The cardinal-weighted pairwise comparison method

by James Green-Armytage*

Abstract. This paper introduces a new Condorcet-efficient voting method that uses ordinal information to determine the direction of pairwise defeats, and cardinal information to determine the strength of pairwise defeats. In this paper, I attempt to explain why Condorcet-efficient methods are desirable, identify the areas in which heretofore-proposed Condorcet-efficient methods could be improved, define the cardinal-weighted pairwise method, provide an example computation, discuss the rationale behind the method (with an emphasis on its resistance to strategic manipulation), briefly list criteria passed and failed, and suggest an application of the cardinal-weighted pairwise principle to CPO-STV.

1. Why majoritarian election methods should be Condorcet-efficient

The Condorcet criterion seems to be the most complete and satisfying definition of majority rule that is available to us. First, if there is one candidate who is preferred by some majority over every other candidate individually, it seems inappropriate to call anyone else a majority winner. For example, if candidate A is a Condorcet winner, and a non-Condorcet-efficient method elects candidate B, a majority will prefer A to B. If there was an election just between these two candidates, A should be expected to win that election.

Second, Condorcet-efficient methods tend towards the political center. For example, if options are neatly arrayed along a one-dimensional spectrum, and voters' preferences are single-peaked, Condorcet-efficient methods will always choose the candidate(s) who are closest to the median voter(s) position on the spectrum.[1] Because of this centripetal property, Condorcet methods, to a greater extent than other majority or plurality methods, should promote stable compromises rather than unstable, off-center outcomes. In turn, this should decrease the polarization of political landscapes, along with the counterproductive rancor that can be caused by such polarization.

Third, if a sincere strong Condorcet winner exists in a particular election, it seems that (using most single-winner methods) only the selection of this candidate can produce a strategic equilibrium such that no group of voters can mutually benefit by strategically altering their ballots. If such a candidate exists, then most plurality or majority methods (including approval

* I thank Nicolaus Tideman, Jobst Heitzig, Chris Benham, and Markus Schulze for their helpful comments on earlier versions of this paper. I am also grateful to many other members of the election methods discussion list for their insight and support.

voting and instant runoff voting) can only reach equilibrium with the selection of this candidate.[2] Hence the election of a Condorcet winner is the most stable result.

Fourth, Condorcet-efficient methods minimize the incentive for the **compromising strategy**, which is “insincerely ranking an alternative higher in the hope of getting it elected.” (Cretney, 2004) For example, if my sincere preferences are R>S>T, a compromising strategy would be to vote S>R>T or R=S>T, in order to improve S’s chances of winning, often at the expense of decreasing R’s chances of winning. All resolvable voting methods that satisfy the majority criterion have a compromising incentive when there is no Condorcet winner. (Schulze, 2004a) But unlike other methods, such as instant runoff voting, voters in Condorcet-efficient methods never have an incentive to use the compromising strategy when there is a Condorcet winner. (Schulze, 2004b) This is an important property because, in the absence of a majority rule cycle, it allows me to vote my R>S preference without worrying that it will undermine my S>T preference. This is a more complete way of curtailing the “lesser of two evils” problem, that is, decreasing the extent to which voters have to worry about earlier choices drawing support away from later choices. Thus, Condorcet-efficient methods allow more candidates to participate on an equal basis, which should lead in turn to a substantially higher level of responsiveness and accountability.

2. Remaining room for improvement in Condorcet-efficient methods

2.a. Considering the relative priority of preferences

Most Condorcet efficient methods that have been proposed so far limit voter input to ordinal rankings. Hence, voters can express preferences between candidates, but they cannot express the relative priority of their preferences. If I worship my first three choices, but detest my fourth and fifth choices, I cannot express this on my ballot, and it is not taken into account when deciding the winner.

Imagine a case where there is a majority rule cycle such that A pairwise beats B, B pairwise beats C, and C pairwise beats A. Imagine that the numerical majority which agrees with the A>B defeat is the smallest of the three, but most of those who prefer A over B consider it to be their most crucial ranking, whereas the C>A defeat is supported by a slightly larger majority, but most of those who prefer C over A are relatively indifferent between the two candidates, and consider other rankings to be far more important. I am interested in developing a system that would drop the C>A defeat rather than the A>B defeat. That is, I am interested in a method that measures defeat strength by taking into account the relative priority of each pairwise preference to each voter.

Cardinal ratings make sense as a way to get information about the relative priority of voter preferences, because they have high resolution while being easy for voters to understand. However, if we are interested in majority rule, it will not suffice to use the simple cardinal ratings method that just takes the sum of cardinal scores to find the winner. This method (as well as many others which resemble it) does not pass the mutual majority criterion, the majority criterion, or the Condorcet criterion,[3] and it is severely vulnerable to a strategic situation known as the cooperation/defection dilemma.[4] So, the challenge is to use cardinal ratings in a subtle way that does not compromise the integrity of the pairwise method.

