On the Axiomatic Characterization of Runoff Voting Rules

Rupert Freeman and **Markus Brill** and **Vincent Conitzer**
Department of Computer Science
Duke University
Durham, NC 27708, USA
{rupert,brill,conitzer}@cs.duke.edu

Abstract

Runoff voting rules such as single transferable vote (STV) and Baldwin's rule are of particular interest in computational social choice due to their recursive nature and hardness of manipulation, as well as in (human) practice because they are relatively easy to understand. However, they are not known for their compliance with desirable axiomatic properties, which we attempt to rectify here. We characterize runoff rules that are based on scoring rules using two axioms: a weakening of local independence of irrelevant alternatives and a variant of population-consistency. We then show, as our main technical result, that STV is the only runoff scoring rule satisfying an independence-of-clones property. Furthermore, we provide axiomatizations of Baldwin's rule and Coombs' rule.

1 Introduction

In the general theory of voting, voters each rank a set of alternatives, and based on these input rankings a voting rule determines an aggregate ranking of the alternatives (or merely a winner). (Due to the possibility of ties, often multiple aggregate rankings or winners are allowed to be returned.) The framework is extremely general and so finds applications in many different settings. The voters can be, e.g., people or software agents; the alternatives can be, e.g., political representatives or joint plans. Given this generality, it is perhaps no surprise that no one voting rule has emerged as the one-size-fits-all best option. When comparing the relative merits of different voting rules in a specific context, various criteria can play a role. In social choice theory, many *axioms* have been defined, which are desirable abstract properties that a voting rule should satisfy. Examples of axioms are *anonymity*, stating (informally) that all voters should be treated equally, and *monotonicity*, stating (informally) that increased support should not harm an alternative. Famous impossibility theorems, like the ones by Arrow (1951) and Gibbard and Satterthwaite (1973; 1975), have demonstrated that every voting rule suffers from some flaws, as certain combinations of desirable properties are incompatible. On the other hand, there are nice axiomatic characterizations that state that, for certain combinations of axioms, one and only one voting rule satisfies all of them simultaneously (e.g., see Young and Levenglick 1978). Besides the compliance with choice-theoretic axioms, other desiderata have been considered. In recent years, there has been an extensive focus on *computational* aspects of voting rules. From this perspective, a good voting rule is one for which determining the outcome is computationally easy, whereas undesirable behavior such as strategic manipulation is computationally intractable (Brandt, Conitzer, and Endriss 2013; Conitzer 2010; Faliszewski, Hemaspaandra, and Hemaspaandra 2010; Faliszewski and Procaccia 2010).

Practical concerns give rise to another criterion that is more difficult to formalize: in order to be adopted for making important collective decisions, a voting rule should be transparent and easy to understand for all participants. This is certainly the case for *(positional) scoring rules* such as *plurality* and *Borda's rule*. Unfortunately, scoring rules are not a panacea. They are not Condorcet-consistent: an alternative that outperforms every other alternative in pairwise comparisons may yet fail to win. They can be easily manipulated, both in an intuitive sense and in some[1] (but not all[2]) precise computational senses. And, they do not satisfy *independence of clones*: the introduction of a new alternative that is an (almost) exact copy of another alternative, so that these two alternatives are ranked adjacent to each other in all votes, can change the winner.[3] Independence of clones may be less important in, for example, an election for a political representative, where a charismatic, high-visibility potential candidate who would be both willing and able to copy another candidate's platform may not be available. On the other hand, independence of clones seems particularly important when voting over plans in AI.[4] For a specific example, consider a planning problem where a robot needs to cross a shallow stream. The robot's two options are to put on rain boots and cross the stream (which could involve putting on the left boot first or the right boot first), or to walk some distance to cross a bridge. The voters here could be different algorithms that

[1]See Bartholdi, III, Tovey, and Trick (1989).

[2]See Conitzer, Sandholm, and Lang (2007), Davies et al. (2011), and Betzler, Niedermeier, and Woeginger (2011).

[3]Tideman (1987) has shown that plurality and Borda's rule are not independent of clones. In the full version of this paper, we generalize these results and show that no non-trivial scoring rule is independent of clones.

[4]Schulze (2011) makes a similar point about voting over plans.

rank the plans with respect to certain criteria, for example the time each option will take or the likelihood that the robot will get wet. If the first option is considered as two separate plans, then one would want to use a voting rule that is independent of clones. Relatedly, in highly anonymous (say, Internet) contexts, independence of clones seems important even when voting over representatives, as it is easy to create an additional online identity. Indeed, cloning has already received some attention in the computational social choice community. For instance, Elkind, Faliszewski, and Slinko (2012) have analyzed the structure of clone sets and Elkind, Faliszewski, and Slinko (2011) have studied the computational complexity of manipulating an election by cloning. We note that voting rules that satisfy independence of clones cannot be manipulated by cloning.

