{j_1}, c_{k_1}), \ldots, (c_{j_s}, c_{k_s})\}$ be a set of s pairs of candidates in C. The sets $\{j_1, k_1\}, \ldots, \{j_s, k_s\}$ are not required to be pairwise disjoint: e.g., we can have $j_1 = j_2$ or $j_3 = k_5$. Consider the set $\mathcal{V}_S$ of all votes where c_{j_ℓ} is ranked above c_{k_ℓ} for all $\ell = 1, \ldots, s$; we will say that $\mathcal{V}_S$ is the set of all votes that are *restricted by* S. This set (viewed as a subset of $\mathcal{G}_{1,m}$) is $\forall$-convex: as we move from one vote in $\mathcal{V}_S$ to another along a shortest path, we only swap pairs of candidates that contribute to the swap distance between the votes, so we never swap c_{j_ℓ} and c_{k_ℓ}, which means that we stay within $\mathcal{V}_S$. Of course, $\mathcal{V}_S$ can be empty; this happens if the constraints imposed by S are incompatible with a linear order over C. This construction can be extended to subsets of $\mathcal{G}_{n,m}$ in a natural way: for each voter $i \in [n]$ we fix a set S^i of pairs of candidates, and consider the set $\mathcal{V}_{(S^1,\ldots,S^n)}$ of all profiles where i's vote is restricted by S^i. Clearly, $\mathcal{V}_{S^1,\ldots,S^n}$ is $\forall$-convex; the argument is the same as in the case of $n = 1$.

This construction produces many examples of $\forall$-convex sets of profiles. In particular, it implies Proposition 6.1: it suffices to set $S^i = \{(a, c) \mid c \in C \setminus \{a\}\}$ for all $i \in [n]$. However, it is easy to see that many other sets of "similar" profiles are not $\forall$-convex. For instance, this holds for the set of all profiles that have a given candidate as their strict majority winner or their Condorcet winner.

Proposition 6.2. *Given a candidate set C and a candidate $a \in C$, let $\mathcal{M}_n(a)$ denote the set of all n-voter profiles over C where more than $n/2$ voters rank a first, and let $\mathcal{C}_n(a)$*

denote the set of all n-voter profiles over C where a is the Condorcet winner. If $|C| \geq 2$ and $n \geq 3$, then for every $a \in C$ the sets $\mathcal{M}_n(a)$ and $\mathcal{C}_n(a)$ are not $\forall$-convex.

Proof. The proof is a straightforward extension of Example 3.3: we can simply 'pad' each vote in $\mathcal{R}$ and $\mathcal{R}'$ with $m-2$ dummy candidates that are ranked in positions $3, \ldots, m$ and are ordered in the same way in each vote. □

Another illustration that $\forall$-convexity is a very demanding notion is provided by the class of profiles single-peaked with respect to a given axis. Indeed, Obraztsova et al. (2013) argue that the set of n-voter m-candidate profiles that are single-peaked with respect to a given axis $<$ is $\exists$-convex for all $n, m \geq 1$. However, by Proposition 5.2 this set is not star $\forall$-convex and hence not $\forall$-convex.

6.2 $\forall$-Convexity of Voting Rules

We will now formalize the intuition that no reasonable voting rule is $\forall$-convex, by characterizing voting rules that satisfy $\forall$-convexity and other well-studied axioms: it turns out that a voting rule is neutral, unanimous and $\forall$-convex if and only if it is a dictatorship, and a voting rule is neutral, anonymous and $\forall$-convex if and only if it outputs the entire set of candidates on every profile.

Recall that a voting rule $\mathcal{F}$ is a *dictatorship* if there exists a voter $i \in [n]$ such that on any profile $\mathcal{R}$ the rule $\mathcal{F}$ outputs the top candidate in R_i.

Theorem 6.3. *A voting rule is neutral, unanimous and $\forall$-convex if and only if it is a dictatorship.*

Proof. It is immediate that a dictatorship is neutral, unanimous and $\forall$-convex; in particular, for every two profiles where a voter i ranks a candidate a first, it holds that in every profile that lies on a shortest path between these two profiles voter i ranks a first.