2.b. Resistance to strategic manipulation

Although Condorcet-efficient methods minimize the incentive for use of the compromising strategy, they are vulnerable to the **burying strategy**. This is defined as "insincerely ranking an alternative lower in the hope of defeating it." (Cretney, 2004) For example, if my sincere preferences are R>S>T, a burying strategy would be to vote R>T>S or R>S=T, in order to hurt S and thereby improve R's chances of winning. For example, imagine that the sincere preferences of the electorate are as follows:

Example 2.1

23: A>B>C[*]	23: B>A>C	25: C>A>B
5: A>C>B	2: B>C>A	22: C>B>A

Given a sincere vote, the pairwise comparisons would be as follows:

A>B : 53-47	A>C : 51-49	C>B : 52-48

A is a sincere Condorcet winner. However, the C>A>B voters have an opportunity for strategic manipulation. If they "bury" candidate A, they can create a false defeat of B over A, which causes a majority rule cycle such that A's defeat of C is dropped in order to resolve the cycle. So, in essence, they can use a false defeat to overrule the genuine defeat of C by A, as follows:

Example 2.2

23: A>B>C	23: B>A>C	19: C>A>B

[*] I intend this notation to indicate that 23 voters prefer A to B and B to C.

5: A>C>B 2: B>C>A 28: C>B>A (6 of 28 are sincerely C>A>B)

Pairwise comparisons:

B>A : 53-47 A>C : 51-49 C>B : 52-48

There is a majority rule cycle now, and A's defeat of C is supported by the fewest votes, so minimax, ranked pairs, and beatpath choose C.

The burying strategy may have the potential to cause substantial trouble in elections that use a Condorcet-efficient method. Some have cited this as a reason not to adopt Condorcet-efficient methods. (Monroe, 2001; Richie and Bouricus, 2004) Unfortunately, Condorcet-efficient methods cannot be completely invulnerable to the burying strategy[5], but perhaps we can identify a Condorcet-efficient method that prevents the most flagrant strategic incursions, that counterbalances strategic incentive against strategic ability, that minimizes strategic disincentives for new candidates to enter the race, and that provides for relatively stable counter-strategies. I will return to this issue in part 6.

3. Definition of the cardinal-weighted pairwise comparison method

3.a. Ballot

1. Voters rank the candidates. Equal rankings are allowed.

2. Voters rate the candidates, e.g. on a scale from 0 to 100. Equal ratings are allowed. If you give one candidate a higher rating than another, then you must also give the higher-rated candidate a higher ranking.[6]

3.b. Tally

1. Pairwise tally, using the rankings information only. Determine the *direction* of the pairwise defeats by using the *rankings* for a standard pairwise comparison tally.

2. Determine the *strength* of the pairwise defeats by finding the weighted magnitude as follows. Let's say that candidate A pairwise beats candidate B, and we want to know the strength of the defeat. For each voter who ranks A over B, and *only* for voters who rank A over B, subtract their rating of B from their rating of A, to get the rating differential. The sum of these individual winning rating differentials is the total weighted magnitude of the defeat. (Note that voters who rank B over A, or rank them equally, do not contribute to the weighted magnitude; hence it is never negative.)

3. Now that the direction of the pairwise defeats have been determined (in step 1) and the strength of the defeats have been determined (in step 2), you can choose from a variety of Condorcet completion methods to determine the winner. I recommend the ranked pairs, beatpath, and river methods.

3.c. Optional, additional provisions

These additional provisions are not an essential part of the cardinal-weighted pairwise method, but they may prove helpful.

1. Maximizing in scale provision:[7] Once a Schwartz (GOCHA) set has been established by the pairwise tally in step 2, it may be a good idea to maximize the voters' rating differentials in scale between the candidates in the set. That is, to change the ratings on each ballot so that the highest-rated Schwartz set candidate is at 100, the lowest-rated Schwartz set candidate is at 0, and the rating differentials between the Schwartz set candidate retain their original ratios. (For example, 50,20,10 would become 100,25,0.) The benefit of this provision is that voters will have equal ballot weight with regard to the resolution of the majority rule cycle in particular. Therefore, voters will not have an incentive against investing priority in sincere gaps that are relatively unlikely to fall within the Schwartz set.

2. Blank ratings option: This allows voters to give one or more candidates a blank rating, such that if I give some candidate a blank rating, my ballot will still affect the direction of pairwise defeats concerning that candidate, but it will not add to the weighted magnitude of such defeats. In the absence of this option, the ratings for unrated candidates on a given ballot would be assigned to candidates according to a default formula.[8]

4. An example computation

Example 4.1

26: Right 100 > Left_B 10 > Left_A 0[*]	22: Right 100 > Left_A 10 > Left_B 0
26: Left_B 100 > Left_A 90 > Right 0	1: Left_B 100 > Right 50 > Left_A 0
21: Left_A 100 > Left_B 90 > Right 0	4: Left_A 100 > Right 50 > Left_B 0

Direction of defeats (using rankings information)

[*] I intend this notation to indicate that 26 voters rank the candidates in the order Right > Left_B > Left_A, and assign the three candidates ratings of 100, 10, and 0, respectively.

Right > $Left_B$: 52-48 $Left_A$ > Right: 51-49 $Left_B$ > $Left_A$: 53-47

Weighted magnitude of defeats (using ratings information)

Right > $Left_B$: (26x(100-10)) + (22x(100-0)) + (4x(50-0)) = 4740

$Left_B$ > $Left_A$: (26x(10-0)) + (26x(100-90)) + (1x(100-0)) = 620

$Left_A$ > Right: (26x(90-0)) + (21x(100-0)) + (4x(100-50)) = 4640

Completion by cardinal-weighted pairwise with **minimax**: Right's worst loss has a weighted magnitude of 4640. $Left_B$'s worst loss has a weighted magnitude of 4740. $Left_A$'s worst loss has a weighted magnitude of 620. $Left_A$'s worst loss is the least bad; *$Left_A$ wins*.

