General Tiebreaking Schemes for Computational Social Choice

Rupert Freeman
Dept. of Computer Science
Duke University
Durham, NC, USA
rupert@cs.duke.edu

Markus Brill
Dept. of Computer Science
Duke University
Durham, NC, USA
brill@cs.duke.edu

Vincent Conitzer
Dept. of Computer Science
Duke University
Durham, NC, USA
conitzer@cs.duke.edu

ABSTRACT

In (computational) social choice, how ties are broken can affect the axiomatic and computational properties of a voting rule. In this paper, we first consider settings where we may have multiple winners. We formalize the notion of parallel universes tiebreaking with respect to a particular tree that represents the computation of the winners, and show that the specific tree used does not matter if certain conditions hold. We then move on to settings where a single winner must be returned, generally by randomized tiebreaking, and examine some drawbacks of existing approaches. We propose a new class of tiebreaking schemes based on randomly perturbing the vote profile. Finally, we show that one member of this class uniquely satisfies a number of desirable properties.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Multiagent Systems; J.4 [**Computer Applications**]: Social and Behavioral Sciences - Economics

Keywords

social choice; voting rules; tiebreaking; perturbation

1. INTRODUCTION

The theory of voting has attracted the interest of multiagent systems researchers because it provides a natural framework for a group of agents with distinct preferences to reach a joint decision. In this theory, each agent (or *voter*) submits a ranking of the available *alternatives*, and a voting rule determines which alternative wins. Many voting rules have been proposed, and which one is most appropriate generally depends on the setting at hand.

Voting rules are often specified in a way that allows for ties to occur, without making it clear how these ties are to be broken. For example, under the plurality rule, multiple alternatives may have the highest score. As a somewhat different type of example, under the STV rule, where in each round we eliminate the alternative with the lowest plurality score, ties may occur *within* a round.

Appears in: *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015), Bordini, Elkind, Weiss, Yolum (eds.), May 4–8, 2015, Istanbul, Turkey.*

Various reasons may be offered for why it is unnecessary to specify a tiebreaking procedure. One is that a tie is unlikely to occur anyway [17, 14, 12, 15]. This is not always a satisfactory excuse. For one, if an election has few voters, then a tie is generally quite likely—and the importance of an election is not always proportional to the number of voters in it. Another issue is that *strategic* voters will base their decision of how to vote purely on scenarios in which they are *pivotal* (a voter is pivotal if changing her vote can change the outcome). Thus, even if such scenarios are unlikely, how they are handled nevertheless completely determines how strategic voters vote. In computational social choice, we see this issue reflected in the fact that how ties are handled affects the complexity of manipulation [22, 21, 23, 2].

Another reason that may be offered for not specifying the tiebreaking procedure is that one can anyway easily specify a "natural" tiebreaking procedure: choosing an alternative at random, choosing a voter at random who breaks the tie, having a fixed tiebreaking order over alternatives, etc. Indeed, in actual elections, usually a simple tiebreaking procedure is specified, such as the Vice President breaking the tie.

In this paper, we pursue an approach to tiebreaking that generalizes across voting rules. The advantages of such an approach are that ad-hoc decisions are avoided, and that it liberates researchers from the burden of specifying tiebreaking procedures for each voting rule.

Before discussing procedures to *break* ties, we consider the question of *which alternatives should be considered tied* in a given election. For multi-stage rules such as STV and ranked pairs, the answer to this questions depends on how intermediate ties, i.e., ties occurring during the execution of the rule, are handled. A common way to handle such ties is *parallel universes tiebreaking* (PUT). Informally, an alternative is a PUT winner if there is some way of breaking ties in the steps of the protocol that makes this alternative the winner. We provide a general definition of PUT, which to our knowledge has previously only been studied on a rule-by-rule basis, and justify the PUT concept by showing that selecting PUT winners is the natural way to extend a rule from non-tied instances to tied ones.

We then move on to actual tiebreaking schemes, i.e., functions that map every tied election to a probability distribution over the set of alternatives that are tied. After reviewing a number of common tiebreaking schemes and their properties, we propose a class of tiebreaking schemes that is based on randomly perturbing the input of the election by a small amount. We conclude the paper by characterizing a particularly attractive member of this class.

2. PRELIMINARIES

Let A be a finite set of m *alternatives* and let N be a finite set of n voters. A *ranking* of A is a permutation of A. The set of all rankings of A is denoted by $\mathcal{L}(A)$ and contains $m!$ rankings. The preference of voter $v \in N$ is represented by a ranking $r(v) \in \mathcal{L}(A)$, and a *preference profile* R contains the preferences of all voters in N.

A *social choice function (SCF)* f associates with every preference profile R a nonempty set $f(R) \subseteq A$ of alternatives. The following three properties are standard assumptions on SCFs that are satisfied by virtually all common SCFs. First, an SCF is *anonymous* if the set of chosen alternatives does not change when the voters are permuted. Second, an SCF is *neutral* if permuting the alternatives in the individual rankings also permutes the set of chosen alternatives in the same way. Third, an SCF is *homogeneous* if $f(R) = f(kR)$ for all $k \in \mathbb{N}_{>0}$. Here, kR is the profile that contains k copies of each voter in R. The conjunction of anonymity and neutrality is also referred to as *symmetry*.

In order to accommodate SCFs that encounter ties during their execution (and that are not equipped with a way of handling these ties), we also define SCFs that are only defined on a subset of preference profiles. A *partially specified social choice function (p-SCF)* f maps every preference profile R to either a nonempty set $f(R) \subseteq A$ of alternatives, or to $\top$. If $f(R) = \top$, we say that f is *undefined* on R. By definition, every SCF is also a p-SCF.

For a (fully specified) SCF f, a tiebreaking scheme specifies the output of f at profiles R with $|f(R)| > 1$. Formally, a *tiebreaking scheme for* f associates with every preference profile R a probability distribution (or *lottery*) $t(R)$ over A such that $t(R)(a) = 0$ whenever $a \notin f(R)$. Here, $t(R)(a)$ denotes the probability of a under lottery $t(R)$. Neutrality, anonymity, and symmetry for tiebreaking schemes are defined analogously to the corresponding properties of SCFs.

Finally, we introduce the SCFs and p-SCFs that we use in our examples. A *scoring rule* is an SCF that is defined by a sequence $s = (s^k)_{k\geq 1}$, where for each $k \in \mathbb{N}$, $s^k = (s^k_1, \dots, s^k_k) \in \mathbb{R}^k$ is a *score vector* of length k. For a preference profile R on m alternatives, the score vector s^m is used to allocate points to alternatives: each alternative receives a score of s^m_j for each time it is ranked in position j by a voter. The scoring rule then selects all alternatives with maximal total score. Prominent examples of scoring rules are *plurality* ($s^k = (1, 0, \dots, 0)$) and *Borda's rule* ($s^k = (k-1, k-2, \dots, 0)$).

Single Transferable Vote (STV) is a p-SCF that works in several stages. At the first stage, the alternative with the lowest plurality score is eliminated and removed from all the votes. This process is repeated, eliminating one alternative at a time, until a single alternative remains. If at any stage, a tie for the lowest plurality score occurs, STV outputs $\top$.

