Soc Choice Welfare (1987) 4:185–206

Social Choice and Welfare

Independence of Clones as a Criterion for Voting Rules

T. N. Tideman*

Virginia Polytechnic Institute and State University, Blacksburg, VA 24061, USA

Received June 16, 1986 / Accepted in revised form February 26, 1987

Abstract. "Independence of clones" is a generalization of the condition of not being subject to the perverse consequences of vote splitting that arise under plurality voting. A new voting rule that is at least "almost always" independent of clones is obtained by the following algorithm: Require the collective ranking of the candidates to be consistent with the paired comparisons decided by the largest and second largest margins, and then, if possible, with the paired comparison decided by the third largest margin, and so on. The advantages of this "ranked pairs" rule over previously proposed voting rules that are independent of clones is that it possesses Condorcet consistency, non-negative responsiveness, and "resolvability" (the property that every tie be within one vote of being broken).

I. Introduction

Example 1. When I was 12 years old I was nominated to be treasurer of my class at school. A girl named Michelle was also nominated. I relished the prospect of being treasurer, so I made a quick calculation and nominated Michelle's best friend, Charlotte. In the ensuing election I received 13 votes, Michelle received 12, and Charlotte received 11, so I became treasurer.

It would be widely agreed that it would be attractive for my stratagem not to be feasible. Indeed, the institution of parties helps insure that candidates with similar constituencies do not split the vote they can attract. But parties also deny the full electorate the opportunity to choose among similar candidates.

The Borda voting rule is subject to a different perverse consequence of having two or more very similar candidates on a ballot. Under the Borda rule, the relative score of one candidate, x, is improved by having on the ballot another candidate, x', who is ranked by all voters immediately below x.

* I am indebted to Jaques Crémer, Amoz Kats, Martin Ricketts, Thomas Schwartz and an anonymous referee for helpful suggestions.

The study these and other perverse results, this paper defines the concept of candidates who are *clones*. A proper subset of two or more candidates, S, is a set of clones if no voter ranks any candidate outside of S as either tied with any element of S or between any two elements of S. This definition is suggested by the idea that if the definition is satisfied, then the voters' rankings are consistent with the hypothesis that the candidates in S are arbitrarily close to one another in some shared perceptual space in which voters locate candidates.

A voting rule is defined to be independent of clones if and only if the following two conditions are met when clones are on the ballot:

1. A candidate that is a member of a set of clones wins if and only if some member of that set of clones wins after a member of the set is eliminated from the ballot.
2. A candidate that is *not* a member of a set of clones wins if and only if that candidate wins after any clone is eliminated from the ballot.

Among previously proposed voting rules, the alternative vote and the GOCHA rule are independent of clones. However, the alternative vote does not possess Condorcet consistency (the property of always selecting the candidate, if there is one, that beats all other candidates in head-to-head contests), or the property of non-negative responsiveness, both of which have been widely held to be important for voting rules to possess. The GOCHA rule possesses these properties, but it calls every election involving a cycle a tie, and therefore lacks "resolvability," the property that every tie be within one vote of being broken. The minimax rule and the Young rule are independent of clones that come in pairs, but these rules are not independent of clones that come in sets of three or more. None of a variety of other previously proposed voting rules that were examined are independent of clones.

This paper introduces a new voting rule, the "ranked pairs" rule, that is at least "almost always" independent of clones. The ranked pairs rule operates by first ranking *the set of pairs* of candidates according to the majorities obtained when the elements of the pairs are compared head-to-head. The collective ranking of all candidates is required to be consistent with the pairings decided by the largest and second-largest majorities, then, if possible, with the pairing decided by the third largest majority, and so on until a unique ranking is determined. The winner is the candidate at the top of this ranking. If there are ties in the ranking of pairs or ties in some pairings, the outcome is a tie among all the candidates that win under some way of breaking these ties.

The ranked pairs rule is definitely independent of clones when no head-to-head contest is a tie and all head-to-head contests, other than those between a set of clones and some other candidate, are decided by different margins. I conjecture that the ranked pairs rule is independent of clones whether or not this condition is met, but that remains to be proved or disproved. The ranked pairs rule also possesses Condorcet consistency, non-negative responsiveness and resolvability.

The issue of clones has not entirely escaped the attention of other writers. The ability of the alternative vote and the single transferable vote to avoid difficulties of vote splitting has long been recognized. One of the notable virtues of approval voting [3] is that it does not require voters to choose among effectively identical candidates. In a paper developed while I was developing this one, Chaudhuri and Roy [4] define a "party or ideological group" the same way that I define clones. They define two "no gain no loss nonproliferation criteria" that correspond to my "independence of clones," and mention, without providing proofs, some of the results regarding independence of clones that appear in Sect. III of this paper. In another concurrently developed paper, concerned primarily with the nature of equilibrium under approval voting when all voters know the preferences of all other voters, Zavist [15] defines "adjacent candidates" and "adjacency preservation" in the way I define "clones that come in pairs" and "independence of clones that come in pairs." Zavist also defines a "sequential function" that corresponds to the "ranked pairs algorithm" defined in this paper.

The paper is organized as follows: Section II provides definitions of required terms. Section III explores previously proposed voting rules with respect to independence of clones. Section IV explores modifications of these rules in search of a rule that is independent of clones and also satisfies other criteria. Section V develops the ranked pairs voting rule. Section VI appraises the value of the ranked pairs rule relative to other rules.

II. Definitions

Ranking-Based Voting Rule

At the most general level, a voting rule is a way of combining input from a number of persons and determining which of a number of previously identified candidates will be "chosen." Under plurality voting, the input from each voter is the candidate, if any, for whom the voter votes. Approval voting [3] requires each voter to specify the set of candidates that he or she "approves of." Sometimes, as in the judging of certain sporting events, the input is scores given to competitors by judges. In the demand-revealing process [12], the input is sacrifices that voters are willing to make to have one candidate chosen instead of another. But voting theorists have been concerned primarily with voting rules in which the input from voters is rankings of the candidates. Sometimes these are required to be strict rankings. That is, voters are not permitted to report ties in their rankings. This paper concentrates on ranking-based voting rules (voting rules in which the input from voters is either strict rankings or rankings that may contain ties), but also discusses plurality and approval voting.