Completion by cardinal-weighted pairwise with **ranked pairs** or **river**: Consider the defeats in the order of descending weighted magnitude.

4740: Right > $Left_B$ keep

4640: $Left_A$ > Right keep

620: $Left_B$ > $Left_A$ skip (would cause a cycle, Right>$Left_B$>$Left_A$>Right)

Kept defeats produce ordering $Left_A$>Right>$Left_B$; *$Left_A$ wins*.

Completion by cardinal-weighted pairwise with **beatpath**: The strength of a beatpath is defined by the defeat along that path with the lowest weighted magnitude.

beatpath Right → $Left_B$: 4740 beatpath $Left_B$ → Right: 620

beatpath $Left_A$ → Right: 4640 beatpath Right → $Left_A$: 620

beatpath $Left_A$ → $Left_B$: 4640 beatpath $Left_B$ → $Left_A$: 620

Complete ordering is $Left_A$>Right>$Left_B$; *$Left_A$ wins*.

5. Defeat strength

For the sake of convenience in the following sections, allow me to use the shorthand terms "ordinal pairwise" and "cardinal pairwise". "Ordinal pairwise" will refer to minimax, ranked pairs, or beatpath, using winning votes or margins. "Cardinal pairwise" is a shorter name for cardinal-weighted pairwise.

Ordinal pairwise measures defeat strength in terms of a sheer number of ballots. Cardinal pairwise attempts to extend the sensitivity of the method by factoring in a measure of how much priority the voters assign to each defeat. The goal is that the weakest defeat in a majority rule cycle should be the one which has the lowest overall combination of these two factors: 1. the number of voters in agreement with the defeat; 2. the relative priority of the defeat to those voters who agree with it.

Example 4.1 above illustrates a type of situation where these two approaches can produce a different result. There are 53 voters who agree with the $Left_B$>$Left_A$ defeat, which is a relatively large number. However, let's consider the ratings expressed by some of these voters. The 26 voters who indicate $Left_B$ 100 > $Left_A$ 90 > Right 0 express three pairwise preferences; these are $Left_B$>$Left_A$, which spans a gap of 10 points, $Left_B$>Right, which spans a gap of 100 points, and $Left_A$>Right, which spans a gap of 90 points. This is intended to indicate that they consider the $Left_B$>$Left_A$ defeat to have the lowest priority among their pairwise comparisons, and that its priority is lower than the others by a substantial margin. In this way, the 26 voters who indicate Right 100 > $Left_B$ 10 > $Left_A$ 0 also place the lowest priority on their $Left_B$>$Left_A$ preference. In light of the fact than 52 of the 53 voters who agree with the $Left_B$>$Left_A$ defeat consider it to be their lowest priority, it shouldn't be surprising that cardinal pairwise assigns to it a much lower weighted magnitude than the other two defeats.

It seems almost axiomatic that, when faced with a majority rule cycle, one should drop the defeat(s) in the cycle that are of least importance to the voters. In doing so, you will minimize the extent to which the supporters of a defeat will be offended by its being overruled. So, the remaining question is how to define the priority of each defeat to each voter, and how to aggregate these individual priorities.

The answer that cardinal pairwise gives to this question is relatively simple. For those who agree with a defeat, we look at the rating differential they expressed between the two candidates being compared. Then we take the sum of these winning rating differentials to find the overall strength of the defeat.

The idea is that the voters will give the candidates ratings such that the rating differentials between pairs of candidates will reflect the relative priority of the preferences between those candidates. The fact that each voter is constrained to the same range of ratings (e.g. 0 to 100) assures that everyone has essentially the same voting "power." The point here is not to do interpersonal comparison of utilities, but to allow voters to prioritize their own preferences relative to one another, using a fluid and simple high-resolution scale.

When learning the cardinal pairwise method, one may wonder why it only looks at the rating differentials of those who agree with a particular defeat, rather than subtracting the losing rating differentials from the winning rating differentials. To begin with, I will say that I am more interested in dropping the defeats which are the least important to voters overall, rather than the defeats which are the closest in terms of the strength of preference on either side. That is, if there is one pairwise comparison that is a very high priority for voters on both sides, I think that it is especially important not to reverse this defeat, even though the rating differentials against the defeat may be equal to or greater than those in support of the defeat. Such high-priority defeats should be regarded as crucial within the election, and the cardinal aspect of the method should be used to defend them rather than to undermine them.

In this way, looking at only the winning rating differentials greatly improves the *stability* of the cardinal pairwise method. Because the defeats which voters place the highest priority on are the most difficult to reverse, the cardinal pairwise method is exceptionally resistant to strategic manipulation. This point will be explored in greater detail in the next section.

6. Strategic manipulation

There are many reasons to think that cardinal pairwise will be more resistant to strategy than most other Condorcet-efficient methods. First, it should tend to prevent the most flagrant strategic incursions. Second, it should tend to balance strategic incentive against strategic ability, so that those who would most like to change the result via strategic incursion should tend to be those who are least able to do so. Third, it should minimize strategic barriers against the entry of new candidates. Fourth, it should create the possibility of more stable counterstrategies than those that are available in ordinal pairwise.