For an introduction to other common SCFs, some of which we will mention throughout this paper, see [4].

3. TIES IN MULTI-STAGE RULES

When a tie occurs in a fully specified SCF like plurality or Borda's rule, it is clear which alternatives are tied for being the winner: the alternatives in $f(R)$. For other rules, however, this is not so clear. For instance, consider multi-stage rules like STV or ranked pairs, where ties can occur at every individual stage. Formally, as already illustrated for STV, we model such rules as p-SCFs that are only defined on the subset of preference profiles that never result in a tie during the execution of the rule. (For these profiles, the p-SCF always outputs a single alternative.) If a tie occurs at a stage (two alternatives having the lowest plurality score in STV, or two pairs of alternatives having the highest majority margin in ranked pairs), the p-SCF outputs $\top$ because there are multiple ways to deal with this and any subsequent ties. This motivates us to look at *tie-handling procedures*, i.e., rules that tell us which alternatives are actually to be considered tied for a certain profile. A tie-handling procedure extends a p-SCF f to a (fully specified) SCF by associating a nonempty subset of alternatives with each profile R for which $f(R) = \top$.[1]

A natural tie-handling procedure is *parallel universes tiebreaking* (PUT)[2] [25, 6]. Under PUT, the set $f(R)$ consists of all alternatives that are chosen by f for *some* way of breaking any ties that occur along the way (see Section 5 for a formal definition).

Example 1. *Consider STV and the 7-voter preference profile R given by*

3	2	2
c	a	b
b	b	a
a	c	c

At the first stage, a and b have the lowest plurality score and are therefore tied to be eliminated. If a is eliminated, b will be chosen (since it beats c in the next stage). If b is eliminated, a will be chosen. The tie-handling procedure PUT would therefore yield $f(R) = \{a, b\}$.

Other tie-handling procedures are conceivable. For example, the following tie-handling procedure is often considered for STV: whenever there is a tie for the lowest plurality score, simultaneously eliminate *all* tied alternatives.[3] For this way of handling ties, c would be the unique STV winner for the profile in Example 1. Even though this tie-handling procedure seems reasonable at first glance, it has undesirable effects. For instance, it violates the intuition that an alternative that is chosen by an SCF on a tied profile should be in some sense 'close' to being the winner in a profile without ties. To illustrate this, consider again the profile in Example 1. Now, suppose that we replace each of the voters with 1000 new voters. The result of the election remains the same: c is still the unique winner because both a and b are eliminated at the first stage. But if even one of the voters changes her most preferred alternative, there are no longer any ties. Either a or b is uniquely eliminated at the first stage, and the one which is not eliminated will be selected because it beats c at the second stage. Therefore, with only a tiny change to the votes, we move from a profile in which c is selected to one in which c is eliminated by a large margin!

[1]The difference between a tie-handling procedure and a tiebreaking scheme is that the former always outputs a nonempty set of alternatives (possibly a singleton), whereas the latter outputs a probability distribution.

[2]According to our terminology, PUT is *not* a tiebreaking scheme; we keep the name for consistency with the existing literature.

[3]This is a tie-handling procedure rather than a tiebreaking scheme because there might be a complete tie in the last round, in which case all remaining alternatives are in $f(R)$.

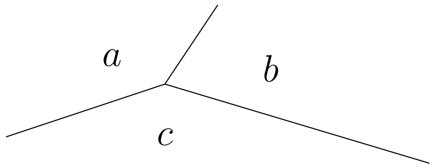

Figure 1: Illustration of a profile space with three separate regions on which a p-SCF f chooses a unique winner. Lines represent profiles at which there is a tie.

In the following section, we formally define properties that prohibit this strange behavior.

4. ANALYTICAL PROPERTIES OF SOCIAL CHOICE FUNCTIONS

We now consider a number of properties of SCFs and p-SCFs that utilize a vector space representation of preference profiles. This allows us to use techniques from real analysis; we are certainly not the first to take this approach [30, 28].

4.1 Profiles as Vectors

Since we only consider anonymous SCFs, we can represent a preference profile R as a vector in $\mathbb{N}_{\geq 0}^{m!}$ as follows. Fix an enumeration $r_1, r_2, \ldots, r_{m!}$ of $\mathcal{L}(A)$ and, for $1 \leq i \leq m!$, let R_i be the number of voters v with $r(v) = r_i$. An SCF then corresponds to a mapping from $\mathbb{N}^{m!}$ to $2^A \setminus \emptyset$ (and a p-SCF corresponds to a mapping from $\mathbb{N}^{m!}$ to $2^A \setminus \emptyset \cup \{\top\}$).

We remark that the domain of a homogeneous SCF or p-SCF can be easily extended to the space $\mathbb{R}_{\geq 0}^{m!}$, and we utilize that fact throughout the paper.[4] (Of course, for most common voting rules, it is quite obvious how to extend them to allow votes with arbitrary real-valued weights.) Let $\mathbf{0} \in \mathbb{R}_{\geq 0}^{m!}$ denote the profile R with $R_r = 0$ for all $r \in \mathcal{L}(A)$. Define the *distance* $d(R, R')$ between two preference profiles R and R' by $d(R, R') = \max_{1 \leq i \leq m!} |R_i - R'_i|$ and the *norm* $|R|$ of a profile R by $|R| = d(R, \mathbf{0}) = \max_{1 \leq i \leq m!} R_i$.

4.2 Continuity and Non-Singularity

Continuity says that if we have a sequence of profiles on which f chooses alternative a, then a is chosen by f at the limit point of that sequence.

Definition 1. *An SCF f is* continuous at R *if for every sequence $(R^k)_{k \in \mathbb{N}}$ with $(R^k) \xrightarrow{k \to \infty} R$ and $a \in f(R^k)$ for all $k \in \mathbb{N}$, $a \in f(R)$. We say f is* continuous *if it is continuous at R for all $R \in \mathbb{R}_{\geq 0}^{m!}$.*

In Figure 1, continuity of f requires that on the tied profiles, f chooses *at least* the alternatives that are chosen in adjacent regions. For example, on the line separating the region where f chooses a and the region where f chooses b, f must choose at least a and b (and possibly more).

Non-singularity requires that if f chooses alternative a at some profile R, then if we randomly choose a profile R' that is 'close' to R, $f(R') = a$ with non-zero probability.

[4]Every homogeneous SCF can be extended to the domain $\mathbb{Q}_{\geq 0}^{m!}$ in a straightforward manner: for $R \in \mathbb{Q}_{\geq 0}^{m!}$, choose $k \in \mathbb{N}$ such that $kR \in \mathbb{N}_{\geq 0}^{m!}$ and define $f(R) = f(kR)$. Homogeneity of f ensures that the definition of $f(R)$ is independent of the factor k. Finally, the gap between $\mathbb{Q}_{\geq 0}^{m!}$ and $\mathbb{R}_{\geq 0}^{m!}$ can be bridged by defining a continuous extension [30].