If voters are permitted to include ties in their rankings, or if the number of voters is a multiple of any number between two and the number of candidates, then it is possible for the outcome to have such symmetry that the result must be declared a tie if no special preference is to be given to any voter or candidate. It is possible to make mathematical provision for multiple winners, but to the extent that purpose of voting is to identify a single plan of coordinated actions as the intentions of a collectivity, the purpose cannot be served until a unique winner is identified. Voting rules sometimes use random processes to resolve ties and sometimes designate particular individuals as tie-breakers. Thus, for purposes of this paper, a ranking-based voting rule has two parts. The first part associates, with any preference profile (list of rankings) of any specified set of candidates, a non-empty subset of the ranked candidates. That subset is the set of winning candidates. If the set of winning candidates contains more than one element, then the second part of the voting rule selects a unique winner on the basis of some random or non-anonymous tie-breaking device.

In providing rationales voting rules, one treats the rankings reported by voters as if they expressed the voters' true preferences. However, there is no presumption that the rankings reported by voters must necessarily represent the voters' true preferences.

Anonymity

A ranking-based voting rule is anonymous if no permutation of the order of the rankings in a preference profile affects the set of winning candidates.

Neutrality

A ranking-based voting rule is neutral if, when the positions of any two candidates are interchanged on all rankings in a preference profile, the winning statuses of the two candidates (relative to the first part of the voting rule) are interchanged.

Condorcet Consistency

A candidate, x, is dominant for a given preference profile if, for every other candidate, y, the number of voters who rank x ahead of y is greater than the number of voters who rank y ahead of x. Condorcet [5] appears to have been the first to suggest that a voting rule ought to select the dominant candidate when there is one, and therefore a ranking-based voting rule is said to be Condorcet consistent if it selects the dominant candidate when there is one.

Non-negative responsiveness

A ranking-based voting rule possesses non-negative responsiveness if a change in the preference profile that moves one candidate, x, up in the ranking of the one voter, i, while the relative positions of all other candidates in i's ranking remain fixed, as do the positions of all candidates in the rankings of all other voters, cannot change x from an element of the set of winning candidates to a losing candidate.

Resolvability

A ranking-based voting rule is resolvable if, when the first part of the rule generates a tie among the candidates in some set, S, then for any candidate x in S, there is some ranking which, when appended to the preference profile, generates a winning set containing only x.

Homogeneity

A ranking-based voting rule is homogeneous if, for any preference profile P, the set of winning candidates associated with P is the same as the set of winning candidates associated with the preference profile obtained by repeating P any finite number of times.

III. Independence of Clones in Previously Proposed Voting Rules

Plurality

Plurality is the most commonly used voting rule: Each voter votes for at most one candidate, and the candidate receiving the most votes wins. Under plurality, the full rankings of voters are generally not known, which means that it is generally not possible to know whether a candidate is dominant or whether a set of candidates are clones. Nevertheless, one can represent plurality as a function of rankings, so these issues may be addressed.

Example 1 (my election as class treasurer) may be expanded to illustrate that when plurality is represented as a function of rankings, it is not independent of clones. Suppose that each of Charlotte's voters ranked Michelle second and each of Michelle's voters ranked Charlotte second. Then Michelle and Charlotte would be clones. With Charlotte removed from the ballot, Michelle becomes the top candidate for all of the voters that had previously voted for Charlotte, and Michelle then wins. Thus plurality is not independent of clones.

If every voter who did not vote for Michelle had ranked her second, then Michelle would have been a dominant candidate. That she could fail to win even though dominant illustrates that plurality is not Condorcet consistent.

Approval Voting

Approval voting is a variation on plurality in which voters are able to give one vote to each of as many candidates on the ballot as they choose [3]. Voters thus divide the candidates between those they "approve" and those they do not approve. Approval voting cannot be expressed as a function of voters' rankings of the candidates, and therefore it is not a ranking-based voting rule. Applying the concept of clones to approval voting is somewhat problematic because clones are defined in terms voters' rankings. Still, in the spirit of independence of clones, it is worth noting that if there were two or more candidates who were so similar that every voter would rank them as tied if given the chance to rank them (perhaps Michelle and Charlotte in Example 1), then under approval voting any voter who approved any one candidate in any such set of "perfect clones" could be expected to approve all candidates in the set, and the number of perfect clones present (as long as there was at least one) would have no effect on whether the perfect clones were in the set of winning candidates under approval voting.

In the event of a tie between a set of perfect clones and one or more other candidates, if the tie were broken by a random process in which all the tied candidates had the same probability of winning, then the number of perfect clones present would affect the chance that a perfect clone would win. To keep this departure from the spirit of independence of clones from occurring, one could use the following tie-breaking procedure for approval voting: Select a random ranking of the voters. Find the first voter who approves of some of the tied candidates and not others. Eliminate all candidates not approved by that voter. Find the next voter who approves of some of the remaining tied candidates and not others. Eliminate all remaining tied candidates not approved by that voter, and so on until there is a unique winner or the list of voters is exhausted. If the list of voters is exhausted before the set of winners is reduced to a single candidate, then the candidates who remain from the winning set form a set among which no voter made any distinction, and a random process in which each of them has the same chance of winning can be applied without violating the spirit of independence of clones.

While approval voting can thus be made independent of perfect clones, approval voting is not generally independent of clones. In the elaboration of Example 1 that was used to show that plurality is not independent of clones, Michelle and Charlotte were clones. If the election had been held under approval voting, it could have happened that each voter voted for just his or her first choice, in which case the result would have been the same under approval as occured under plurality. Thus Example 1 illustrates the lack of independence of clones in approval voting as well as in plurality. Since Michelle could have been a dominant candidate even though I could have won under approval voting, Example 1 also shows that approval voting is not Condorcet consistent.

Alternative Vote

The alternative vote is a special case of the form of proportional representation known as the single transferable vote [2, pp. 72–74]. When there is just one candidate to be selected by a given electorate, the name "alternative vote" is used. Under the alternative vote, each voter submits a strict ranking of the candidates. The votes are first sorted according to the candidate named first on them. If no candidate has a majority, the candidate with the fewest votes is eliminated, and these votes are transferred to the candidate named second on them. If there is still no candidate with a majority, the remaining candidate with the fewest votes (including those received by transfer) is eliminated, and these votes are transferred to the still-remaining candidates ranked highest on them. This process continues until a candidate accumulates a majority of the votes, and that candidate wins.