6.a. Flagrant strategic incursions

I define a flagrant strategic incursion as one that causes a very high-priority defeat to be overruled by a false defeat. Here is an example of a flagrant incursion:

Example 6.1: Sincere votes

46: A 100 > B 10 > C 0	44: B 100 > A 10 > C 0
5: C 100 > A 50 > B 0	5: C 100 > B 50 > A 0

A is a Condorcet winner. Clearly, the primary contest is between A and B, as C is opposed by a mutual majority of 90% of the voters. However, using ordinal pairwise, the B voters can use C’s candidacy to alter the outcome of the election, by burying A under C, that is, by voting B>C>A.

Example 6.2: Altered votes

46: A 100 > B 10 > C 0	44: B 100 > C 85 > A 0 (sincerely B 100 > A 10 > C 0)
5: C 100 > A 50 > B 0	5: C 100 > B 50 > A 0

Pairwise comparisons:

A > B : 51-49 C > A : 54-46 B > C : 90-10

Amazingly, B is now the winner in ordinal pairwise, which considers the A>B defeat to be the weakest of the three. The B voters have effectively used C as a "weapon" against A.

It is impossible for a flagrant incursion of this particular type to succeed within the cardinal pairwise method. To illustrate the reason for this, let's look at the weighted magnitudes of the three defeats in example 6.2:

A > B : 4390 C > A : 4490 B > C : 1120

The B>C defeat is now by far the weakest, so C is elected, and the B voters' strategy has completely backfired. Furthermore, even with perfect information and coordination, there is no way for the B voters to make their strategy a success. The reason is that if they spend more of the weight of their ballots on the B>C defeat, then they will have less weight to apply to the C>A defeat. The C voters already apply 500 points of weight toward C>A, and the A voters apply 460 towards B>C, but the rest has to come from the B voters. The total weight of their ballots is a finite number, 4800, and the same weight cannot be applied towards C>A and B>C simultaneously; it must be distributed among them. To bring both defeats to the level of the A>B defeat would take a total weight of (4390-500)+(4390-460) = 7820, so 4800 is insufficient, and the B voters do not have a strategic opportunity.

In general, with 100 voters and a ballot with ratings from 0-100, if there is any defeat that has a weight of more than $3333^1/_3$ (that is, $^1/_3$ of the highest possible weighted magnitude), there is absolutely no way for that defeat to be the weakest in a three-candidate cycle. In a three-candidate cycle, no point of weight can count towards more than one of the defeats, and so it would be impossible for all three defeats to have more than $^1/_3$ of the maximum weight. Therefore, when any candidate suffers a defeat with a weighted magnitude greater than $3333^1/_3$, those who disagree with that defeat cannot overrule it via the creation of a three-candidate cycle.

With larger cycles (four candidates and above, e.g. A>B>C>D>A), the $3333^1/_3$ limit does not apply, but the principle of finite weight remains in effect. Let's say that there is a candidate B, who is pairwise-beaten by a candidate A. In order for B to win, if the beatpath method is used in step 3 of the tally, there must be a chain of defeats from B to A (e.g. B>C>D>A), such that every defeat along that chain has a weighted magnitude that is greater than the A>B defeat. The minority who prefer B to A will have a finite amount of weight to distribute along the B>C>D>A chain. A given point of weight can count towards two defeats in this four-candidate chain (e.g.

the gap in B 1 > D 1 > C 0 > A 0 counts towards B>C and D>A), but it cannot count towards more than two.[9]

Cardinal pairwise, unlike ordinal pairwise, does not allow a voter to apply the maximum weight to all of their pairwise preferences; there is a certain kind of scarcity of weight which forces the choice as to which preferences should receive a higher priority. What we are seeing here is that this scarcity produces excellent anti-strategic effects, by placing a limit on the extent to which a strategizing group of voters can build up the weight of multiple pairwise defeats at the same time, in order to manipulate the overall result.

In general, flagrant incursions are much less likely to be successful in cardinal pairwise than in ordinal pairwise, because the chances of overruling an A>B defeat decrease as more voters assign a higher priority to the A>B defeat. I hope that my definition of a flagrant incursion can be seen to have value, and that it can be agreed upon that relatively high-priority defeats should be harder to overrule. Consider that when a defeat of A over B is given a very high priority, we can generally expect B to be very *different* from A (in the eyes of the voters), relative to differences with the other candidates in the election. In order to quantify this difference, we can look at both the average A>B rating differential and the average B>A rating differential for individual voters.

I think it is crucial that we make it as difficult as possible for strategic voters to alter an election result in such a way that the actual winner is considered by the voters to be extremely different from the sincere Condorcet winner, if there is one. Consider how seriously it would undermine the legitimacy of the voting system, if it was found that partisan supporters had pulled off a successful burying strategy which won the election for a candidate who was the ideological polar opposite of the sincere Condorcet winner. Ordinal pairwise unfortunately cannot offer much protection against this disturbing possibility, but cardinal pairwise can.