As it turns out, some of the drawbacks of scoring rules can be remedied by moving to *runoff* rules, in which weak alternatives are repeatedly eliminated, until only a single alternative, the winner, remains. Examples are *single transferable vote* (STV), in which the alternative with the weakest plurality score is repeatedly eliminated, and *Baldwin's rule*, in which Borda's rule is used rather than plurality. (There is also *Nanson's rule*, in which all alternatives with a below average Borda score are eliminated.) Unlike Borda's rule from which it derives, Baldwin's rule is Condorcet-consistent (as is Nanson's). Similarly, STV is known to possess desirable properties that plurality does not. Notably, it satisfies independence of clones, and this will be the focus of much of this paper.

Our contribution. The contribution of this paper is threefold.

- We characterize runoff scoring rules using two axioms. The first of these axioms is a weakening of *local independence of irrelevant alternatives* (LIIA), which was used by Young (1988) to characterize Kemeny's rule. The second is a variant of *consistency*, which features prominently in an important axiomatic characterization of scoring rules (Smith 1973; Young 1975).[5]
- We show that STV is the only runoff scoring rule that is independent of clones. Thus, STV is characterized by the combination of three axioms. (We are aware of one other axiomatic characterization that uses an independence-of-clones criterion, namely by Laslier (2000), who uses it to characterize an SCF known as the *essential set*.)
- We demonstrate the versatility of our approach by deriving a number of further axiomatic characterizations of specific runoff scoring rules, on the backs of earlier characterizations of specific scoring rules.

2 Preliminaries

Let A be a finite set of *alternatives*. For a subset $B \subseteq A$, a (strict) *ranking* of B is a permutation of B. The set of all rankings of B is denoted by $\mathcal{L}(B)$. We will often focus on the alternative that is ranked last in a ranking. For a given ranking $r = (a_1, \dots, a_m)$, the last-ranked alternative a_m is called the *bottom element* of r, denoted $bottom(r)$. For a set of rankings $\{r_1, \dots, r_\ell\}$, $bottom(\{r_1, \dots, r_\ell\}) = \bigcup_{i=1}^{\ell} bottom(r_i)$ is the set of all alternatives that are the bottom element of some ranking in the set.

Let N be a finite set of voters. The preferences of voter $i \in N$ are represented by a ranking $R_i \in \mathcal{L}(A)$. A *preference profile* is a list $R \in \mathcal{L}(A)^N$ containing a ranking R_i for each voter $i \in N$. For $B \subseteq A$, let $R|_B$ denote the preference profile in which all elements in $A \setminus B$ have been removed from all rankings in R. Furthermore, the *rank distribution* of an alternative a under R is defined as the vector $d(a, R) \in \mathbb{N}^{|A|}$ whose j-th entry is given by the number of voters that rank a in position j.

We distinguish between two types of aggregation functions. A *social choice function* (SCF) f associates with every preference profile R a non-empty set $f(R) \subseteq A$ of alternatives. A *social preference function* (SPF) f associates with every preference profile R a non-empty set $f(R) \subseteq \mathcal{L}(A)$ of rankings of A. Note that ties are allowed due to the fact that the output set can have size greater than 1.

An SCF or SPF is *neutral* if permuting the alternatives in the individual rankings also permutes the set of chosen alternatives, or the set of chosen rankings, in the exact same way. An SCF or SPF is *anonymous* if the set of chosen alternatives, or the set of chosen rankings, does not change when the voters are permuted. An SCF or SPF is called *symmetric* if it is both neutral and anonymous. Throughout this paper, we only consider symmetric SCFs and SPFs.

For social preference functions, we will also consider the following basic properties that we consider quite mild.

Definition 1. *An SPF f satisfies*

- weak unanimity *if $R_i = r$ for all $i \in N$ implies $r \in f(R)$;*
- weak decisiveness *if there exists a preference profile R on at least two alternatives with $|f(R)| = 1$; and*
- continuity at the bottom *if for any two preference profiles R and R' with $bottom(f(R)) = \{a\}$, there exists a natural number k such that $bottom(f(kR \cup R')) = \{a\}$.*[6]