That the alternative vote is independent of clones may be shown as follows: Suppose an election is held by the alternative vote over a set of candidates that includes a set of clones, C. Whenever a candidate from C is eliminated, that candidate's votes will be reallocated to other elements of C, as long as elements of C remain uneliminated. And every vote reallocated to an element of C is a vote that would have gone to some other member of C, if any were present, if the member of C who received the vote had not been present. Therefore the number of votes going to the last remaining element of C and to every other remaining candidate when just one element of C remains, as well as the positions of the last element of C and all other contenders on all votes at that time, will be independent of which non-empty subset of the elements of C is included on the ballot. Apart from the treatment of ties, this demonstrates that the alternative vote is independent of clones.

It is customary when conducting an election by the alternative vote to break any ties for fewest votes that arise during the elimination process at the time that they arise. This would not be in keeping with the definition of a voting rule given earlier, in which all of the counting occurs before any of the tie-breaking. To make the alternative vote consistent with the definition of a ranking-based voting rule (in a way that would in fact be impractical unless the counting were done by a computer) one could specify that in the event of a tie for the fewest votes, the course of the election under each breaking of the tie would be examined. The set of winning candidates would be those who could win under some breaking of the ties, with the unique winner selected from these by some tie-breaking device. To keep the probability that the winner will be from a set of clones independent of the number of clones on the ballot, one can use a variation on the tie-breaking device described in connection with approval voting: Select a voter at random and declare as the final winner the tied candidate ranked highest by that voter. The alternative vote, so conceived, is independent of clones.

While independence of clones is an attractive property of the alternative vote, its failure to satisfy Condorcet consistency [2, pp. 73–74] and non-negative responsiveness [7] represent serious shortcomings.

GOCHA

The GOCHA rule (general optimal choice axiom, [10, Chap. 6] was devised by Schwartz to describe the limits within which collective choice in the presence of cycles could be circumscribed by a presumption that pairings under majority rule yield appropriate rankings when cycles do not contradict this. Specification of the GOCHA rule requires a series of definitions. Given a preference profile for a set of candidates S, a non-empty subset of S, s, is undominated if no element of $S \setminus s$ (those elements of S not in s) beats any element of s in head-to-head comparisons by majority rule. If s is undominated and no proper subset of s is undominated, then s is minimum undominated. The set of winners under the GOCHA rule is the union of minimum undominated subsets of S. The removal of a clone cannot affect either the undominatedness of a set, or its minimality if it is undominated. Therefore the GOCHA rule is independent of clones.

The GOCHA rule also possesses Condorcet consistency and non-negative responsiveness, but it lacks resolvability, because it declares all elections with cycles at the top to be ties.

Coombs

The Coombs voting rule [6, p. 399] is a variation on the alternative vote in which the candidate eliminated in each round is the one with the most last-place votes rather than the one with the fewest first-place votes. When successive ballots are used and voters are asked on each ballot only for their last choices, this rule is known as exhaustive voting [2, pp. 69–72].

Example 2 below shows that the Coombs voting rule is not independent of clones. In this and subsequent examples, a number at the top of a column of letters represents the number of voters who rank the candidates in the order given by the column.

Example 2.

9	8	15	16
v	w	x	y
w	v	z	z
x	x	y	v
y	z	w	w
z	y	v	x

Note that v and w are clones in Example 2, as are y and z. Under the Coombs rule, the order in which the candidates are eliminated is x, w, v, z, with y therefore winning. But without z, y would be the first candidate eliminated. Therefore the Coombs rule is not independent of clones.

Example 3 shows that the Coombs rule possesses neither Condorcet consistency nor non-negative responsiveness.

Example 3.

3	2	2
x	z	z
y	y	x
z	x	y

Applying the Coombs rule to Example 3, z is eliminated and x wins. But z is dominant. Thus the Coombs rule is not Condorcet consistent. And if the two voters whose rankings are expressed in the third column change their rankings and put y ahead of x, then x is eliminated and z wins. Thus the rises in the position of y result in y losing, showing that the Coombs rule does not possess non-negative responsiveness.

Borda

Under the Borda rule [2, pp. 59–66], each candidate receives one point for every other candidate ranked below him or her on any ballot. To allow properly for ties in rankings, Black proposed that in addition, each voter lose one point for each candidate ranked above him or her on any ballot [2, p. 62]. The rationale Borda gave for his rule was that each candidate ranked below a given candidate by any voter should add equally to the estimate of the given candidate's merit [2, p. 158].

To see that Example 2 shows that the Borda rule is not independent of clones, it is useful to compute a "matrix of majorities" for the example. The matrix of majorities is an anti-symmetric matrix, of order equal to the number of candidates that voters have ranked, in which the xy^{th} component is the difference between the number of voters who rank x ahead of y and the number who rank y ahead of x. For Example 2, the matrix of majorities is:

	v	w	x	y	z
v	0	2	18	−14	−14
w	− 2	0	18	−14	−14
x	−18	−18	0	16	16
y	14	14	−16	0	2
z	14	14	−16	− 2	0

Under the variation of the Borda rule proposed by Black to take account of ties in rankings, the Borda score of any candidate x is the sum of row x of the matrix of majorities. Thus in Example 2, the Borda winner is y, but without w the winner would be x. Thus the Borda rule is not independent of clones.

The Borda rule has also been faulted for not being Condorcet consistent [2, pp. 60–61].

Black

Black felt that when there was a dominant candidate, that candidate should be chosen, and that otherwise the Borda rule (revised to allow for ties) was appropriate [2, p. 66]. This is the Black rule. Since there is no dominant candidate in Example 2, the Black rule yields the same result as the Borda rule for this example, and therefore the Black rule is not independent of clones.