For example, if an electorate is very heavily divided between the Liberal and Labor parties (such that most voters place their largest ratings gap between the Liberal and Labor candidates), and there is a set of Liberal candidates who pairwise beat every Labor candidate (as well as any candidates from neither party), then the winner of a cardinal pairwise election should nearly always be a member of such a set. In ordinal pairwise methods, Labor voters could gain an advantage by burying Liberal candidates from this set under other Liberal candidates, but in cardinal pairwise, that is unlikely to work, because the defeats from one Liberal candidate to another are likely to be relatively weak, meaning once again that the Labor voters' strategy will probably either fail or backfire.

6.b. Strategic incentive and strategic ability

There are impossibility theorems that state that strategic manipulation cannot be completely avoided in any reasonable election method (Gibbard, 1973; Satterthwaite, 1975; Hylland, 1980), but I'm not aware of a theorem that says that we can't find a method that distributes strategic ability in roughly inverse proportion to strategic incentive.

Let's say that a sincere winner, A, pairwise beats another candidate, B. If the B>A voters place a high priority on their B>A preference, then the average B>A rating differential will be high. However, we should also expect the average A>B rating differential to be high. In general, it is reasonable to expect that the average rating differentials on either side of a defeat will tend to be strongly correlated with one another, because the perceived intensity of difference between candidates should be largely independent of a voters' ranking of those candidates.[10]

So, if the B>A voters' preference for B is strong, the average A>B rating differential is likely to be high. If the average A>B rating differential is high, the weighted magnitude of the A>B defeat will also be high. Finally, if the weighted magnitude of the A>B defeat is high, it will be difficult to overrule, and it is unlikely that a burying strategy by the B>A voters will succeed. Hence, cases where the voters have a strong incentive to bury A for the benefit of B will also tend to be cases where such a strategy is unlikely to work.

On the other hand, if there is a candidate C, such that C also has a pairwise loss to A, but such that the average A>C rating differentials are relatively small, it may be more feasible for supporters of C to overrule the A>C defeat using a burying strategy. However, conversely to the situation with B and A above, we can expect that the average C>A rating differential will be relatively small, and that the C>A voters will not place a high priority on the preference. If so, it is likely that A and C will be somewhat allied with each other, in that the A and C voters will be relying on each other's later preferences to beat other candidates (like B, perhaps). Hence, C supporters may have more of an opportunity than the B supporters to steal the election from A, in that the C>A defeat may be easier to overrule, but they will probably have less to gain by doing so, and more to lose if the strategy backfires.

6.c. Minimizing strategic barriers to candidate entry

In example 4.1 above, Left_B and Left_A can be considered to be relatively similar candidates, in that there is a low average rating differential placed on the comparison between them, going in both directions. If only Left_A and Right were candidates, Left_A would probably win, since he has a pairwise win over Right. In cardinal pairwise, the entry of Left_B does not change this result. However, the winner will change to Right in ordinal pairwise, which defines his 49-51 pairwise loss as the weakest in the cycle. In general, it is much harder in cardinal pairwise for the entry of a new, non-winning candidate to do harm to a similar candidate. The reason for this is that if the new candidate beats the similar candidate, but does not win, this defeat will be relatively likely to be overruled in the event of a cycle.

In ordinal pairwise, a voter who would otherwise support a potentially-entering candidate might have some anxiety that this candidate could hurt a similar candidate whom that voter also supports. Because the potentially-entering candidate's support base may feel ambivalent about his presence in the race, entry of the candidate may not occur. Thus, the method retains a certain strategic barrier to entry of new candidates. Cardinal pairwise minimizes this barrier to entry, in that the entry of a new candidate is extremely unlikely to affect the result in retrograde to the will of his would-be supporters.

6.d. Stable counterstrategies

If a group of voters try to coordinate a strategic incursion, and other voters learn about this and consider it to be undesirable, they may attempt to coordinate a counterstrategy, in order to make the initial strategy unsuccessful. One hopes that counterstrategy will rarely or never be needed, but it is nevertheless to the credit of cardinal pairwise that it provides for somewhat more-stable counterstrategies than ordinal pairwise. Actually, the availability of stable and effective counterstrategies may be a part of the balance that prevents strategic incursion from arising in the first place. There are two basic counterstrategy replies to the burying strategy, and I will illustrate them both as potential responses to a single example.

Example 6.3: Sincere votes

28: A 100 > B 60 > C 0	23 C 100 > A 40 > B 0
27: B 100 > A 60 > C 0	22: C 100 > B 40 > A 0

A is a sincere Condorcet winner, but supporters of B can gain an advantage using the burying strategy.

Example 6.4: Altered votes

28: A 100 > B 60 > C 0 23 C 100 > A 40 > B 0

17: B 100 > A 60 > C 0 22: C 100 > B 40 > A 0

10: B 100 > C 100 > A 0 (sincerely B 100 > A 60 > C 0)

Pairwise comparisons

A > B : 51-49 C > A : 55-45 B > C : 55-45

Using ordinal pairwise, the weakest defeat is A>B, and therefore B wins. B also wins with cardinal pairwise, with weighted magnitudes as follows:

A > B : 2040 C > A : 4580 B > C : 3380

So, in example 6.4, a successful strategic incursion has taken place. However, if the other voters become aware of the strategizing voters' intentions ahead of the election (which would be likely in a large electorate, since strategic coordination on that scale would be difficult to keep a secret), they have the option of engaging in counterstrategy. The two counterstrategies that are possible in this situation can be termed the compromising counterstrategy and the deterrent/burying counterstrategy.