Definition 2. *An SCF f is* singular at R *if $a \in f(R)$ and there exists an $\varepsilon > 0$ such that the set $\{R' : d(R, R') < \epsilon$ and $a \in f(R')\}$ has (Lebesgue) measure zero. We say f is* non-singular *if there is no $R \in \mathbb{R}_{\geq 0}^{m!}$ such that f is singular at R.*

In Figure 1, non-singularity of f requires that on the tied profiles, f chooses *at most* the alternatives that are chosen in adjacent regions. For example, on the line separating the region where f chooses b and the region where f chooses c, f must not choose a (but f can choose b, c, or both).

We consider continuity and non-singularity very natural properties; indeed, they are satisfied by almost all common voting rules. However, STV appended with the simultaneous-elimination tie-handling procedure from the previous section violates both continuity and non-singularity. In order to see that continuity is violated, consider the profile R from Example 1 and observe that it is easy to find profiles arbitrarily close to R at which a is selected—but $a \notin f(R)$. In order to see that non-singularity is violated, observe that f is singular at R because (1) the only profiles close to R for which c is chosen are those which have a and b tied in the first stage; and (2) the set of profiles with this property is a measure zero subset of the entire space of profiles. By contrast, STV appended with PUT is continuous and non-singular.

4.3 Essential Resoluteness

For every p-SCF f, we can partition the set $\mathbb{R}_{\geq 0}^{m!}$ of all profiles into the set of *non-tied profiles* $\mathcal{NT}(f)$ and the set of *tied profiles* $\mathcal{T}(f)$ as follows. $\mathcal{NT}(f)$ is given by the set of all profiles for which f is defined and outputs a single alternative, i.e., $\mathcal{NT}(f) = \{R \in \mathbb{R}_{\geq 0}^{m!} : f(R) \subseteq A$ and $|f(R)| = 1\}$. Accordingly, the set $\mathcal{T}(f) = \mathbb{R}_{\geq 0}^{m!} \setminus \mathcal{NT}(f)$ contains all profiles for which f outputs either $\top$ or a subset of A that contains more than one alternative.

A p-SCF is *resolute* if $\mathcal{T}(f)$ is empty (or, equivalently, $\mathcal{NT}(f) = \mathbb{R}_{\geq 0}^{m!}$). No symmetric p-SCF can be resolute (resoluteness implies that the p-SCF is defined everywhere, and symmetry implies that $f(\mathbf{0}) = A$), but many common SCFs satisfy the weaker notion of *essential* resoluteness, which requires that *almost all* profiles are non-tied.

Definition 3. *A p-SCF f is* essentially resolute *if $\mathcal{T}(f)$ has measure zero in $\mathbb{R}_{\geq 0}^{m!}$.*

Equivalently, $\mathcal{NT}(f)$ has full measure. Examples of essentially resolute rules include all scoring rules, Kemeny's rule, ranked pairs, and STV.[5] Notable exceptions include tournament solutions such as Copeland and Slater. The reason for this is that those SCFs depend only on pairwise comparisons between alternatives. Therefore, if a given preference profile R is tied, then so is any other profile that gives rise to the same pairwise comparisons (which, provided that none of the pairwise comparisons themselves are tied, will be the case for any R' with $d(R, R') < \epsilon$ for a sufficiently small $\epsilon > 0$—a set of profiles with strictly positive measure).

The p-SCF depicted in Figure 1 is essentially resolute, since ties occur only on the lines, which represent a measure zero subset of profiles.

[5]All scoring rules and Kemeny's rule are also symmetric, homogeneous, continuous, and non-singular. The same holds for the p-SCFs ranked pairs and STV when appended with PUT.

1	1	1	1
a	c	b	c
b	b	a	a
c	a	c	b

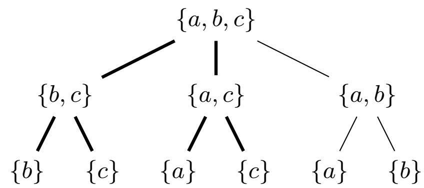

Figure 2: Example of a profile R and a tree T representing f = STV. Nodes are labelled with the set of alternatives that have not yet been eliminated at the corresponding state of computation. Edges of the subtree $T(R)$ are shown in bold. Thus, $PUT_{f,T}(R) = \{b\} \cup \{c\} \cup \{a\} \cup \{c\} = \{a, b, c\}$.

5. PARALLEL UNIVERSES TIEBREAKING

In this section, we formalize the notion of parallel universes tiebreaking (PUT) and provide a justification of this tie-handling procedure using continuity.

5.1 A Formal Definition of PUT

We model the execution of a (multi-stage) p-SCF using a computational tree. Paths from the root to the leaves correspond to different branches of computation of the p-SCF, and each leaf is labelled with an alternative that wins at that leaf. At every non-leaf node x (corresponding to a given execution state), a preference profile specifies a nonempty subset of the children as *feasible* (corresponding to feasible successor states). A tie occurs if more than one child node is feasible. PUT winners can be found by following all feasible paths from the root to the leaves and collecting all labels of the leaves encountered.

For STV, the children of the root of T could correspond to states in which one alternative has been eliminated. For a given profile, the set of feasible children would then be given by those children corresponding to states in which an alternative *with the lowest plurality score* has been eliminated (see Figure 2). We remark that other trees for STV are conceivable, and that, in general, every p-SCF f can be modeled as a tree of depth 1 (such that a leaf labelled with a is feasible for profile R if and only if $a \in f(R)$).

Formally, let T be a tree with root $r(T)$. Denote the set of leaf nodes of T by $L(T)$, and for every non-leaf node x, let $N(x)$ denote the set of children of x. Let w be the function which maps every leaf node ℓ to its label $w(\ell) \in A$. A preference profile R maps every non-leaf node x to a nonempty subset $N_R(x) \subseteq N(x)$ of feasible children. Let $T(R)$ be the minimal subtree of T with the following property: $T(R)$ contains $r(T)$ and all nodes x such that $x \in N_R(x')$ for some node x' in $T(R)$. The set $w(L(T(R)))$ corresponds to the set of PUT winners for profile R.

Given a p-SCF f, we say that a tree T *represents* f if $f(R) = w(L(T(R)))$ for all non-tied profiles $R \in \mathcal{NT}(f)$. We can now formally define PUT winners for a p-SCF f with respect to a tree T.

Definition 4. *Let T be a tree that represents a p-SCF f. For a given profile $R \in \mathbb{R}^{m!}_{\geq 0}$, the* PUT winners of f with respect to T *are defined as the alternatives corresponding to the leaves of $T(R)$, i.e.,*

$$PUT_{f,T}(R) = w(L(T(R))).$$

The definition of $PUT_{f,T}$ depends on the choice of the tree T. We go on to identify conditions (on f and on T) that allow us to do away with this dependence. The following properties of trees are defined analogously to the corresponding properties of SCFs (see Section 4).