A *scoring rule* is an SCF that is defined by a sequence $s = (s^n)_{n \geq 1}$, where for each $n \in \mathbb{N}$, $s^n = (s^n_1, \dots, s^n_n) \in \mathbb{Q}^n$ is a *score vector*[7] of length n. For a preference profile R on k alternatives, the score vector s^k is used to allocate points to alternatives: the score $s^k(a)$ of a under R is given by $s^k(a) = s^k \cdot d(a, R)^\intercal$, where $d(a, R)^\intercal$ is the transpose of the rank distribution of a under R. (Note that if two alternatives have the same rank distribution, they have the same score for every score vector.) The outcome of the scoring rule is the alternative (or set of alternatives in the case of a tie) with maximal score. Examples of scoring rules are *plurality* ($s^n = (1, 0, \dots, 0)$), *veto* ($s^n = (0, \dots, 0, -1)$), and

[5]Indeed, Smith (1973) was already looking ahead to runoff scoring rules, stating that "It would be interesting, but perhaps very difficult, to characterize [runoff scoring rules]".

[6]The preference profile $kR \cup R'$ consists of k copies of R and one copy of R'.

[7]Note that we do *not* impose the condition that $s^n_1 \geq \ldots \geq s^n_n$. This will result in us having to deal with some unintuitive rules in Section 4, but we obtain a more complete result this way, which will also be useful later on.

Borda's rule ($s^n = (n-1, n-2, \dots, 0)$). We consider two scoring rules s and t identical if for each $n \in \mathbb{N}$, the score vector t^n is an affine transformation of s^n. For instance, the scoring rule given by $s^n = (-1, -2, \dots, -n)$ is identical to Borda's rule.

Every scoring rule s gives rise to a *(one-at-a-time) runoff scoring rule* as follows. As long as the number of alternatives is greater than or equal to two, iteratively eliminate one alternative with the *lowest* score (according to s) from all rankings. The runoff scoring rule corresponding to s is the SPF that outputs the ranking in which alternatives appear in reverse elimination order (the alternative that is eliminated first is ranked last, and so on).

An important issue is how ties are handled. In this paper, we use *parallel-universes tie-breaking (PUT)*, i.e., the outcome of a runoff scoring rule is the union of all rankings that result from the iterative procedure described above for *some* way of breaking the ties. The following three SPFs will be of particular interest in this paper: *Single transferable vote (STV)* is the runoff scoring rule that is based on plurality; *Coombs' rule* is the runoff scoring rule that is based on veto; and *Baldwin's rule* is the runoff scoring rule that is based on Borda's rule.[8]

Example 1. *Consider the following preference profile with four voters and three alternatives.*

$$1 \times (a, b, c) \qquad 1 \times (b, c, a)$$
$$1 \times (a, b, c) \qquad 1 \times (c, a, b)$$

Breaking ties using PUT, it can be easily checked that STV selects the three following rankings: (a, b, c), (a, c, b), *and* (c, a, b). *Coombs' rule and Baldwin's rule do not encounter any ties for this profile and uniquely select* (a, b, c).

3 Main Axioms

In this section, we introduce three axioms that we will use to characterize STV. The first axiom states that collective rankings can be constructed recursively by excluding last-ranked alternatives. For a ranking $r \in \mathcal{L}(B)$ and an alternative $a \notin B$, let (r, a) denote the ranking of $B \cup \{a\}$ that agrees with r on B and has $bottom((r, a)) = a$.

Definition 2. *An SPF f satisfies* independence of bottom alternatives *if for all preference profiles R,*

$$f(R) = \bigcup_{a \in bottom(f(R))} \{(r, a) : r \in f(R|_{A\setminus\{a\}})\}.$$

Independence of bottom alternatives is implied by local independence of irrelevant alternatives as introduced by Young (1988).

The second axiom is a variant of consistency (Smith 1973; Young 1975) that focuses on last-ranked (instead of first-ranked) alternatives: If two groups of voters, N and N', collectively rank sets of alternatives C and D last, respectively, then the set of alternatives ranked last by $N \cup N'$ is precisely $C \cap D$, if this is non-empty.

[8] Nanson's rule is not a runoff scoring rule in our sense, as it generally eliminates more than one alternative at a time. See Niou (1987) for the difference between Nanson's rule and Baldwin's rule.

Definition 3. *An SPF f satisfies* consistency at the bottom *if for all $R \in \mathcal{L}(A)^N$ and $R' \in \mathcal{L}(A)^{N'}$ with $N \cap N' = \emptyset$,*

$$bottom(f(R \cup R')) = bottom(f(R)) \cap bottom(f(R'))$$

whenever the set on the RHS is non-empty.