Condorcet

Condorcet [5, pp. 126–127] made a proposal for a voting rule that for two centuries seemed indecipherably enigmatic (Black [2, pp. 159–178]). However , Young [14] has recently traced through the logic of Condorcet's argument and has deduced that Condorcet proposed a rule equivalent to the following: Score each possible *ordering of the candidates* according to the sum of the upper triangle of the matrix of majorities for that ordering of the row and columns. The winner is the candidate ranked first in the ordering with the greatest score. This rule has been proposed independently by Kemeny [8]. For Example 2, the ordering with the greatest Condorcet score is $yzvwx$, yielding the following matrix of majorities:

	y	z	v	w	x
y	0	2	14	14	−16
z	− 2	0	14	14	−16
v	−14	−14	0	2	18
w	−14	−14	− 2	0	18
x	16	16	−18	−18	0

Thus the winner for Example 2 under the Condorcet rule is y. However, if the clones w and z are eliminated, then the ordering with the greatest Condorcet score is vxy, yielding the following matrix of majorities:

	v	x	y
v	0	18	−14
x	−18	0	16
y	14	−16	0

Thus the Condorcet rule is not independent of clones.

Copeland

Under the Copeland rule, candidates are scored according to the number of other candidates they beat in head-to-head contests [9, pp. 170–171]. From the matrix of majorities for Example 2 it can be seen that under the Copeland rule the winner for this example is y. But without w the result would be a tie between x and y. Since the tie could be resolved in favor of x, the Copeland rule is not independent of clones. Because the Copeland rule calls any cycle among three candidates a tie, it also lacks resolvability.

Dodgson

Dodgson did not actually propose the rule that has been given his name. Rather, he used it implicitly to criticize other rules [2, pp. 227–228]. Dodgson appeared to suggest that any given candidate should be scored (negatively) according to the number of inversions of adjoining candidates in individual rankings that would be needed to make the given candidate dominant. This has been called the Dodgson rule [9, pp. 172–173]. Such "Dodgson scores" are quite cumbersome to compute. A simplification that is almost equivalent is to score each candidate (negatively) according to the sum of the majorities by which he or she is beaten by other candidates. This would be equivalent to Dodgson's suggestion if a tie could be overcome by an arbitrarily small fraction of a vote instead of by a full vote, and if all contending candidates appeared in enough rankings just below the candidates that beat them that they could be made dominant by being advanced above these without having to be advanced above any others to reach these. For the simplified Dodgson rule, the score for any candidate, x, is the sum of the non-positive components of row x of the matrix of majorities, and the winners are the candidates with the greatest (least negative) scores.

For Example 2 under the simplified Dodgson rule, the winner is y. But without z the winner would be v. Thus the simplified Dodgson rule is not independent of clones.

Nanson

Nanson suggested a procedure of successive eliminations in which, in each round, the candidates with below-average Borda scores, calculated with respect to the remaining candidates, would be eliminated [1, p. 95]. Since the average Borda score is 0 in the version of the Borda rule described here, Nanson's rule is equivalent to eliminating in each round the candidates with negative Borda scores, calculated with respect to the remaining candidates.

For Example 2 under the Nanson rule, candidates v, w, and x are eliminated in the first round, and z in the second round, leaving y the winner. But without w the winner would be x. Thus the Nanson rule is not independent of clones. The Nanson rule also lacks non-negative responsiveness [11].

Minimax

The minimax rule, named by Young [13], is one of the suggestions made by Black [2, p. 175] in the effort to decipher Condorcet's proposed voting rule. Under the minimax rule, the score for any candidate, x, is the least (most negative) component of row x of the matrix of majorities, and the winner is the candidate with the greatest (least negative) score.

The minimax rule is independent of "identical twins" (clones that come in pairs). This is because the addition of a clone of a candidate that had no previous clone replicates a row and a column of the matrix of majorities, with only the performance of the two identical twins against each other as new information. At least one of any pair of identical twins has a non-negative score against the other, and therefore a minimax score that is unaffected by the presence of the other twin. And the minimax score of every candidate other than the pair of identical twins is unaffected by the addition of the twin. Therefore the only possible effects of the addition of an identical twin on the set of winning candidates are the replacement of one twin by the other or the inclusion of both instead of one in the set of winning candidates, both of which are consistent with independence of clones.

The minimax rule is not independent of clones that come in sets of three or more. This is shown by Example 4.

Example 4.

6	5	4	5	4	3
w	x	y	z	z	z
x	y	w	w	x	y
y	w	x	x	y	w
z	z	z	y	w	x

The matrix of majorities for Example 4 is:

	w	x	y	z
w	0	9	− 5	3
x	−9	0	13	3
y	5	−13	0	3
z	−3	− 3	− 3	0

In Example 4, w, x, and y are clones. Under the minimax rule, z is the winner, but without w the winner would be x. Thus the minimax rule is not independent of clones.

Young

The Young voting rule came from Young's initial effort [13] to decipher Condorcet's intended voting rule. Under the Young rule, the score of any candidate, x, is the cardinality of the largest subset of rankings for which x is a dominant candidate, and the candidate with the largest score wins. To give his rule homogeneity, Young specifies that any voter's ranking may be divided into fractions of a ranking for the purpose of identifying the largest subset of rankings for which a candidate is dominant.

Like the minimax rule, the Young rule is independent of identical twins. A twin that ties or beats the other has a Young score that is unaffected by the removal of the other twin, and all candidates other than the twins have Young scores that are unaffected by the removal of a twin.

Under the Young rule, z is the winner in Example 4, but without w the winner would be x. Thus the Young rule, like the minimax rule, is not independent of clones that come in sets of three or more.

Summary

Of the 13 voting rules examined, only the alternative vote and the GOCHA rule are independent of clones, and both of these lack other important properties. The alternative vote lacks Condorcet consistency and non-negative responsiveness; the GOCHA rule lacks resolvability. The minimax rule and the Young rule are independent of clones that come in pairs, but they are not independent of clones that come in larger sets. The Coombs, Borda, Black, Condorcet, Copeland, simplified Dodgson, and Nanson rules are not independent of clones. The concept of independence of clones cannot be applied directly to plurality or approval voting, because these voting rules are not based on rankings. However, when rankings are associated with votes under these rules in a straightforward way, they are not independent of clones, although approval voting is independent of "perfect clones," candidates so similar that no voter makes any distinction among them. Both plurality and approval voting lack Condorcet consistency. Table 1 summarizes the properties that are lacking in the rules that have been discussed.

IV. Modifications of Previously Proposed Rules

Is there a simple modification of one of these rules that is independent of clones and also possesses the properties of Condorcet consistency, non-negative responsiveness and resolvability? Because approval voting is not a ranking-based voting rule, it seems unlikely that any modification of it would be independent of clones or Condorcet consistent.