In ordinal pairwise, the **compromising counterstrategy** would entail the C>A>B voters weakening or reversing the defeat against A by voting C=A>B. Starting from example 6.4, it would take 5 such equalizations to make C>A the weakest defeat using a winning votes method, and 9 equalizations using a margins method. This would result in the election of A; a countering of the initial strategic incursion.

In cardinal pairwise, the method above could still be used, but the C>A>B voters also have an opportunity to get the same result by raising A's rating up to 100, hence weakening the C>A defeat and strengthening the A>B defeat. If all 23 of the C>A>B voters were to follow this strategy, by voting C 100 > A 100 > B 0, and the rest of example 6.4 remained unchanged, then A would win, with the following pairwise defeats and weighted magnitudes:

A > B : 3420 C > A : 3200 B > C : 3380

The cardinal pairwise counterstrategy is more stable than the ordinal pairwise counterstrategy, in that it does not risk a change in the winner of the A-C pairwise comparison. This makes it a less perilous choice for the C>A>B voters.

The **deterrent/burying counterstrategy** would entail the A>B>C voters weakening or reversing B's defeat of C, such that the B>A>C voters' burying of A could only backfire by electing C. In ordinal pairwise, this would require some A>B>C voters to equalize or reverse the

B>C preference, thus voting A>B=C or A>C>B. Again, starting from example 6.4, it would take 5 equalizations with winning votes, and 9 equalizations with margins.

In cardinal pairwise, it is possible for the A>B>C voters to get a similar deterrent effect by voting A 100 > B 0 > C 0. For example, if all A>B>C voters were to make this change, and the other preferences remained constant with example 6.4, C would win, with defeats and weighted magnitudes as follows:

A > B : 3720   C > A : 4580   B > C : 1700

With the deterrent/burying counterstrategy in general, the counterstrategizers are unlikely to know for sure whether the original strategizers will carry out their incursion or not, until the votes have already been cast. Therefore it is important to have an effective counterstrategy that they can use without severely destabilizing the result, in case the original strategy is not carried out. In this respect, the cardinal pairwise version of the counterstrategy is preferable, in that it does not alter the direction of any pairwise defeats, and therefore will not interfere with the identification of a Condorcet winner.

Of course, the existence of more-stable counterstrategies in cardinal pairwise does not mean that strategy will never be a problem. However, it suggests to me that the threat of a strategic incursion, should it arise, is less likely to spiral out of control.

7. Satisfied and unsatisfied criteria

For the sake of convenience in this section alone, assume that the method being evaluated is cardinal pairwise, completed by beatpath (in step 3 of the tally), without any of the additional provisions. Assume also that beatpath ties are resolved via the use of a tie-breaking ranking of the options.

The method satisfies the mutual majority criterion, the Condorcet criterion, the Schwartz GOCHA criterion, the Pareto criterion, the resolvability criterion, and the monotonicity criterion. (Green-Armytage, 2004)

The method does not satisfy criteria that are incompatible with Condorcet efficiency, such as participation (Moulin, 1988), consistency (Young, 1975), later-no-harm, and later-no-help (Woodall, 1997). Also, it does not satisfy independence of clones, unless the definition of the criterion is adapted to require that candidates must be given the same rating as one another in order to qualify as clones. Also, the maximizing in scale provision can cause the method to fail monotonicity. (Green-Armytage, 2004)

8. Application to CPO-STV

Comparison of pairs of outcomes by the single transferable vote, or CPO-STV, is a successful integration of the pairwise comparison principle with the STV proportional representation principle. While ordinary STV methods reduce to instant runoff voting in the single-winner case, CPO-STV reduces to a pairwise comparison method; hence it extends the benefits of pairwise analysis to proportional representation. The following is my attempt to apply the cardinal-weighted pairwise principle to CPO-STV.

For each comparison of outcomes, the options within the outcomes should be divided into **common options** and **contested options**. Common options are in both outcomes, and contested options are only in one outcome or the other.

When you get to the final stage of the STV tally for a given outcome comparison, voters have their votes (or fractions thereof) invested in different places. When calculating the weighted magnitude of pairwise defeats among outcomes, look only at the ballots, or the fractions thereof, that are invested in contested options. The ballot fractions invested in contested options within the winning outcome are the ones that we will look at to determine the weighted magnitude.

When comparing two outcomes, define the rating of an outcome for a given voter as the rating they give to their most-preferred contested option.[11] Define the overall rating differential as the difference between the overall ratings of the outcomes being compared. Define the weighted magnitude of the defeat as the sum of the overall ratings differentials expressed by ballot fractions that are invested in a contested option within the winning outcome.