Definition 5. *Let T be a tree and R a preference profile.*

- *T is* continuous at R *if for every sequence $(R^k)_{k\in\mathbb{N}}$ with $(R^k) \xrightarrow{k\to\infty} R$ and $\ell \in L(T(R^k))$ for all $k \in \mathbb{N}$, $\ell \in L(T(R))$. We say T is* continuous *if it is continuous at R for all $R \in \mathbb{R}^{m!}_{\geq 0}$.*
- *T is* singular at R *if $\ell \in L(T(R))$ and there exists $\epsilon > 0$ such that $\{R' : d(R, R') < \epsilon$ and $\ell \in L(T(R'))\}$ has measure zero. We say T is* non-singular *if there is no $R \in \mathbb{R}^{m!}_{\geq 0}$ such that T is singular at R.*

Theorem 1. *Let f be an essentially resolute p-SCF and let T_1 and T_2 be two continuous, non-singular trees that both represent f. Then,*

$$PUT_{f,T_1}(R) = PUT_{f,T_2}(R)$$

for all $R \in \mathbb{R}^{m!}_{\geq 0}$.

Proof. Since both T_1 and T_2 represent f, $PUT_{f,T_1}(R) = PUT_{f,T_2}(R)$ for all *non-tied* profiles. Therefore, consider a tied profile $R \in \mathcal{T}(f)$ and let $a \in PUT_{f,T_1}(R) = w(L(T_1(R)))$. Then there is some leaf $\ell \in L(T_1(R))$ with $w(\ell) = a$. And since T_1 is non-singular, there exists a sequence of profiles $(R_i)_{i\in\mathbb{N}}$ such that $(R_i) \to R$ and $l \in L(T_1(R_i))$ for all i. By non-singularity of T_1 and the fact that f is essentially resolute, we can find a sequence of profiles $(R^k)_{k\in\mathbb{N}}$ on which f is resolute such that $(R^k) \xrightarrow{k\to\infty} R$ and $\ell \in L(T_1(R^k))$ for all k. Therefore, $f(R^k) = \{a\}$ for all k. Since T_2 represents f, we have $w(L(T_2(R^k)) = \{a\}$ for all k. By finiteness of T_2, there is a subsequence of $(R^k)_{k\in\mathbb{N}}$, say $(R^{k_j})_{j\in\mathbb{N}}$, and $m \in L(T_2)$ with $w(m) = a$, such that $m \in L(T_2(R^{k_j}))$ for all j. By continuity of T_2, $m \in L(T_2(R))$ and thus $a \in w(L(T_2(R))) = PUT_{f,T_2}(R)$. Therefore, $PUT_{f,T_1}(R) \subseteq PUT_{f,T_2}(R)$. The opposite inclusion holds by a symmetric argument. $\square$

For an essentially resolute p-SCF f, we simply write PUT_f for $PUT_{f,T}$, where T is some tree that satisfies the conditions of Theorem 1. Theorem 1 guarantees that PUT_f is well-defined. Note that the SCF PUT_f agrees with f on all non-tied profiles, and that continuity and non-singularity of a tree T imply continuity and non-singularity of $PUT_{f,T}$ (and thus of PUT_f).

5.2 A Justification of PUT

We conclude this section with an axiomatic justification of PUT. For a given p-SCF f, one can think of the set $\mathcal{NT}(f)$ as the set of profiles where a clear-cut winner exists, whereas profiles in $\mathcal{T}(f)$ give rise to some ambiguity. Given an essentially resolute p-SCF (i.e., a p-SCF for which the set $\mathcal{T}(f)$ has measure zero), a natural approach for reasoning about potential outcomes for profiles in $\mathcal{T}(f)$ is by focusing on outcomes for profiles in $\mathcal{NT}(f)$ and to extrapolate. Assuming that f is continuous and non-singular at all profiles in $\mathcal{NT}(f)$, it turns out that there is a unique way of extending f to $\mathcal{T}(f)$ such that these properties are retained, and that this extension coincides with PUT_f.

THEOREM 2. *Let f be an essentially resolute p-SCF that is continuous and non-singular at all $R \in \mathcal{NT}(f)$. Then, PUT_f is the only SCF that is continuous, non-singular, and agrees with f on $\mathcal{NT}(f)$.*

PROOF SKETCH. By definition, PUT_f agrees with f on $\mathcal{NT}(f)$. Moreover, PUT_f is continuous and non-singular.

Now let f' be an SCF that is continuous, non-singular, and agrees with f on $\mathcal{NT}(f)$. Non-singularity of f' implies $f'(R) \subseteq PUT_f(R)$ for all $R \in \mathbb{R}_{\geq 0}^{m!}$, and continuity of f' implies $PUT_f(R) \subseteq f'(R)$ for all $R \in \mathbb{R}_{\geq 0}^{m!}$. $\square$

6. EXISTING TIEBREAKING SCHEMES

Whereas the previous sections have been concerned with the question of *which alternatives are tied* at certain profiles, we now focus on methods to *break ties*. This section briefly surveys some existing tiebreaking schemes, and prepares the ground for the tiebreaking schemes that will be introduced in Section 7.[6]

6.1 Fixed-Order Tiebreaking

A very common way to break ties works by first fixing a tiebreaking ordering, either by some external authority or by copying the preferences of a distinguished voter. Then, this ordering is used to break any ties that occur, either during the execution of the SCF, or after computing the set of PUT winners. These schemes violate neutrality (in the case where the ordering is externally given) or anonymity (in the case where the preferences of a given voter are used). Since we consider symmetry (i.e., the conjunction of neutrality and anonymity) a basic fairness axiom for tiebreaking schemes, we do not consider fixed-order tiebreaking schemes any further in this paper.

6.2 PUT-Based Tiebreaking

Symmetric variants of fixed-order tiebreaking schemes can be constructed by randomization. The following two schemes both rely on the notion of PUT winners (which are defined independently of the tree under the conditions of Theorem 1). In *PUT random order* tiebreaking, a ranking $r \in \mathcal{L}(A)$ is selected uniformly at random and the winner is defined as the highest ranked alternative in r that is among the PUT winners. This is equivalent to choosing one alternative from the set of PUT winners uniformly at random. *PUT random vote* tiebreaking works similarly, except that the ranking r is selected uniformly at random from the set of all votes that have been cast. Equivalently, each PUT winner is assigned a probability proportional to the plurality score they would get in an election among only the PUT winners. In the example in Figure 2, PUT random order tiebreaking gives the lottery $t(R) = \frac{1}{3}a + \frac{1}{3}b + \frac{1}{3}c$, and PUT random vote tiebreaking gives $t(R) = \frac{1}{4}a + \frac{1}{4}b + \frac{1}{2}c$.

6.3 Tree-Based Tiebreaking

Multi-stage SCFs that are modeled as trees (see Section 5.1) allow for a different approach. Namely, one can define tiebreaking schemes by choosing only a single path from the root to the leaves of the tree. The most obvious way to do this would be to choose a leaf of $T(R)$ uniformly at random, and declare the label of that leaf the winner. Call this method *random leaf* tiebreaking. Another option is to find a path by choosing, at each non-leaf node x, a child node in $N(x)$ uniformly at random among all feasible children. We call this *random child* tiebreaking. For the example specified in Figure 2, random leaf and random child tiebreaking both give the lottery $t(R) = \frac{1}{4}a + \frac{1}{4}b + \frac{1}{2}c$ (in general, these schemes may produce different lotteries).