Table 1. Properties that voting rules lack (X indicates the lack of a property)

Rule	Property			
	Condorcet Consistency	Non-negative Responsiveness	Resolvability	Independence of Clones
Plurality[a]	X			X
Approval[a]	X			X
Alternative vote	X	X		
GOCHA			X	
Coombs	X	X		X
Borda	X			X
Black				X
Condorcet				X
Copeland			X	X
Simplified Dodgson				X
Nanson		X		X
Minimax				X
Young				X

[a] When the properties are suitably generalized to take account of the fact that this is not a ranking-based voting rule

Consider the alternative vote, which is independent of clones but does not possess Condorcet consistency or non-negative responsiveness. One could modify the alternative vote in the way that the Black rule modifies the Borda rule, to give it Condorcet consistency: Specify that if there is a dominant candidate then that candidate wins, and only otherwise is the alternative vote used. One could similarly modify some other voting rule to give it independence of clones, by specifying that the first step in counting votes would be to examine the rankings to see if any subsets of the candidates were clones. If any clones were discovered, a contest would be held among the clones first, and all but one clone in any set would be eliminated before proceeding with whatever rule might be desired.

While such devices could be used to develop rules that were Condorcet consistent and independent of clones, the rules would not be satisfying. They would lack what might be called "coherence." This objection may be made more explicit in the following way: When a voting rule is anonymous, as all voting rules examined in this paper are, its domain may be conceived as the space of vectors whose components are the numbers of votes of each permitted type. If a voting rule is also homogeneous, as again all voting rules examined in this paper are, its domain can be projected from the set of vectors whose components are non-negative integers onto the rational subset of the unit simplex. That is, the result may be expressed as a function of the fractions of the votes that are of each permitted type. In this domain, a rule of looking first for clones, or for dominant candidates, divides the domain into a region where the test is met (clones are found, or a dominant candidate is found) and a region where the test is not met. The subsets of the domain that are mapped to different outcomes will therefore generally have ragged edges where the transition occurs between the region where the first test is met and the region where it is not.

These ragged edges may be regarded as evidence that the two-stage allocation of the domain to winning candidates is not appropriate. For this reason it is not satisfying to deal with the issues of clones and dominant candidates by searching for them first.

In looking for a modification of the minimax rule that would be independent of clones, one might think of confining the minimax rule to the set of winners under the GOCHA rule. If the minimax rule is applied in this way, then in Example 4 the winner would be w, and Example 4 would not be a counterexample to independence of clones. However, Example 5 shows that such a rule would not be independent of clones.

Example 5.

7	3	6	3	5	3
v	z	y	w	z	y
w	y	z	x	x	x
x	v	w	v	v	v
y	w	x	z	w	w
z	x	v	y	y	z

The matrix of majorities for Example 5 is:

	v	w	x	y	z
v	0	9	−7	3	−1
w	−9	0	11	3	−1
x	7	−11	0	3	−1
y	−3	−3	−3	0	5
z	1	1	1	−5	0

In Example 5, v, w, and x are clones, and the set of GOCHA winners consists of all five candidates. A minimax rule confined to this set would select y, but if v were withdrawn, the winner would be w. Thus a minimax rule confined to the set of GOCHA winners is not independent of clones.

What is happening in Example 5 is that the clones v, w, and x form a "deep cycle," while each of the clones forms a "shallow cycle" with candidates y and z. This is shown diagrammatically in Fig. 1.

In Fig. 1 the candidates of Example 5 are represented as points, and $v \xrightarrow{9} w$ means that the pairing of v and w yields a majority of 9 for v. The three clones, v, w, and x all lose to z by just one vote, but because of the way that they lose to one another, y wins under a minimax rule confined to the set of GOCHA winners, if all three are present. If a voting rule is to be independent of clones, the depth of the cycle among any clones must not be considered when deciding whether a clone or some other candidate is the winner. And one cannot count on cycles being as tidy as they are in Example 5. A voting rule must be capable of coping with cycles of any complexity among any number of candidates.

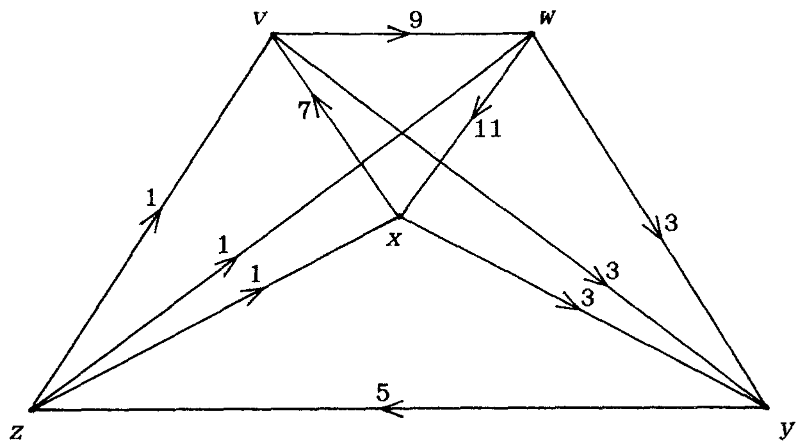

Fig. 1. Pairings for Example 5

V. The Ranked Pairs Voting Rule

The way the ranked pairs rule copes with complex cycles is akin to applying the minimax rule to orderings rather than to candidates. The ranked pairs rule attends not to the most negative component of a *row* of the matrix of majorities, but rather to the most negative component of the *upper triangle* of the matrix of majorities, for a given ordering of the candidates. When two orderings generate the same most-negative upper-triangle component, the second most negative components are examined. Because the treatment of ties that is needed to achieve independence of clones is somewhat different than might be expected, if it is useful to present the rule as an algorithm before elaborating on the rule as a function.

An Algorithm

Start with the pairings decided by the largest and second largest majorities, and require that the orderings they specify be preserved in the final ranking of all candidates. (Require that the ordering of the candidates be such that when the rows and columns of the matrix of majorities are so ordered, the positive expressions of these majorities are in the upper triangle of the matrix, and the negative expressions in the lower triangle.) Seek next to preserve the pair-ordering decided by the third largest majority, and so on. When a pair-ordering is encountered that cannot be preserved while also preserving all pair-orderings with greater majorities, disregard it and go on to pair-orderings decided by smaller majorities. Stop when a unique ranking of all candidates is determined, and declare as the winner the candidate at the top of that ranking.