Finally, we consider *independence of clones*. This property was introduced by Tideman (1987) and Zavist and Tideman (1989) for SCFs, and we adapt it to SPFs. We say an SPF is independent of clones if, after ignoring all clones that are not ranked highest, cloning operations do not affect the set of collective rankings. More precisely, a *cloning operation* $\mathcal{C}$ transforms a preference profile R into another preference profile $R^{\mathcal{C}}$ in which one of the alternatives, say a, has been replaced by a set of *clones*. In every ranking R_i, the clones keep the relative position of a with respect to all other alternatives; the ranking among the clones themselves is arbitrary. For a ranking r of the set of all alternatives, including any clones that have been introduced by cloning operation $\mathcal{C}$, we let $\lceil r \rceil_{\mathcal{C}}$ denote the ranking of the set of all original alternatives that is obtained from r by deleting all clones and inserting a at the position where the highest-ranked clone of a was ranked in r. Then:

Definition 4. *An SPF f satisfies* independence of clones *if for all preference profiles R and all cloning operations $\mathcal{C}$,*

$$\lceil f(R^{\mathcal{C}}) \rceil_{\mathcal{C}} = f(R).$$

Example 2. *In the profile given in Example 1, replace alternative a by two clones a and a' such that the profile becomes*

$$1 \times (a, a', b, c) \qquad 1 \times (b, c, a, a')$$
$$1 \times (a', a, b, c) \qquad 1 \times (c, a, a', b)$$

It is easily verified that Baldwin's rule selects (c, a, a', b), among other rankings. Since $\lceil (c, a, a', b) \rceil_{\mathcal{C}} = (c, a, b)$ was not selected by Baldwin's rule in the original profile, this example shows that Baldwin's rule is not independent of clones.

4 Axiomatic Characterization of STV

Our main result is a characterization of STV in terms of the axioms that were introduced in the previous sections. We will show that STV is the only symmetric SPF that satisfies the three axioms introduced in Section 3 and the basic properties in Definition 1. Due to space constraints, several proofs are omitted. They can be found in the full version of this paper.

We start by checking that STV indeed satisfies the conditions. It is easy to check that STV is symmetric and satisfies the basic properties in Definition 1. Being a runoff scoring rule, STV also satisfies independence of bottom alternatives and consistency at the bottom (see Lemma 1). It is left to show that STV is independent of clones. Tideman (1987) has proved that STV is independent of clones when interpreted as an SCF. We recap the argument and extend it to also show clone independence in the SPF setting.

Recall that STV is based on the plurality scoring rule. Since only first-ranked alternatives receive points under plurality, in the first round of STV, the only effect of a clone operation is that the points of the cloned alternative are distributed among its clones, while the number of points remains the same for all other alternatives. This leaves only two possibilities for the first alternative to be eliminated: either it is the alternative that would have been eliminated in the original profile, or some clone is eliminated. In the former case, the elimination order is not affected. In the latter case, the points of the eliminated clone transfer exclusively to other clones. Applying this reasoning iteratively, it follows that (1) as long as there is more than one clone left, the elimination order does not change (when ignoring the elimination of clones), and (2) the last remaining clone accumulates all the points of the other clones, ending up with exactly the number of points the cloned alternative received in the original profile. Therefore, from the point where only one clone is left, the elimination procedure proceeds exactly the same as in the original profile. This shows that the SPF version of STV satisfies independence of clones.

The remainder of this section is devoted to showing that STV is indeed the *only* SPF satisfying the aforementioned set of axioms. We start by characterizing runoff scoring rules.

Lemma 1. *Let f be a symmetric SPF satisfying continuity at the bottom. Then f satisfies independence of bottom alternatives and consistency at the bottom if and only if it is a runoff scoring rule.*

Proof. We first show that every runoff scoring rule satisfies independence of bottom alternatives and consistency at the bottom. Let s be the sequence of score vectors on which f is based. Independence of bottom alternatives follows immediately from the definition of a runoff scoring rule (and from our assumption that PUT is used in case of ties). As for consistency at the bottom, observe that the set of last-ranked alternatives $bottom(f(R))$ coincides with the set of *winners* for the scoring rule $\bar{s}$, which is given by $\bar{s}_j^k = -s_j^k$ for all $1 \leq j \leq k$. All scoring rules are consistent (Smith 1973; Young 1975), and consistency of $\bar{s}$ implies consistency at the bottom of f.

For the other direction, let f be a symmetric SPF satisfying continuity at the bottom, independence of bottom alternatives, and consistency at the bottom. Independence of bottom alternatives implies that f is uniquely determined by the SCF, call it g, that selects $bottom(f(R))$. Symmetry of f implies that g is symmetric as well. Furthermore, continuity at the bottom of f implies continuity of g and consistency at the bottom of f implies consistency of g. We can therefore apply the result by Smith (1973) and Young (1975), which states that a symmetric SCF is continuous and consistent if and only if it is a scoring rule. It follows that the bottom elements of $f(R)$ are selected by a scoring rule. Therefore, f must be the corresponding runoff scoring rule. □

Since Smith (1973) has shown that no runoff scoring rule is monotonic, Lemma 1 yields the following impossibility result.