For the converse direction, fix a voting rule $\mathcal{F}$ that is neutral, unanimous and $\forall$-convex. Consider a candidate $a \in C$. We claim that for each $n > 1$ there exists an n-voter profile where a is nominated, but is not the unique winner under $\mathcal{F}$. Indeed, consider a profile $\mathcal{R}$ where some voters rank a first and the remaining voters rank b first. If a is the unique winner in $\mathcal{R}$ then by neutrality b must be the unique winner in $\pi(\mathcal{R}_{a \leftrightarrow b})$, yet some voters in $\mathcal{R}_{a \leftrightarrow b}$ rank a first.

Let k be the largest integer such that there exists a profile $\mathcal{R}$ where k voters rank a first, yet a is not the unique winner under $\mathcal{F}$; we have argued that $k \geq 1$, and, by unanimity, $k < n$. Fix some such profile $\mathcal{R}^a$ and let $S \subseteq [n]$ be the set of voters in $\mathcal{R}^a$ who rank a first. As $\mathcal{F}(\mathcal{R}^a) \neq \{a\}$, there exists a candidate $b \in \mathcal{F}(\mathcal{R}^a) \setminus \{a\}$. Note that, by neutrality, for any candidate $c \in C$ and any profile where more than k voters rank c first it holds that c is the unique winner under $\mathcal{F}$.

We will now argue that $k = n - 1$. Suppose for the sake of contradiction that $k \leq n-2$. Consider a profile $\mathcal{R}'$ where some voter $i \in [n] \setminus S$ ranks a first, while all other voters rank b first. In $\mathcal{R}'$ there are $n - 1 > k$ voters who rank b first, so $\mathcal{F}(\mathcal{R}') = \{b\}$. Thus, $b \in \mathcal{F}(\mathcal{R}') \cap \mathcal{F}(\mathcal{R}^a)$. Now, consider the profile obtained from $\mathcal{R}^a$ by shifting a to the top of the i-th vote. In this profile, $k + 1$ voters rank a first, so a is the unique winner under $\mathcal{F}$. On the other hand, this profile lies on a shortest path from $\mathcal{R}^a$ to $\mathcal{R}'$, so we obtain a contradiction with $\forall$-convexity of $\mathcal{F}$.

Now, let i be the unique voter in $[n] \setminus S$. Consider an arbitrary profile $\mathcal{R}$ where all voters in S rank a first, and i ranks some candidate $c \in C$ first. We claim that $\mathcal{F}(\mathcal{R}) \subseteq \{a, c\}$. Indeed, suppose that $d \in \mathcal{F}(\mathcal{R})$ for some $d \in C \setminus \{a, c\}$. Then by neutrality we have $d \in \mathcal{F}(\mathcal{R}_{a \leftrightarrow c})$. But there is a shortest path from $\mathcal{R}$ to $\mathcal{R}_{a \leftrightarrow c}$ that goes through a profile where all voters rank a first. The unique winner at that profile is a by unanimity, a contradiction with $\forall$-convexity of $\mathcal{F}$. By the same argument, $\mathcal{F}(\mathcal{R}) \neq \{a, c\}$. Indeed, if $\mathcal{F}(\mathcal{R}) = \{a, c\}$, then by neutrality we have $\mathcal{F}(\mathcal{P}_{a \leftrightarrow c}) = \{a, c\}$, yet there is a shortest path from $\mathcal{R}$ to $\mathcal{R}_{a \leftrightarrow c}$ that goes through a profile where all voters rank a first.

The argument in the previous paragraph implies that in $\mathcal{R}^a$ voter i ranks b first, and $\mathcal{F}(\mathcal{R}^a) = \{b\}$. Hence by neutrality $\mathcal{F}(\mathcal{R}^a_{a \leftrightarrow b}) = \{a\}$. But this means that b is the unique winner in any profile $\mathcal{R}$ where all voters in S rank a first, but i ranks b first: indeed, we already know that $\mathcal{F}(\mathcal{R}) \subseteq \{a, b\}$, and if $a \in \mathcal{F}(\mathcal{R})$, we get a contradiction since there is a shortest path from $\mathcal{R}$ to $\mathcal{R}^a_{a \leftrightarrow b}$ that goes through a profile where all voters rank b first (and where b is the unique winner by unanimity). By neutrality, it holds that in any profile where all voters in S rank some candidate c first, but voter i ranks another candidate d first, d is the unique winner, i.e., voter i is a dictator for all profiles of this type.