Consider how this "ranked pairs" rule would be applied to Example 5. The ranking of pair-orderings for Example 5 is shown in Table 2.

Applying the ranked pairs rule to the information in Table 2, the first criterion of the final ranking is that w be ranked ahead of x, and the second is that v be ranked ahead of w. Because a ranking must be transitive, these two taken together

Table 2. The ranking of pair-orderings for Example 5

Rank	Pair-ordering	Majority
1	w over x	11 votes
2	v over w	9 votes
3	x over v	7 votes
4	y over z	5 votes
5 (3-way tie)	v over y	3 votes
	w over y	3 votes
	x over y	3 votes
8 (3-way tie)	z over v	1 vote
	z over w	1 vote
	z over x	1 vote

imply that v be ranked ahead of x, which means that the third-ranked pair-ordering, x over v, cannot be achieved. It is therefore disregarded. The fourth criterion is that y be ranked ahead of z. This is not inconsistent with any higher-ranked criterion, but it leaves unresolved the way that y and z are ranked relative to v, w, and x. This is specified by the three pair-orderings tied for the fifth position: v, w, and x are all ahead of y and z. Thus the final ranking is $vwxyz$, and v is the winner.

The ranked pairs rule may be described as implementing lexicographic preferences for maintaining consistency with pair-orderings decided by larger majorities at the expense of maintaining consistency with pair-orderings decided by smaller majorities. One way of making the rule precise is through the following algorithm: Given r candidates, consider the $r!$ possible rankings of these candidates. Eliminate the rankings that are not consistent with the first and second pair-orderings. When one reaches the third and subsequent pair-ordering, it is possible, as in Example 5, that none of the remaining rankings are consistent with the pair-ordering under consideration. In that case, ignore that pair-ordering and proceeds to the next. Continue until just one ranking of the candidates remains. The winner is the candidate at the top of that ranking.

In the event of ties in the ranking of pairs, proceed as follows: Whenever p of the pairs have the same majority and none of the remaining orderings are consistent with the p pair-orderings implied by these pairs, consider each of the $p!$ ways of breaking the tie in the ranking of pairs. For each way of breaking that and any subsequent ties, there will be a final ranking of the candidates. The election is a tie among all candidates that are at the top of a ranking that the algorithm generates for some way of breaking the ties in the ranking of pairs. If the majority for any pair is 0 and the algorithm proceeds to the point where majorities of 0 would be considered, then all rankings remaining at that point are in the set of tied winning rankings, and all candidates at the tops of such rankings are in the set of winning candidates.

When there is a tie for the winning candidate, choosing among the tied candidates with equal probability would reward the nomination of clones. Ties can be broken in a way that neither rewards nor penalizes cloning, by picking a voter at random and selecting as the sole winner the candidate among those who are tied that is ranked highest by the selected voter. If the selected voter ranks two or more of the tied candidates as tied at the top of his or her ranking of those candidates, use a random process to select among them.

The Ranked Pairs Rule as a Function

The ranked pairs rule is presented above as an algorithm. To examine the properties of the rule it is useful to express it as a function. This can be accomplished as follows:

Map any preference profile to the corresponding matrix of majorities. Given a matrix of majorities, M, and a strict ranking of the candidates, R, let the function F map the pair (M, R) to the greatest M_{xy} such that x is ranked below y in R. Let the function G map the pair (M, R) to the subset of ordered pairs of candidates (x, y) such that x is below y in R and $M_{xy}=F(M, R)$. Define ranking R to dominate ranking R' for matrix M if and only if either $F(M, R)<F(M, R')$ or $[F(M, R)=F(M, R')$ and $G(M, R)\subset G(M, R')]$. Define the set of winners under the ranked pairs rule as the set of candidates that are ranked first in undominated rankings.

To show that this function produces the same outcome as the algorithm for the ranked pairs rule, one must show that:

1. If R is undominated, then R is produced by the algorithm for some way of breaking ties in the ranking of pair orderings, and
2. If R is dominated, then R is not produced by the algorithm for any way of breaking ties in the ranking of pair orderings.

Taking these in order, suppose that R is undominated. Let $W=F(M, R)$. (W is the magnitude of the worst problem with ordering R.) Among pairs $\{x, y\}$ with majorities of W, break any ties in the ranking of pairs by placing pairs with x ahead of y in R ahead of any pairs with y ahead of x and R. Then it can be shown that, given any other ranking R', for the selected way of breaking ties the algorithm cannot eliminate R before it eliminates R'.

Let H be the greatest M_{xy} such that x is ahead of y in R' and not in R. (H is the magnitude of the most powerful challenge of R' to R). It cannot be that $H>W$, because that would contradict the assumption that R is undominated. If $H=W$ and $W>0$, then there must be some $M_{xy}=W$ with x ahead of y in R but not in R', because otherwise R' would dominate R. The algorithm and the breaking of ties are such that the pair $\{x, y\}$ is encountered before any pair with a majority of W for which R disagrees with the majority, so that R' is discarded by the algorithm before there is any occasion to discard R. If $H=W$ and $W=0$, then neither R nor R' will be discarded by the algorithm. If $H=W$ and $W<0$, then R is never discarded. And if $H<W$, then R' is discarded before there is any occasion to discard R. Thus, since R is never discarded before any R', for the selected way of breaking ties, R must be among the final rankings yielded by the algorithm.

Now suppose R is dominated by R'. If $H>W$, then R is discarded before there is any occasion to discard R', so the algorithm cannot choose R. If $H=W$, then it must be that $G(M, R')\subset G(M, R)$, because otherwise the assumption that R is dominated by R' would be contradicted. Since $G(M, R')\subset G(M, R)$, when the algorithm is dealing with pairs with majorities of W it cannot discard R' before it discards R, and it will discard R when it reaches an element of $G(M, R) \setminus G(M, R')$, so the algorithm cannot wind up with R. It cannot be that $H < W$, because then the assumption that R is dominated by R' would be contradicted. Thus R cannot ever be selected by the algorithm if it is dominated. Thus the description of the ranked pairs rule as a function is equivalent to the description of the rule as an algorithm.