Corollary 1. *There does not exist a symmetric SPF that satisfies independence of bottom alternatives, consistency at the bottom, continuity at the bottom, and monotonicity.*

From now on, we will identify a runoff scoring rule f with the sequence $s = (s^k)_{k \in \mathbb{N}}$ that defines the scoring rule on which f is based. Call a score vector s^k *trivial* if $s_i^k = s_j^k$ for all $1 \leq i, j \leq k$. The following lemma shows that a weakly decisive runoff scoring rule can only be independent of clones if all score vectors are non-trivial.

Lemma 2. *Let s be a weakly decisive runoff scoring rule. If s is independent of clones, then s^n is non-trivial for all $n \geq 2$.*

We go on to show that, under mild conditions, it is possible to construct a preference profile, including one pair of clones, such that two alternatives (disjoint from the clone set) obtain an arbitrarily lower score than all other alternatives. This construction will later be useful because it enables us to push any two alternatives to the bottom of the ranking, no matter how many points these alternatives received in the original profile.

Lemma 3. *Let s^n be a non-trivial score vector with $n \geq 5$ and $s^n \neq (1, 0, 1, 0, 1)$. For every $M \in \mathbb{N}$, there exists a preference profile R on a set A of n alternatives with the following properties:*

- *there are two alternatives a and a' that appear in consecutive positions in every ranking and that have the same rank distribution,*
- *there are two alternatives c and d with $\{c, d\} \cap \{a, a'\} = \emptyset$ that have the same rank distribution, and*
- *$s^n(c) + M = s^n(d) + M < \min\{s^n(b) : b \in A \setminus \{c, d\}\}$.*

We will use the construction in Lemma 3 to show that the score vectors of a clone-independent runoff scoring rule must have a particularly simple form.

Definition 5. *A score vector s^n of length n is a* plurality/veto combination *if $s_i^n = s_j^n$ for all $2 \leq i, j \leq n-1$.*

By definition, all score vectors of length three or smaller are plurality/veto combinations. When dealing with plurality/veto combinations, we will always assume that the score vector is normalized so that $s_2^n = ... = s_{n-1}^n = 0$.

The following is our key lemma.

Lemma 4. *Let s be a weakly decisive runoff scoring rule. If s is independent of clones, then s^n is a plurality/veto combination for all $n \in \mathbb{N}$.*

Proof. Let s be a weakly decisive runoff scoring rule that is independent of clones. Trivially, s^n is a plurality/veto combination for all $n \leq 3$. We first consider the case $n \geq 5$, and deal with the four-alternative case later.

Let $n \geq 5$ and suppose that s^n is not a plurality/veto combination. Then there exists $2 \leq k \leq n-2$ such that $s_k^n \neq s_{k+1}^n$. We will show that s is not independent of clones.

Assume for now that $s^n \neq (1, 0, 1, 0, 1)$. We construct a two-voter profile such that two alternatives, say c and d, have the same rank distribution in the case of $n-1$ alternatives but obtain different scores in the n alternative profile that results from cloning one alternative, say a, into clones a and a'.

In the case $k < n-2$, consider the $n-1$ alternative profile

$$1 \times (a, ..., c, ..., b, d) \qquad 1 \times (b, ..., d, ..., a, c)$$

where c and d are each once ranked k-th and once n-th. After cloning a, the profile becomes

$$1 \times (a, a', ..., c, ..., b, d) \qquad 1 \times (b, ..., d, ..., a, a', c)$$

Here, c obtains one $(k+1)$-th ranking and one n-th ranking, and d obtains one k-th ranking and one n-th ranking.

In the case $k = n-2$, consider the $n-1$ alternative profile

$$1 \times (c, b, ..., d, a) \qquad 1 \times (d, a, ..., c, b)$$

and clone a to obtain

$$1 \times (c, b, ..., d, a, a') \qquad 1 \times (d, a, a', ..., c, b)$$

so that d receives a first and a k-th ranking and c receives a first and a $k+1$-th ranking. In both cases, $s_k \neq s_{k+1}$ implies that c and d have different scores after cloning.

We now apply Lemma 3 to append a profile of votes such that c and d have the same rank distribution but score the lowest of all alternatives in the cloned profile. Hence we ensure that the same alternative (either c or d) is ranked uniquely last in every universe after cloning, whereas before cloning both were ranked last in at least one universe or neither was ranked last in any universe. This contradicts independence of clones.