There are also tree-based schemes that are similar to the PUT-based schemes described in the previous paragraph, in that a ranking r is selected (uniformly at random from either the set of all rankings or the set of all votes) and used to break any ties that occur in the execution of the tree.[7] Tree-based tiebreaking schemes are highly sensitive to the structure of the tree T chosen to represent f. For example, these schemes reduce to their corresponding PUT-based schemes in the case where f is represented by the trivial tree of depth 1.

6.4 Discussion

It is straightforward to show that the tiebreaking schemes considered in this section may yield different results. A major drawback of PUT-based tiebreaking schemes is that the set of PUT winners may be NP-hard to compute. In particular, this is the case for STV [7], ranked pairs [5], and a number of other SCFs [18]. Therefore, we cannot efficiently sample from the probability distributions defined by PUT-based schemes.

Random child tiebreaking, on the other hand, does not suffer from this problem, as long as it is easy to follow a random path down $T(R)$, which tends to be the case for natural trees. There is also a conceptual advantage of tree-based schemes. Intuitively, under rules such as STV, on a tied profile, some alternatives have more routes to victory than others, and it feels natural that they should have greater probability of having the tie broken in their favor (see, e.g., [26]). Tree-based tiebreaking achieves this, in the sense that an alternative that would win under more tiebreaking paths (feasible paths in $T(R)$) has an advantage. The downside is that (unlike PUT) the alternatives' winning probabilities depend on the precise tree chosen—and if a trivial (depth 1) tree is chosen, this will obscure the fact that an alternative has multiple ways to win (and therefore not take it into account).

Is there a way to get the best of both worlds? In what follows, we take a different approach that is independent of the tree: we randomly *perturb* the original profile and compute the winner for the perturbed profile. As we will see, there are many ways in which we can perturb profiles, with different axiomatic and computational properties.

[6]While most tiebreaking schemes considered in this section are formulated in terms of multi-stage rules and/or trees, we remark that all the schemes are also applicable to fully specified (single-stage) SCFs such as Borda's rule: for the schemes in Sections 6.1 and 6.2, $f(R)$ can take over the role of the set of PUT winners; and the schemes in Section 6.3 can be applied to a trivial (depth 1) tree representation of f.

[7]Unfortunately, while these schemes are well suited to some multi-stage rules such as STV, there is no guarantee that intermediate ties are as simple as ties between alternatives. For example, in the ranked pairs rule, intermediate ties are between pairs of alternatives. For this case, ties could be broken lexicographically (where the lexicographic order is given by r; see [31]), but in general the intermediate ties could take any form.

7. RANDOM PERTURBATION TIE-BREAKING

In this section, we propose the class of *random perturbation* tiebreaking schemes. The idea is that given a tied profile, we add a small random perturbation to the profile which will break the tie with high probability. This technique is inspired by (symbolic) perturbation methods that are used in algorithm design to handle degenerate input instances,[8] but its application to voting is, to our knowledge, new.

To motivate perturbation tiebreaking in the context of voting, suppose that we know that the preferences of the voters are going to gradually change over time, but we have no idea how. Then it makes sense to choose from the set of tied outcomes in a way that reflects the distribution of voters' preferences as they change slightly.

7.1 Definition of Perturbation Tiebreaking

Random perturbation tiebreaking works as follows. Let f be an essentially resolute SCF and consider a profile $R \in \mathcal{T}(f)$. We choose a *perturbation* $P \in \mathbb{R}^{m!}$ randomly according to some probability distribution D (depending on R) and add P to the profile R. The corresponding tiebreaking scheme t_D is then given by

$$t_D(R)(a) = Pr_{P\sim D}(f(R+P) = \{a\}).$$

Before analyzing appropriate candidates for the choice of the probability distribution D, we consider general properties of random perturbation schemes. In order to be a well-defined tiebreaking scheme, t_D should be such that $t_D(R)(a) = 0$ for all $a \notin f(R)$. Intuitively, this can be guaranteed by using a perturbation P with sufficiently small norm $|P|$. Thus, after choosing P according to D, we will multiply all entries of P by a sufficiently small number (call this operation *scaling*). A natural worry is that the choice of the multiplier influences the resulting probabilities of the perturbation scheme. We go on to show that, under mild conditions, this is not the case: for each profile R, there exists $\epsilon_R > 0$ such that t_D is independent of the exact scaling of P, as long as the scaling guarantees that $|P| < \epsilon_R$.

The mild conditions we need to prove this scaling independence result are continuity (see Section 4) and *finite local consistency (FLC)* [28]. FLC is not very restrictive: with the exception of Dodgson's rule (which is not homogeneous and therefore not extendable to $\mathbb{R}^{m!}_{\geq 0}$ anyway), every common SCF satisfies this property. For $S \in \mathbb{R}^n$, let $\overline{S}$ denote the closure of S.[9]

Definition 6. *Let $S \subseteq \mathbb{R}^{m!}_{\geq 0}$ be a set of profiles. An SCF f is* locally consistent *on S if for any $R_1, R_2 \in S$ with $f(R_1) = f(R_2)$, we have $R_1 + R_2 \in S$ and $f(R_1 + R_2) = f(R_1) = f(R_2)$. f is* finitely locally consistent *(FLC) if there exists a finite set of subsets $\{S_1, \ldots, S_t\}$ of $\mathbb{R}^{m!}_{\geq 0}$ such that f is locally consistent on S_i for all $1 \leq i \leq t$ and $\overline{\bigcup_i S_i} = \mathbb{R}^{m!}_{\geq 0}$.*

We will now show that the scaling parameter does not influence the tiebreaking scheme, provided it is sufficiently small (Theorem 3). Some intuition for the proof is provided in Figure 3. Finite local consistency allows us to find a finite set of hyperplanes separating regions with different winners.

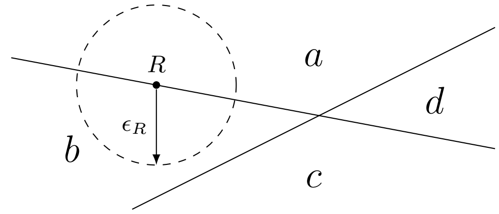

Figure 3: Illustration of the proof of Theorem 3. Lines represent separating hyperplanes and the chosen alternative within each region is specified by a, b, or c. By continuity, $f(R) = \{a, b\}$. We can see that the choice of ϵ_R satisfies the conditions of Theorem 3.

Theorem 3. *Let f be an essentially resolute SCF satisfying continuity and FLC. For every $R \in \mathbb{R}^{m!}_{\geq 0}$, there exists an $\epsilon_R > 0$ such that the following events have probability 1 for all $P \in \mathbb{R}^{m!}_{\geq 0}$ with $0 < |P| < \epsilon_R$ (irrespective of the choice of D):*

(i) $f(R + P) = f(R + \alpha P)$ for all $0 < \alpha \leq 1$, and

(ii) $f(R + P) \subseteq f(R)$.