Independence of Clones in the Absence of Certain Ties

Define a matrix of majorities, M, to be differentiated apart from clones if no non-diagonal $M_{xy}=0$, and $M_{xy}=M_{vw}$ only if x is either equal to or a clone of v, and y is either equal to or a clone of w. It can be shown that the ranked pairs rule is independent of clones if matrices of majorities are differentiated apart from clones.

Note first that if matrices of majorities are differentiated apart from clones, then no ranking can be undominated if it puts one or more other candidates between two clones. Suppose there was an undominated ranking R containing a sequence $(c, v, \ldots z, d)$, with c and d elements of a set of clones, and $v, \ldots, z$ not elements of that set. Let $m = \text{Max}\ (|M_{dv}|, |M_{dw}|, \ldots, |M_{dz}|)$. From the assumption that M is differentiated apart from clones, $m > 0$, and there is only one candidate such that $|M_{dx}| = m$. Let x be that candidate. Either $M_{dx} = m$ or $M_{dx} = -m$. Suppose first that $M_{dx} = m$. Then consider the ranking R' that leaves the other candidates in their existing order and replaces the sequence above by $(c, d, v, \ldots z)$. R' dominates R because the pairs $M_{dv}, M_{dw}, \ldots, M_{dz}$ are the only ones that the two rankings place in different orders, and R' agrees with the majority for the one of these with the greatest majority. Thus R is dominated. If $M_{dx} = -m$, then replace the second sequence with $(v, \ldots z, c, d)$, and a similar proof holds.

Next, note that if a matrix of majorities is differentiated apart from clones, there must be just one undominated ranking. Suppose that both R and R' are undominated. Let K equal the greatest M_{xy} such that x is ahead of y in one of R and R' and not the other. If there is just one pair ordering such that $M_{xy} = K$, then, when that pair ordering is reached in the algorithm, one ranking will be kept and the other eliminated, so they are not both undominated. The only way that there can be a set $\{M_{x_1 y_1}, M_{x_2 y_2}, \ldots, M_{x_n y_n}\}$ of components of M such that $M_{x_i y_i} = K$, with M differentiated apart from clones, is for the x_i's or the y_i's or both to be clones. And in that event any ranking that agrees with the majority on one of these $M_{x_i y_i}$'s must agree with the majority on all of them, in which case the algorithm will still keep just one of R and R' when it reaches the pairings with majorities of K. Thus there can be only one undominated ranking when the matrix of majorities is differentiated apart from clones.

To show that the ranked pairs rule is independent of clones when matrices of majorities are differentiated apart from clones: Let M be a matrix of majorities, differentiated apart from clones, for the set of candidates D. Let R be the unique undominated ranking for M. Let c be an element of a set of clones, C. Let M^c be the matrix formed by deleting row c and column c from M. Let R^* be the unique undominated ranking of the elements of $C \setminus c$ among themselves. Let R^c be the ranking that places the elements of $C \setminus c$ in the order specified by R^*, and the other candidates, relative to each other and to the elements of $C \setminus c$, in the order specified by R. Suppose some R' dominates R^c. Let R'' be the ranking that places C in the order specified by R and $D \setminus C$, relative to each other and to C, in the order specified by R'. Let K = the greatest M_{xy} such that x is ahead of y in one of R and R' and not the other. Since R^c places the elements of C in their uniquely undominated order, it cannot be that the majority of K arises from a comparison of two elements of C. Since R' has been hypothesized to dominate R^c, it must be that R' agrees with the majority for the pair or pairs with a majority of K. But since R' ranks all pairs not involving two elements of C in the same way that R'' does, and the same is true of R and R^c, in that event R'' would dominate R. Since R is the uniquely undominated ranking for M, that cannot be. Thus there could not have been an R' that dominated R^c, so that R^c is the unique undominated ranking of D for the matrix of majorities M^c. By the construction of R^c, if an element of $D \setminus C$ is first in R, that same element of $D \setminus C$ is first in R^c, and if an element of C is first in R, some element of C is first in R^c. Thus the ranked pairs rule is independent of clones if matrices of majorities are differentiated apart from clones.

Independence of Clones when there are Ties

The assumption that the matrix of majorities is differentiated apart from clones plays an important part in the above proof, and yet there does not seem to be a good reason why the ranked pairs rule should fail to be independent of clones if the matrix of majorities is not differentiated apart from clones. I have searched without success for an example in which independence of clones is not satisfied. I invite others to look for either a proof or a counterexample for the conjecture that the ranked pairs rule is independent of clones even if matrices of majorities are not differentiated apart from clones.

One indication of the complications that arise when a matrix of majorities is not differentiated apart from clones is provided by Example 6, in which y and z are clones.

Example 6:

1	1	1
w	x	z
x	y	y
y	z	w
z	w	x

The matrix of majorities for Example 6 is:

	w	x	y	z
w	0	1	−1	−1
x	−1	0	1	1
y	1	−1	0	1
z	1	−1	−1	0

Since all pairings are decided by the same majority, the pairings may be considered in any order by the ranked pairs algorithm. If the pairs $\{z, w\}$, $\{w, x\}$, and $\{x, y\}$ are considered before any other pairs, then the resulting ranking will be $zwxy$, which shows that a ranking can be undominated even if clones are neither adjacent nor in an undominated order.

Condorcet Consistency

A simplification in the use of the ranked pairs rule arises from the fact that if there is a set of candidates, D, that dominates all others, in the sense that a majority prefers every candidate in D to every candidate not in D, then any undominated ranking puts every element of D ahead of every element not in D. The reason for this is that any ranking that puts an element of D immediately below a candidate that is not in D is dominated by the ranking in which these two are reversed while the positions of all other candidates are maintained. Furthermore, any ranking of all candidates that did not put D in an undominated ranking would be dominated by one that did, and any ranking of all candidates that did put D in an undominated ranking would be undominated if it put the other candidates in the order in which they were found in some undominated ranking. Therefore only the minimum dominant set of candidates need be examined to determine the winner or winners under the ranked pairs rule. This implies immediately that if there is a dominant candidate, that candidate is the sole winner under the ranked pairs rule. Thus the ranked pairs rule is Condorcet consistent.