Next we deal with the case $s^5 = (1, 0, 1, 0, 1)$ (which does not allow us to apply Lemma 3). By the above reasoning we may assume that s^n is a plurality/veto combination for all $n > 5$. Write $s^6 = (w, 0, 0, 0, 0, z)$, $w, z \in \mathbb{Q}$. Consider the ten voter, five alternative profile

$$\begin{array}{ll} 2 \times (a, e, c, d, b) & 2 \times (a, d, b, c, e) \\ 2 \times (a, b, d, c, e) & 2 \times (c, a, d, b, e) \\ & 2 \times (c, a, d, e, b) \end{array}$$

Note that $s^5(\cdot) = 6$ for all five alternatives. Therefore, every alternative is eliminated first in some parallel universe. We will now show that s is not independent of clones. The clone operation will depend on the values of w and z.

If $w > 0$, we clone a in such way that a and a' have the same rank distribution. In the cloned profile, $s^6(a) = s^6(a') > 0 = s^6(d)$. If $w < 0$ then we similarly clone a, and now $s^6(a) = s^6(a') = 3w > 4w = s^6(c)$. In both cases, a is no longer the first alternative eliminated in any universe, contradicting independence of clones. Lastly suppose $w = 0$. Then Lemma 2 implies $z \neq 0$, as otherwise s^6 would be trivial. We can now clone e and apply reasoning that is completely analogous to before, in both the $z < 0$ and $z > 0$ cases.

Finally, we consider the case $n = 4$. By the above, we may assume s^n is a plurality/veto combination for all $n > 4$. Assume for the sake of contradiction that s^4 is not a plurality/veto combination, so $s^4_2 \neq s^4_3$. Consider one voter with preferences (a, b, c, d) over four alternatives, so that $s^4(b) \neq s^4(c)$. Now we clone a so that the voter's preferences become (a, a', b, c, d). Since s^5 is a plurality/veto combination, we know $s^5(b) = s^5(c)$. We can apply Lemma 3 to append a profile of votes such that b and c score lower than all other alternatives, and hence both b and c are first eliminated in at least one universe. Before cloning, however, one of b and c was *not* eliminated first in any universe. This contradicts independence of clones. □

The previous lemma shows that independence of clones requires score vectors of the form $s^n = (x_n, 0, ..., 0, y_n)$. A closer analysis allows us to narrow down the possibilities even further: the following lemma implies that both x_n and y_n must be non-negative, and in the case where there exist $n', n'' \geq 3$ such that $x_{n'}$ and $y_{n''}$ are positive, the fraction $\frac{x_n}{y_n}$ has to be constant for all $n \geq 3$ (since they can never both be zero by weak decisiveness).

Lemma 5. *Let s be a weakly decisive runoff scoring rule that satisfies independence of clones. Then, for all $n \geq 3$, $s^n = (x_n, 0, ..., 0, y_n)$ with $x_n \geq 0$, $y_n \geq 0$, and $x_n y_{n+1} = x_{n+1} y_n$.*

We finally arrive at a full characterization of the class of scoring rules that give rise to clone-independent runoff scoring rules.

Lemma 6. *A weakly decisive runoff scoring rule s is independent of clones if and only if one of the following four cases holds:*

- $s^2 = (1, 0)$ *and* $s^k = (x, 0, ..., 0, 1)$ *for all* $k \geq 3$ *and some* $x \geq \frac{1}{2}$,
- $s^2 = (0, 1)$ *and* $s^k = (1, 0, ..., 0, x)$ *for all* $k \geq 3$ *and some* $x \geq \frac{1}{2}$,
- $s^k = (1, 0, ..., 0)$ *for all* $k \geq 2$, *or*
- $s^k = (0, ..., 0, 1)$ *for all* $k \geq 2$.

We are now ready to prove our main result.