Example 8.1: 3 seats to be decided. 400 voters. 6 candidates: A, B, C, X, Y, and Z. Newland-Britton quota = 400÷(3+1) = 100 votes. The votes are cast as follows:

150: A 100 > B 100 > C 100 > X 0 > Y 0 > Z 0

60: A 100 > B 90 > C 80 > X 0 > Y 0 > Z 0

190: X 100 > Y 90 > Z 80 > A 0 > B 0 > C 0

We know from a regular CPO-STV count that the outcome {AXB} is a Condorcet winner, and hence the relative strength of the defeats will not come into play, but just for example, we will calculate the winning rating differential of {AXB} over {AXY}. Here is the STV-tally component of the outcome comparison:

	A	X	B	Y
Initial count	210	190	0	0

Transfer	-110	-90	+110	+90
Final count	100	100	110	90

There are 210 ballots which rank B over Y, but each of those ballots only contributes **52%** of its weight to B in the final stage of the tally, because that is the transfer fraction for A's surplus. For a single ballot in each group, we find the rating differential by subtracting Y's rating from B's rating, as follows:

150 ballots: A 100 > B 100 > C 100 > X 0 > Y 0 > Z 0

B - Y = 100 - 0 = **100**

60 ballots: A 100 > B 90 > C 80 > X 0 > Y 0 > Z 0

B - Y = 90 - 0 = **90**

So, we have 150 ballots at 52% weight, each expressing an overall rating differential of 100, and we have 60 ballots at 52% weight, each expressing an overall rating differential of 90. We can find the overall weighted magnitude by calculating $(150 \times 0.52 \times 100) + (60 \times 0.52 \times 90)$. The winning rating differential of {AXB} over {AXY} is **10685.7**.

Note that, if desired, the ratings information in a cardinal-weighted CPO-STV election can be maximized in scale to some degree. The ratings that are important are the ones for candidates who are in some of the Schwartz set outcomes but not all of them. You can have a rule that automatically maximizes the rating differentials between these candidates on each ballot, such that the highest-ranked becomes 100, the lowest-ranked becomes 0, and the different gaps in between retain their original ratios.

9. Conclusion

I do not intend cardinal-weighted pairwise as a frivolous academic exercise or a mathematical show-off. I intend it as a realistic proposal, one that I sincerely prefer over other existing proposals. I recognize that it adds an extra layer of complexity, but I feel that the compensation far outweighs the cost.

I strongly approve of Condorcet-efficient voting methods, but I feel that they can be significantly improved-upon in two ways. One, that the relative priority of voters' pairwise preferences should be taken into account, and two, that the method should be more resistant to the burying strategy. I find it serendipitous that the same principle can achieve both benefits simultaneously.

I find both of these potential improvements to be quite significant, but perhaps the strategic issue is the more pressing of the two. I am convinced that voting methods that aim for majority rule should be Condorcet-efficient, but I suspect that the burying strategy could prove to be a serious problem for Condorcet-efficient methods in contentious elections. It is important to have a method that, in addition to recognizing a Condorcet winner when one is clearly expressed, works to protect sincere Condorcet winners from being obscured by strategic incursion. I believe that cardinal-weighted pairwise accomplishes this to an unusual degree.

So, I ask those who read this paper to seriously consider this method, and to afford it further discussion. Perhaps it will be some time before its implementation will be practical on a large scale, but if we allow ourselves to imagine, and to ask what kind of voting methods will be used in a far more democratic and collectively-intelligent future, I suggest that they might resemble cardinal-weighted pairwise.

Definitions

Approval voting: Voters can give each candidate a rating of either one or zero. The candidate with the highest rating sum wins.

Beatpath method: A beatpath is a series of pairwise defeats that form a path from one candidate to another. For example, if A beats B, and B beats C, then there is a beatpath from A to C. The strength of a beatpath is defined as the strength of its weakest defeat. If the strongest beatpath from X to Y is stronger than the strongest beatpath from Y to X, then X has a beatpath win over Y. The winner of the beatpath method will be a candidate such that no other candidate has a beatpath win against it. See Schulze (2003).

Cardinal ratings: For example, a ballot which asks voters to rate candidates on a scale from 0 to 100.

Comparison of pairs of outcomes by the single transferable vote (CPO-STV): A proportional representation method that uses the STV principle to make pairwise comparisons between multiple-winner outcomes. See Tideman (1995). See also Tideman and Richardson (2000).

Condorcet criterion: If a Condorcet winner exists, it should be elected with certainty.

Condorcet winner: A candidate that wins all of its pairwise comparisons. Also called a ‘dominant candidate.’

Instant runoff voting: Voters submit ranked ballots. The candidate with the fewest top-choice votes is eliminated, and the ballots are recounted. This process repeats until one candidate remains. Also known as 'the alternative vote.'

Majority criterion: If a majority of the voters ranks one candidate in first place, over all other candidates, then this candidate should win.

Majority rule cycle: A circular series of pairwise comparisons (e.g. A beats B, B beats C, C beats A) that leaves no single candidate unbeaten.

Margins: Shorthand term for any pairwise count method in which the strength of a defeat is defined by the number of ballots on the winning side of the defeat, minus the number of ballots on the losing side of the defeat.

Minimax: The winner is the candidate whose worst pairwise loss (if any) is the least severe. Equivalent to a method where you drop the weakest pairwise defeat until a single candidate is undefeated.

Mutual majority criterion: If there is a single majority of the voters who rank every candidate in a set S1 over every candidate outside S1, then the winner should certainly be a member of S1.

Newland-Britton quota: The number of votes cast divided by one more than the number of positions to be filled. See Newland and Britton (1976).

Pairwise comparison: Given two candidates A and B, A wins its pairwise comparison against B if and only if A is ranked above B on more ballots than B is ranked above A. If the number of A>B ballots is equal to the number of B>A ballots, then there is a pairwise tie between A and B.