Proof. Since f satisfies FLC and essential resoluteness, there exist finitely many regions $\{S_1, ..., S_k\}$ such that f is locally consistent and resolute within each region, and these regions are convex (this is implied by local consistency). For each pair of regions (S_i, S_j), we can find a separating hyperplane $(H_{i,j})$ by the hyperplane separation theorem. Now let $d_{i,j}$ be the distance from R to $H_{i,j}$, where the distance $d(R, H)$ from a vector R to a hyperplane H is defined as $d(R, H) = \min_{h\in H} d(R, h)$. Let $\epsilon_R = \min_{d_{i,j}\neq 0} d_{i,j}/2$. First we show that event (i) has probability 1. Let $|P| < \epsilon_R$ and suppose $f(R + P) = \{a\}$ but $f(R + \alpha P) = \{b\}$ for some $0 < \alpha \leq 1$ (the case where $|f(R+P)| > 1$ or $f(R+\alpha P)| > 1$ occurs with probability 0). Either $R + P$ and $R + \alpha P$ are in different regions, and there must be a separating hyperplane passing between them, or one (or both) of $R + P$ and $R + \alpha P$ lie on a separating hyperplane. In the first case, we have contradicted our choice of ϵ_R. The second case occurs with probability 0 since each of the finite number of hyperplanes have measure zero.

Now we consider event (ii). Let $S \subseteq \mathbb{R}^{m!}_{\geq 0}$ be the set of profiles P for which (i) holds. It suffices to show that $f(R+P) \subseteq f(R)$ for all $P \in S$. Suppose there exists $P \in S$ for which this is not the case, then we could construct a sequence $(R^k)_{k\in\mathbb{N}}$ defined by $R_k = R + \frac{1}{k}P$, so that $(R^k) \xrightarrow{k\to\infty} R$, but there exists an $a \in f(R + P) = f(R_k)$ with $a \notin f(R)$. □

This gives rise to the following general approach. We draw the perturbation P from any distribution D, and then multiply the entries of P by a sufficiently small constant so that $|P| < \epsilon_R$. We will say that any tiebreaking scheme of this form is a *perturbation tiebreaking* scheme.

To illustrate this, we exhibit a suitable choice of ϵ_R for plurality. For simplicity, assume that R is an integral preference profile, i.e., $R \in \mathbb{N}^{m!}_{\geq 0}$.[10]

[8]See [16] and the references therein. In particular, perturbation methods have been extensively studied in mathematical programming [11, 13, 27, 1] and computational geometry [29, 8, 9, 20, 24].

[9]Our definition of FLC diverges slightly from the definition in [28] in order to accommodate set-valued SCFs.

[10]While it is possible to derive ϵ_R for non-integral profiles, it is not possible to write it down in such a clean manner.

Proposition 1. *Let $R \in \mathbb{N}_{\geq 0}^{m!}$. For plurality, $\epsilon_R = \frac{1}{2m!}$ satisfies the conditions of Theorem 3.*

Proof. Let $a \in f(R)$, $b \notin f(R)$, and let P be a profile with $|P| < \epsilon_R$. To show that ϵ_R satisfies condition (ii) of Theorem 3, we must show that $b \notin f(R+P)$. To this end, note that the difference in the plurality score of a and that of b is at least one, since R is integral. And the change in plurality score of any alternative as a result of adding P to R is at most $m!|P| < m!\frac{1}{2m!} = \frac{1}{2}$. Therefore, it is impossible for b to have a plurality score as high as that of a in the profile $R+P$. Finally, to see that ϵ_R satisfies condition (i) of Theorem 3, let $0 < \alpha < 1$ and suppose that $a \in f(R+P)$. Then the plurality score of a according to P, and therefore also according to αP, is at least as high as that of any other alternative in $f(R)$. Therefore $a \in f(R+\alpha P)$. $\square$

It is straightforward to derive formulas for ϵ_R for a number of other common SCFs, e.g., $\epsilon_R = \frac{1}{2(m-1)m!}$ works for Borda's rule and $\epsilon_R = \frac{1}{2m!}$ for both STV and ranked pairs.[11]

7.2 Perturbation Distributions

We now turn to the question from which probability distribution D the perturbation should be drawn. Let $e_r \in \mathbb{N}_{\geq 0}^{m!}$ denote the profile that consists of a single vote r (i.e., $[e_r]_r = 1$ and $[e_r]_s = 0$ for all $s \in \mathcal{L}(A) \setminus \{r\}$). The following are three natural choices for the distribution D.

- *Uniform Independent Ranking Perturbation (UIRP):* For every ranking $r \in \mathcal{L}(A)$, choose p_r from $[0,1]$ uniformly at random and let $P = \sum_{r \in \mathcal{L}(A)} p_r e_r$.
- *Uniform Independent Voter Perturbation (UIVP):* For every voter v, choose p_v from $[0,1]$ uniformly at random and let $P = \sum_v p_v e_{r(v)}$.
- *Uniform Proportional Perturbation (UPP):* For every ranking $r \in \mathcal{L}(A)$, choose p_r from $[0,1]$ uniformly at random and let $P = \sum_{r \in \mathcal{L}(A)} R_r p_r e_r$.

UIRP perturbs each entry of a profile R according to the same distribution, regardless of the size of that entry. UIVP and UPP, on the other hand, perturb each entry according to a distribution that depends on the number of voters casting that entry as their vote. UIVP does so by perturbing the weight of each *voter* independently, while UPP draws a perturbation for each entry of R and multiplies it by the number of voters with that ranking as their vote.

Example 2. *Consider again the preference profile R in Figure 2 and let f be Borda's rule. Observe that $f(R) = \{a,b,c\}$. Since $m = 3$, we let $\epsilon = \frac{1}{2 \cdot 2 \cdot 3!} = \frac{1}{24}$ (see the paragraph after Proposition 1). Breaking the tie with UIRP results in the lottery $\frac{1}{3}a + \frac{1}{3}b + \frac{1}{3}c$, and UIVP and UPP both yield $\frac{29}{96}a + \frac{29}{96}b + \frac{19}{48}c$.*

It is not trivial to derive these numbers; we used Mathematica to compute these probabilities. However, deriving the exact probabilities is not actually necessary in order to execute the tiebreaking scheme—sampling from the distribution is sufficient.

[11]The value ϵ_R is not necessarily independent of the number of voters, even for integral profiles. For instance, ϵ_R does depend on n for the scoring rule given by $(\sqrt{2}, 1, 0, \ldots, 0)$.

It is not a coincidence that UIRP gives the same lottery as PUT random order tiebreaking in this example. It can be easily shown that the two schemes coincide for all scoring rules (however, they differ for other SCFs). Notice that both UIVP and UPP have $t(R)(c) = \frac{19}{48} > \frac{1}{3}$. The reason for this is that c is a polarizing alternative among these four voters; two voters prefer c to both a and b and two voters place c at the bottom of their ranking. So, when we perturb only rankings that appear as votes, c has both a higher probability of doing very well, and of doing very badly, after perturbation. This explains the high probability of c winning under these schemes (and also the high probability of c doing the worst, which we are not interested in here).