Theorem 1. *STV is the unique symmetric SPF satisfying independence of bottom alternatives, consistency at the bottom, independence of clones, and the properties in Definition 1.*

Proof. We have already discussed that STV satisfies all the properties. As to whether any other SPF could satisfy all the properties, Lemma 1 implies that such an SPF must be a runoff scoring rule. Then, Lemma 6 characterizes runoff scoring rules that are independent of clones. We finally argue that among these, only STV satisfies weak unanimity, as follows. We claim that if a runoff scoring rule s satisfies weak unanimity, then $s^n_{n-1} \geq s^n_n$ for all $n \geq 2$. Suppose this is not true for some $n \in \mathbb{N}$ and consider a profile consisting of a single voter with preferences $(a_1, ..., a_n)$ over the n alternatives. Then, a_n is not the first alternative eliminated in any parallel universe, since $s^n_{n-1} < s^n_n$. This violates weak unanimity. Hence all rules listed in Lemma 6 other than STV violate weak unanimity. □

5 Other Characterization Results

We can use Lemma 1 to obtain characterizations of two additional runoff scoring rules, leveraging existing characterizations of scoring rules.[9]

Definition 6. *An SPF satisfies the* bottom-majority criterion *if $bottom(f(R)) = bottom(r)$ whenever more than half of the voters agree on a ranking r.*

[9] For an overview of such characterizations, see Merlin (2003).

Definition 7. *An SPF f satisfies the* ranking-majority criterion *if $f(R) = \{r\}$ whenever more than half of the voters agree on a ranking r.*

Theorem 2. *Coombs' rule is the unique symmetric SPF satisfying independence of bottom alternatives, consistency at the bottom, the ranking-majority criterion, and continuity at the bottom.*

Proof. Lemma 1 requires that the rule is a runoff scoring rule. Lepelley (1992) has shown that plurality is the only scoring rule that satisfies the majority criterion. Equivalently, veto is the only scoring rule such that $bottom(f(R)) = \{x\}$ whenever a majority of voters rank x last. We now show that veto is also the only scoring rule which satisfies the (weaker) bottom-majority criterion. Clearly veto does satisfy this. Now consider some other scoring rule g. Take a profile R in which a majority of voters rank a in last place but $bottom(g(R)) \neq a$. Such a profile exists because g is not veto. Define $b = bottom(g(R))$ and modify R in the following way. For all those voters who ranked a last, move b to a position which scores the lowest out of all the positions not occupied by a. This does not increase the score of b—in particular, b still scores less than a. Now place the other alternatives such that all voters who rank a last now vote identically, so that more than half the voters now agree on their ranking of the alternatives. The loser may have changed but it still can not be a since the score of b is strictly smaller than the score of a. Therefore, g does not satisfy the bottom-majority criterion.

Hence, among all runoff scoring rules, Coombs' rule uniquely satisfies the ranking-majority criterion. because it is the only runoff scoring rule guaranteed to eliminate the alternative ranked in the bottom position by a majority (should one exist) at every step. □

In a similar fashion, we can characterize Baldwin's rule.

Definition 8. *An SPF f is* Condorcet-consistent *if a Condorcet winner is ranked first in all rankings in $f(R)$.*

Theorem 3. *Baldwin's rule is the unique symmetric SPF satisfying independence of bottom alternatives, consistency at the bottom, continuity at the bottom, and Condorcet-consistency.*

Nanson's rule is also Condorcet-consistent. However, it is not a runoff scoring rule according to our definition.

6 Non-Scoring Runoff Rules

So far, we have only considered runoff rules where a scoring rule is used to determine which alternative is ranked last. If we drop the criterion of consistency at the bottom, we can also choose the last alternative using another rule. One particularly remarkable way of doing so is to use *the inverted STV rule* to choose which alternative is ranked last. The inverted STV rule is the runoff scoring rule given by $s^n = (0, \ldots, 0, 1)$. (Equivalently, it is STV applied to the inverted votes.) Eliminate that alternative, and repeat. We call this *nested runoff voting*, as it involves a runoff rule within a runoff rule.

Example 3. *For the preference profile given in Example 1, nested runoff voting selects (a, b, c) and (b, c, a).*

As it turns out, this rule is independent of clones.

Proposition 1. *Nested runoff voting satisfies independence of bottom alternatives and independence of clones (but not consistency at the bottom).*

Proof. It fails consistency at the bottom because inverted STV, which is used to determine the bottom alternative, is not a scoring rule. It satisfies independence of bottom alternatives because it is a runoff rule. By Lemma 6, inverted STV is independent of clones, so cloning does not affect which alternative is ranked last at any point (with the possible exception of a clone being ranked last instead of the original alternative that it cloned). As a result, nested runoff voting also satisfies independence of clones. □

Therefore, nested runoff voting serves as an example of an SPF that satisfies independence of bottom alternatives and independence of clones, but not consistency at the bottom. Furthermore, any runoff scoring rule other than those characterized by Lemma 6 (such as Baldwin's or Coombs' rule) satisfies independence of bottom alternatives and consistency at the bottom, but not independence of clones. SPFs that satisfy consistency at the bottom and independence of clones but not independence of bottom alternatives appear harder to find.