It remains to show that i is also a dictator when not all voters in S rank the same candidate first. To this end, consider an arbitrary profile $\mathcal{R}$ and assume that i ranks candidate a first, but $c \in \mathcal{F}(\mathcal{R})$ for some $c \in C \setminus \{a\}$. Let $\mathcal{R}'$ be the profile where i ranks c first, but all voters in S rank a first. As argued in the previous paragraph, we have $\mathcal{F}(\mathcal{R}') = \{c\}$. Consider the profile obtained by shifting a to the top of the i-th vote in $\mathcal{R}'$. In this profile a is the unique winner, but it lies on a shortest path from $\mathcal{R}$ to $\mathcal{R}'$, a contradiction with $\forall$-convexity of $\mathcal{F}$. This completes the proof. □

For our second characterization result we give up unanimity, but require anonymity. The resulting rule is the trivial rule that always outputs the entire set C.

Theorem 6.4. *If the number of voters is at least 3, then a voting rule is neutral, anonymous and $\forall$-convex if and only if it always outputs the entire set of candidates C.*

Theorems 6.3 and 6.4 can be seen as impossibility results, since they show that $\forall$-convex voting rules are bound to be either dictatorships or completely non-decisive. Note, however, that similar conclusions follow from, e.g., requiring that a voting rule is strategy-proof, in the sense that no voter can improve the election outcome from their perspective by casting a vote that differs from their true preference order (Gibbard 1973; Satterthwaite 1975). Arguably, then, our negative results suggest that $\forall$-convexity is, in fact, a very attractive notion. It is well-known that there are restricted domains that admit strategyproof voting rules (Black 1958); is it perhaps the case that there are attractive voting rules that are $\forall$-convex on restricted domains?

7 Conclusions

We have built on the idea of $\exists$-convexity of voting rules, introduced by Obraztsova et al. (2013). We have introduced $\forall$-convexity, which strengthens $\exists$-convexity, star $\exists$-convexity, which relaxes it, and star $\forall$-convexity, which is incomparable to $\exists$-convexity. We have shown that $\forall$-convexity is such a demanding notion that no reasonable voting rule satisfies it, but nonetheless we have found a few natural sets of elections that are $\forall$-convex. On the other hand, we have shown that many natural voting rules are star $\exists$-convex; indeed, we have linked star $\exists$-convexity to the notion of feeble monotonicity, which improves upon the result of Obraztsova et al. (2013) that establishes a connection between feeble monotonicity and connectivity.

There is a number of research directions that stem from our work. First, we would like to build on the results of Section 4 to obtain a complete characterization of anonymous, neutral, unanimous and $\exists$-convex voting rules. Further, it is remarkable that most of the rules considered in Section 4 are not star $\forall$-convex, and it would be interesting to identify further examples of star $\forall$-convex voting rules.

8 Acknowledgments

We are grateful to the reviewers for their helpful and encouraging comments. S. Obraztsova was supported by Tier-1 grant of MOE (Singapore) 2018-T1-001-118, E. Elkind was supported by ERC grant number 639945 (ACCORD), P. Faliszewski was supported by the funds assigned to the AGH University by the Polish Ministry of Science and Higher Education.

References

Arrow, K.; Sen, A.; and Suzumura, K., eds. 2002. *Handbook of Social Choice and Welfare, Volume 1*. Elsevier.

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

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

Gibbard, A. 1973. Manipulation of voting schemes. *Econometrica* 41(4):587–601.

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

Nehring, K., and Puppe, C. 2007. The structure of strategy-proof social choice—part I: General characterization and possibility results on median spaces. *Journal of Economic Theory* 135(1):269–305.

Obraztsova, S.; Elkind, E.; Faliszewski, P.; and Slinko, A. M. 2013. On swap-distance geometry of voting rules. In *Proceedings of AAMAS-13*, 383–390.

Preparata, F. P., and Shamos, M. I. 1985. *Computational geometry: an introduction*. New York, USA: Springer-Verlag.

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

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

Satterthwaite, M. 1975. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory* 10(2):187–217.