One basic property that we would want a tiebreaking scheme to satisfy is homogeneity, stating that—just like the corresponding property for SCFs—the result of the perturbation does not change if the profile is multiplied by a constant. Formally, tiebreaking scheme t_D is *homogeneous* if it satisfies $t_D(R) = t_D(kR)$ for all $R \in \mathbb{N}_{\geq 0}^{m!}$ and all $k \in \mathbb{N}_{>0}$.

Unfortunately, UIVP fails homogeneity. This is perhaps not particularly surprising; given that, unlike the others, the definition of UIVP relies on an integral number of voters, we would not expect to be able to sensibly extend UIVP to non-integral profiles. If we double the profile in Example 2, the lottery returned by UIVP becomes $\frac{48739}{161280}a + \frac{48739}{161280}b + \frac{31901}{80640}c$. Intuitively, as we add more voters, by the Central Limit Theorem, P begins to look more and more like its entries were drawn from a normal distribution, since each entry is the sum of i.i.d. random variables.[12] We can avoid this problem by choosing p_v from a normal distribution to begin with. This idea gives rise to the following distribution.

- *Normal Independent Voter Perturbation (NIVP):* For every voter v, draw p_v according to a normal distribution with mean 0 and standard deviation 1. Define $P = \sum_{v \in N} p_v e_{r(v)}$.

Recall that we will scale the profile P to be sufficiently small; since, additionally, we only perturb votes that have been cast, there will not be any negative entries in the perturbed profile $R+P$. Also, the choice of mean and standard deviation is arbitrary; the same lottery over alternatives would result from any other choice. Moreover, we note that NIVP is equivalent to drawing P_r from a normal distribution with mean 0 and standard deviation $\sqrt{R_r}$. This latter definition does not require the profile to be integral, allowing us to naturally extend NIVP to profiles with real-valued weights. The remainder of this section will be devoted to characterizing NIVP and illustrating its advantages over the other perturbation tiebreaking schemes presented.

7.3 Characterization of NIVP

We identify four properties of perturbation distributions.

- **Property 1:** $Pr(P_r = 0 | R_r = 0) = 1$ for all $r \in \mathcal{L}(\mathcal{A})$. This property states that the perturbation gives positive weight only to those rankings that have been cast by at least one voter. An appealing consequence of this property is that the perturbed profile preserves structural characteristics of the original profile, such as

[12]Note that we do not need the original distribution to be uniform for this argument to work. For example, perturbing by simply removing each vote with some probability [19] will also work.

membership in restricted domains (e.g., single-peaked preferences) or the existence of clones.

- **Property 2:** For all R, R_1, R_2 with $R = R_1 + R_2$, we have

$$Pr(P|R) = \int_{P_1, P_2 : P_1 + P_2 = P} Pr(P_1|R_1) Pr(P_2|R_2).$$

This property states the distribution over perturbations does not change if we first partition the profile and then apply the perturbation scheme to each part separately.

- **Property 3:** The resulting tiebreaking scheme t_D is homogeneous.

- **Property 4:** For all profiles R and rankings $r \in \mathcal{L}(\mathcal{A})$, the marginal distribution for P_r has finite variance.

It is easy to show that NIVP satisfies all four of these properties. Interestingly, it can also be shown that no other symmetric perturbation tiebreaking scheme achieves this.

THEOREM 4. *NIVP is the only symmetric perturbation tiebreaking scheme that satisfies Properties 1, 2, 3, and 4.*

PROOF SKETCH. Let $R \in \mathbb{N}_{\geq 0}^{m!}$ be a profile consisting of voters $v_1, v_2, \ldots, v_n$. For this proof only, for the sake of brevity, we will also use v_i to denote voter i's ranking (which is usually denoted by $r(v_i)$). By Property 2, we can write

$$Pr(P|R) = \int_{\substack{P_{v_1}, P_{R \setminus \{v_1\}}: \\ P_{v_1} + P_{R \setminus \{v_1\}} = P}} Pr(P_{v_1}|v_1) Pr(P_{R \setminus \{v_1\}}|R \setminus \{v_1\}).$$

Continuing along the same lines on $Pr(P_{R \setminus \{v_1\}}|R \setminus \{v_1\})$, we see that $Pr(P|R)$ can be expressed entirely in terms of $Pr(P_{v_i}|v_i)$. And, by Property 1, the perturbation vector P_{v_i} contains $m! - 1$ zeroes (for all rankings not equal to v_i). So we can generate P_{v_i} by drawing a single number. By symmetry, we require that P_{v_i} is drawn from the same distribution for all votes. Lastly, homogeneity implies that P_{v_i} is drawn from a normal distribution: Suppose that this were not the case and consider a profile R, with $R_r > 0$ for some ranking r, for which $|f(R)| > 1$. Consider multiplying R by a large constant k. By the Central Limit Theorem, the distribution of P_r will converge to a normal distribution for large k since P_r is the sum of many individual draws from a distribution with finite variance (from Property 4). Thus, if P_{v_i} is not drawn from a normal distribution, homogeneity is violated. And we have already noted that any perturbation tiebreaking scheme that draws each P_{v_i} from the same normal distribution is equivalent to NIVP. $\square$

We note that the properties are independent. UIVP satisfies Properties 1, 2, and 4 but not Property 3. UPP satisfies Properties 1, 3, and 4 but not Property 2. For a rule which satisfies Properties 2, 3, and 4 but violates Property 1, suppose that for every voter we draw a vector P_v as follows: For each $1 \leq i \leq m!$, draw $(P_v)_i$ from a normal distribution with mean 0 and standard deviation 1. We let $P = \sum_{v \in V} P_v$. In the same way as for NIVP, we can naturally extend this scheme to fractional profiles by simply drawing each entry in P from a normal distribution with mean 0 and standard deviation $\sqrt{n}$. Lastly, we note that we can satisfy Properties 1, 2, and 3 while violating Property 4 by defining P in the same way as for NIVP, but drawing p_v from a stable distribution with infinite variance, for example the Cauchy distribution.

7.4 Advantages of NIVP

We have already mentioned that perturbation tiebreaking schemes satisfying Property 1 respect structural characteristics such as membership in a restricted domain or the existence of clone sets. One reason why this is desirable is that many standard computational problems get much harder when leaving a well-structured domain such as single-peaked preferences [10, 3]. With respect to clone sets, i.e., sets of alternatives that are ranked consecutively by each voter, a desirable property of tiebreaking schemes is *independence of clones* [25]. This property requires that, whenever clones of an alternative a are introduced, the only effect this has on the tiebreaking probabilities is that the probability mass of alternative a is divided among its clones (and all other alternatives still have the same probability as in the original profile). It can be shown that NIVP satisfies independence of clones, while UIRP and UPP do not (even though UPP satisfies Property 1).

Another advantage of NIVP is that the distribution is easy to sample from. All we have to do is draw one random number for every ranking that is cast as a vote (see Section 7.2). By contrast, the straightforward implementation of UIRP tiebreaking requires us to generate $m!$ random numbers to generate a perturbation vector P.

8. CONCLUSION

It is well understood that how the ties of a voting rule are broken can significantly impact how voters vote strategically, the complexity of standard problems in computational social choice, whether the rule satisfies certain axiomatic properties, etc. In spite of this, often an ad-hoc tiebreaking scheme is used, or none is specified at all.