7 Conclusion

In this paper, we focused on the axiomatic characterization of runoff scoring rules. While we gave axiomatizations of Baldwin and Coombs' rules as well, our main contribution was an axiomatic characterization of STV based on the independence of clones criterion.

As far as we are aware, STV is the only known example of a neutral voting rule that is both independent of clones and NP-hard to manipulate by a single manipulator.[10] This is so because the only neutral variant of ranked pairs is the variant using PUT (Brill and Fischer 2012), and this variant is not independent of clones (by Example 1 in Zavist and Tideman 1989); Schulze's rule is easy to manipulate by a single manipulator (Parkes and Xia 2012); and we are not aware of any other non-trivial rules that are independent of clones. In our view, independence of clones is a very important criterion, and we do not yet have a good understanding of what properties make a rule independent of clones. Hence, characterizing rules that satisfy independence of clones is an important direction for future research.

If we were to view STV as a social *choice* function rather than an SPF, and relax our definition of independence of clones accordingly, we do not know whether other non-trivial runoff scoring rules would satisfy the property. Indeed, our proofs relied heavily on being able to alter *some* position in the ranking. However, we are not aware of any runoff scoring rules that are independent of clones in the SCF context, other than STV.

[10] We do not conjecture that it is the unique rule with these properties. For example, nested runoff voting may very well be NP-hard to manipulate as well.

Acknowledgements

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

Arrow, K. J. 1951. *Social Choice and Individual Values*. New Haven: Cowles Foundation, 1st edition. 2nd edition 1963.

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

Betzler, N.; Niedermeier, R.; and Woeginger, G. 2011. Unweighted coalitional manipulation under the borda rule is np-hard. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, 55–60.

Brandt, F.; Conitzer, V.; and Endriss, U. 2013. Computational social choice. In Weiß, G., ed., *Multiagent Systems*. MIT Press, 2nd edition. 213–283.

Brill, M., and Fischer, F. 2012. The price of neutrality for the ranked pairs method. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI)*, 1299–1305. AAAI Press.

Conitzer, V.; Sandholm, T.; and Lang, J. 2007. When are elections with few candidates hard to manipulate? *Journal of the ACM* 54(3):Article 14, 1–33.

Conitzer, V. 2010. Making decisions based on the preferences of multiple agents. *Communications of the ACM* 53(3):84–94.

Davies, J.; Katsirelos, G.; Narodytska, N.; and Walsh, T. 2011. Complexity of and algorithms for Borda manipulation. In *Proceedings of the 25th AAAI Conference on Artificial Intelligence (AAAI)*, 657–662. AAAI Press.

Elkind, E.; Faliszewski, P.; and Slinko, A. 2011. Cloning in elections: Finding the possible winners. *Journal of Artificial Intelligence Research* 42:529–273.

Elkind, E.; Faliszewski, P.; and Slinko, A. 2012. Clone structures in voters' preferences. In *Proceedings of the 13th ACM Conference on Electronic Commerce (ACM-EC)*, 496–513.

Faliszewski, P., and Procaccia, A. D. 2010. AI's war on manipulation: Are we winning? *AI Magazine* 31(4):53–64.

Faliszewski, P.; Hemaspaandra, E.; and Hemaspaandra, L. 2010. Using complexity to protect elections. *Communications of the ACM* 53(11):74–82.

Gibbard, A. 1973. Manipulation of voting schemes. *Econometrica* 41:587–602.

Laslier, J.-F. 2000. Aggregation of preferences with a variable set of alternatives. *Social Choice and Welfare* 17:269–282.

Lepelley, D. 1992. Une caractérisation du vote à la majorité simple. *RAIRO. Recherche opérationnelle* 26(4):361–365.

Merlin, V. 2003. The axiomatic characterizations of majority voting and scoring rules. *Mathématiques et sciences humaines. Mathematics and social sciences* (163).

Niou, E. M. S. 1987. A note on Nanson's rule. *Public Choice* 54:191–193.

Parkes, D., and Xia, L. 2012. A complexity-of-strategic-behavior comparison between Schulze's rule and ranked pairs. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI)*, 1429–1435. AAAI Press.

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

Schulze, M. 2011. A new monotonic, clone-independent, reversal symmetric, and Condorcet-consistent single-winner election method. *Social Choice and Welfare* 36(2):267–303.

Smith, J. H. 1973. Aggregation of preferences with variable electorate. *Econometrica* 41(6):1027–1041.

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

Young, H. P., and Levenglick, A. 1978. A consistent extension of Condorcet's election principle. *SIAM Journal on Applied Mathematics* 35(2):285–300.

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

Young, H. P. 1988. Condorcet's theory of voting. *The American Political Science Review* 82(4):1231–1244.

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