Non-Negative Responsiveness

Let z be a winning candidate for preference profile P, and assume that: z is ranked higher in the i^{th} ranking of profile P' than in the i^{th} ranking of P, all other candidates are ranked in the same order in the i^{th} rankings of P' and P, and for j not equal to i, all candidates are ranked in the same order in the j^{th} ranking of profile P' as in the j^{th} ranking of profile P. Let M and M' be the matrices of majorities for P and P' respectively. The only differences between M and M' are that in row z some of the components of M' are greater than the corresponding components of M by 1 or 2, and the corresponding components of column z are smaller by 1 or 2. Let m be the greatest M'_{zx} that is greater than the corresponding M_{zx}. Let R be a ranking that puts z first and is undominated for matrix of majorities M. Let R' be any other ranking. By hypothesis, R' does not dominate R for matrix of majorities M. Let K be the greatest M_{xy} such that x is ahead of y in one of R and R' and not the other. If $m > K$, then R dominates R' for matrix M'. And if $m \leq K$, then R is not dominated by R' for matrix M. Thus the ranked pairs rule possesses non-negative responsiveness.

Resolvability

Suppose the set of wining candidates for preference profile P contains more than one element. Let M be the matrix of majorities associated with P, and let R and R' be two undominated rankings for M. Let M' be the matrix of majorities that results when a vote with the ranking described by R is appended to preference profile P. For all ordered pairs (x, y) with x ahead of y in R, $M'_{xy} = M_{xy} + 1$. Let K = the greatest M_{xy} such that x is ahead of y in one of R and R' and not the other. Let T be the set of pairs $\{x, y\}$ with majorities of K that are ranked oppositely by R and R'. Since R and R' are both undominated, there is at least one element of T for which R agrees with the majority. For such pairs, the majority under M' is $K+1$. All other pairs with majorities of $K+1$ or more under M have the same majorities under M'. Therefore for the augmented profile R dominates R'. Therefore the ranked pairs rule is resolvable.

A Conjectured Characterization of the Ranked Pairs Rule

It has been shown that the ranked pairs rule is independent of clones when matrices of majorities are differentiated, and that it possesses the properties of Condorcet consistency, non-negative responsiveness and resolvability. There may be other voting rules with all of these properties, but I conjecture that the ranked pairs rule is the only voting rule with these properties that is a function of the matrix of majorities and possesses the properties of anonymity and neutrality.

One reason this characterization may be interesting is that, for choices between pairs of options, majority rule can be characterized as the only voting rule that possesses anonymity, neutrality, non-negative responsiveness and resolvability [11]. Thus if the conjectured characterization holds, it is possible to single out the ranked pairs rule as the only ranking-based voting rule that, as a binary rule, satisfies this characterization of majority rule, and also is (1) a function of the matrix of majorities, (2) Condorcet consistent, and (3) independent of clones.

VI. Appraisal

Is the ranked pairs rule a voting rule that collectivities would reasonably want to use? How do its strengths and weaknesses compare, overall, with those of other voting rules?

The failure to achieve independence of clones is quite serious for plurality because of vote splitting and for the Borda rule because of the positive strategic value of including clones on the ballot. Including clones on the ballot also has positive strategic value under the Condorcet, Copeland, simplified Dodgson, Nanson and Black rules, although the seriousness of the lack of independence of clones for these rules is reduced by their Condorcet consistency. The alternative vote and the Coombs rule lack not only Condorcet consistency but also non-negative responsiveness.

It may not be appropriate to criticize approval voting for not being Condorcet consistent, since approval voting is not a ranking-based voting rule. The most serious flaw in approval voting is probably that it requires voters to decide whether to use their votes to attempt to discriminate among the likely winners or to help ensure that only truly qualified candidates will win. This is an unfortunate choice to impose on voters.

The GOCHA rule, in a sense, is only half a voting rule. It does not address the issue of what should be done to resolve cycles. Young's rule, as he acknowledges, is extremely complex to calculate. This disposes of all of the competing rules that have been mentioned except the minimax rule.

The minimax rule and the ranked pairs rule share the property of "lexicographic preference" for satisfying a single pair-ranking with a majority of m, even if the cost is not satisfying any number of pair rankings with majorities of $m-1$. The ranked pairs rule merely does this in a more systematic fashion. While this lexicographicness might seem inappropriate, it seems to be required for achieving independence of clones, if a voting rule is to be a function of the matrix of majorities.

The only instances in which the minimax rule has been shown to be not independent of clones are ones in which a cycle among clones is either embedded within another cycle or combined with the clones beating another candidate by a smaller majority than the weakest link in their cycle. These seem like events with extremely low probabilities. It may be that for all but the most sophisticated electorates, the small sacrifice of full independence of clones from using the minimax rule rather than the ranked pairs rule would be worth the saving in additional complexity of the voting rule. However, if full independence of clones is desired, the ranked pairs rule can, at least "almost always," provide it.

References

1. Arrow KJ (1963) Social choice and individual values. Yale University Press, New Haven
2. Black D (1958) The theory of committees and elections. Cambridge University Press, Cambridge
3. Brams SJ, Fishburn PC (1978) Approval voting. Am Pol Sci Rev 72:831–47
4. Chaudhuri A, Roy D (1984) A nonproliferation criteria for a good voting scheme. Hooghly Mohsin College, Chinsurah, W.B., India
5. Condorcet (1785) Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix. Paris
6. Coombs CH (1964) A theory of data. Wiley, New York
7. Doron G, Kronick R (1977) Single transferable vote: an example of a perverse social choice function. Am J Pol Sci 21:303–11
8. Kemeny J (1959) Mathematics without numbers. Daedalus 88:571–91
9. Fishburn PC (1973) The theory of social choice. Princeton University Press, Princeton
10. Schwartz T (1986) The logic of collective choice. Columbia University Press, New York
11. Tideman TN (1986) A characterization of majority rule with multiple extensions. Soc Choice Welfare 3:17–30
12. Tideman TN, Tullock G (1976) A new and superior process for making social choices. J Pol Econ 84:1145–59
13. Young HP (1977) Extending Condorcet's rule. J Econ Theory 16:335–53
14. Young HP (undated) Condorcet's rule. School of Public Affairs, Univeristy of Maryland, College Park
15. Zavist T (1986) Toward consensus. 2415 Tangley, #10, Houston, Texas 77005, USA