In this paper, we addressed these issues by investigating tiebreaking schemes that can be generally applied. We first analyzed how intermediate ties in multi-stage rules can be handled. In particular, we investigated the notion of parallel universes tiebreaking (PUT), which is a generally applicable tie-handling procedure but has, to our knowledge, not previously been defined in a general manner. We define it in a way that is dependent on a tree used to compute the outcome of the voting rule, and show that in fact any correct tree will lead to the same definition provided certain properties hold.

We then moved on to methods that actually break ties. We reviewed some standard tiebreaking schemes and pointed out some of their drawbacks, such as high computational complexity and dependence on the tree used to compute the outcome. In the remainder, we focused on schemes that break ties by perturbing the profile in various ways. Most notably, we introduced the NIVP perturbation scheme, which can be efficiently executed and is the unique perturbation scheme satisfying several desirable properties.

In this paper, we have paid close attention to axiomatic properties of how we break ties, as well as the computational complexity of executing the tiebreaking scheme. Future research could be devoted to studying other properties of our tiebreaking schemes, such as their effect on strategic voting and on the complexity of problems such as manipulation, control, bribery, etc.

Acknowledgments

This work was supported by NSF and ARO under grants CCF-1101659, IIS-0953756, CCF-1337215, W911NF-12-1-0550, and W911NF-11-1-0332, and by a Feodor Lynen research fellowship of the Alexander von Humboldt Foundation.

REFERENCES

[1] A. Archer, R. Rajagopalan, and D. Shmoys. Lagrangian relaxation for the k-median problem: New insights and continuity properties. In G. Di Battista and U. Zwick, editors, *Algorithms - ESA 2003*, volume 2832 of *Lecture Notes in Computer Science*, pages 31–42. Springer Berlin Heidelberg, 2003.

[2] H. Aziz, S. Gaspers, N. Mattei, N. Narodytska, and T. Walsh. Ties matter: Complexity of manipulation when tie-breaking with a random vote. In *Proc. of 27th AAAI Conference*, pages 74–80. AAAI Press, 2013.

[3] F. Brandt, M. Brill, E. Hemaspaandra, and L. Hemaspaandra. Bypassing combinatorial protections: Polynomial-time algorithms for single-peaked electorates. In *Proc. of 24th AAAI Conference*, pages 715–722. AAAI Press, 2010.

[4] F. Brandt, V. Conitzer, and U. Endriss. Computational social choice. In G. Weiß, editor, *Multiagent Systems*, chapter 6, pages 213–283. MIT Press, 2nd edition, 2013.

[5] M. Brill and F. Fischer. The price of neutrality for the ranked pairs method. In *Proc. of 26th AAAI Conference*, pages 1299–1305. AAAI Press, 2012.

[6] V. Conitzer, M. Rognlie, and L. Xia. Preference functions that score rankings and maximum likelihood estimation. In *Proc. of 21st IJCAI*, pages 109–115. AAAI Press, 2009.

[7] V. Conitzer, M. Rognlie, and L. Xia. Preference functions that score rankings and maximum likelihood estimation. In *Proc. of 21st IJCAI*, pages 109–115. AAAI Press, 2009.

[8] I. Emiris and J. Canny. A general approach to removing degeneracies. *SIAM Journal on Computing*, 24:650–664, 1995.

[9] I. Z. Emiris, J. F. Canny, and R. Seidel. Efficient perturbations for handling geometric degeneracies. *Algorithmica*, 19(1-2):219–242, 1997.

[10] P. Faliszewski, E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. The shield that never was: Societies with single-peaked preferences are more open to manipulation and control. In *Proc. of 12th TARK Conference*, pages 118–127. ACM Press, 2009.

[11] R. Fletcher. Resolving degeneracy in quadratic programming. *Annals of Operations Research*, 46-47(2):307–334, 1993.

[12] W. Gehrlein and D. Lepelley. On some limitations of the median voting rule. *Public Choice*, 117(1-2):177–190, 2003.

[13] K. George and M. Osborne. On degeneracy in linear programming and related problems. *Annals of Operations Research*, 46-47(2):343–359, 1993.

[14] R. Gillett. Collective indecision. *Behavioral Science*, 22(6):383–390, 1977.

[15] R. Gillett. The comparative likelihood of an equivocal outcome under the plurality, Condorcet, and Borda voting procedures. *Public Choice*, 35(4):483–491, 1980.

[16] T. Jacobs. Analytical aspects of tie breaking. *Theoretical Computer Science*, 465:1–9, 2012.

[17] T. Marchant. The probability of ties with scoring methods: Some results. *Social Choice and Welfare*, 18:709–735, 2001.

[18] N. Mattei, N. Narodytska, and T. Walsh. How hard is it to control an election by breaking ties? In *Proc. of 21st ECAI*, pages 1067–1068, 2014.

[19] M. Messner and M. K. Polborn. Strong and coalition-proof political equilibria under plurality and runoff rule. *International Journal of Game Theory*, 35(2):287–314, 2007.

[20] M. Neuhauser. Symbolic perturbation and special arithmetics for controlled handling of geometric degeneracies. In *Proc. of 5th International Conference in Central Europe on Computer Graphics and Visualization*, pages 386–395, 1997.

[21] S. Obraztsova and E. Elkind. On the complexity of voting manipulation under randomized tie-breaking. In *Proc. of 22nd IJCAI*, pages 319–324. AAAI Press, 2011.

[22] S. Obraztsova, E. Elkind, and N. Hazon. Ties matter: Complexity of voting manipulation revisited. In *Proc. of 10th AAMAS Conference*, pages 71–78. IFAAMAS, 2011.

[23] S. Obraztsova, Y. Zick, and E. Elkind. On manipulation in multi-winner elections based on scoring rules. In *Proc. of 12th AAMAS Conference*, pages 359–366. IFAAMAS, 2013.

[24] R. Seidel. The nature and meaning of perturbations in geometric computing. In P. Enjalbert, E. Mayr, and K. Wagner, editors, *STACS 94*, volume 775 of *Lecture Notes in Computer Science*, pages 1–17. Springer Berlin Heidelberg, 1994.

[25] T. N. Tideman. Independence of clones as a criterion for voting rules. *Social Choice and Welfare*, 4(3):185–206, 1987.

[26] B. Wichmann. Tie breaking in STV. *Voting Matters*, (19):1–4, November 2004.

[27] P. Wolfe. A technique for resolving degeneracy in linear programming. *Journal of the Society for Industrial and Applied Mathematics*, 11(2):205–211, 1963.

[28] L. Xia and V. Conitzer. Finite local consistency characterizes generalized scoring rules. In *Proc. of 21st IJCAI*, pages 336–341. AAAI Press, 2009.

[29] C.-K. Yap. Symbolic treatment of geometric degeneracies. *Journal of Symbolic Computation*, 10:349–370, 1990.

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

[31] T. M. Zavist and T. N. Tideman. Complete independence of clones in the ranked pairs rule. *Social Choice and Welfare*, 6(2):167–173, 1989.