Ranked pairs method: Defeats are considered in descending order of strength. They are locked in place unless they make a cycle with already-locked defeats, in which case they are skipped. The winner will be a candidate who is undefeated after all the defeats have been considered. See Tideman (1987).

River method: Similar to ranked pairs, but it does not lock more than one defeat against the same candidate. See Heitzig (2004).

Schwartz set: In this paper, I refer to Schwartz's GOCHA set as the Schwartz set. An undominated set is a non-empty set of candidates such that no candidate within the set is pairwise beaten by a candidate outside the set. A minimal undominated set is an undominated set that doesn't contain a smaller undominated set. The union of all minimal undominated sets is called the GOCHA set (for Generalized Optimal-CHoice Axiom). See Schwartz (1986).

Sincere Condorcet winner: A candidate who would be a Condorcet winner if all ballots were cast sincerely.

Strong Condorcet winner: A candidate whose pairwise victories are each supported by more than one half of the ballots.

Tie-breaking ranking of the options (TBRO): Choose a random ballot and engrave all of the pairwise preferences expressed in that ballot onto the TBRO. Pairs that are not yet evaluated can be evaluated by choosing more random ballots, one by one, until strict preference relations have been established between all of the candidates, or until all ballots have been used. If there is still an unevaluated pair (which requires that absolutely none of the voters express a preference between the pair), the order can be determined randomly. See Zavist and Tideman (1989).

Winning votes: Shorthand term for any pairwise count method in which the strength of a defeat is defined by the number of ballots on the winning side of the defeat.

Notes

1. We assume that the candidates are located on a one-dimensional spectrum such that only candidates ranked in first place on a given ballot will be ranked above all adjacent candidates on the spectrum. By the definition of a median, a candidate who is at the median will be preferred by a majority over a candidate who is away from the median; hence such a candidate cannot lose a pairwise comparison.
2. It should be clear that, using these methods and most other single-winner methods, if there is a majority of the voters who fully coordinate their efforts, they can elect whomever they choose. If a sincere strong Condorcet winner existed but was not chosen by a given method, and there was an opportunity for people to coordinate their intentions and recast their votes, then in theory, such a majority would have both the incentive and ability to elect the Condorcet winner instead, because that majority would clearly prefer the Condorcet winner over the current winner.
3. The following is an example where this simple cardinal ratings method simultaneously fails the majority, mutual majority, and Condorcet criteria: 100 voters and 3 candidates: A, B, and C. 55 people vote A 100, B 30, C 0; 45 people vote B 100, A 0, C 0. A receives 5500 points, B receives 6150 points, and B wins.
4. This is a strategic chicken game between supporters of similar candidates, caused by the fact that if I prefer the candidate in second place to the candidate in first place, I will have an incentive to insincerely lower my rating of the first place candidate, but if others follow the same incentive, the situation tends toward a car-crash disequilibrium. See “approval voting” in Green-Armytage (2003).
5. This follows from the fact that Condorcet-efficiency is incompatible with the later-no-help criterion, which is shown in Woodall (1997).
6. There are many possible ballot designs that can fit these specifications, but ballot design will not be discussed in this paper.

7. This provision was suggested to me by Chris Benham.
8. For example, a candidate ranked in first place could be given a default rating of 100, a candidate ranked in last place could be given a rating of 0, and remaining default ratings could be spaced evenly within the constraints imposed by surrounding ratings.
9. A full discussion of large cycle (cycles of more than three candidates) strategy is beyond the scope of this introductory paper. In short, I suggest that manufactured large cycles are not a severe problem, but that the possibility is worth keeping an eye on. In some cases, of course, four-candidate cycles can be legitimate, and hence not a strategic problem. On the other hand, if supporters of a candidate B try to use another candidate X (who is far from being a member of a sincere majority rule cycle) to help form a four candidate cycle to overrule an A>B defeat, there are several preconditions that have to be met for it to succeed, execution can be prohibitively complex, and there is a relatively easy and stable counterstrategy whereby some or all of the A>B voters give a blank rating to candidate X.
10. It will be helpful to support this with real voting data, but in the meantime, it seems intuitive in theory.
11. This method of assigning a rating to a multiple-candidate outcome was suggested to me by Nicolaus Tideman.

References

Cretney, B. (2004). Electoral methods: definitions. http://condorcet.org/emr/defn.shtml

Gibbard, A. (1973). Manipulation of voting schemes: A General Result. Econometrica 41: 587-601.

Green-Armytage, J. (2003). A survey of basic voting methods. http://fc.antioch.edu/~jarmyta@antioch-college.edu/voting_methods/survey.htm

Green-Armytage, J. (2004). Cardinal pairwise criteria. http://fc.antioch.edu/~jarmyta@antioch-college.edu/voting_methods/cwp_criteria.htm

Heitzig, J. (2004). River method -- updated summary. Discussion list entry. October 6. http://lists.electorama.com/pipermail/election-methods-electorama.com/2004-October/013971.html

Hylland, A. (1980). Strategy proofness of voting procedures with lotteries as outcomes and infinite sets of strategies. University of Oslo.

Moulin, H. (1988) Condorcet's principle implies the no show paradox. Journal of Economic Theory 45: 53-64.

Monroe, B. (2001) Raising Turkeys: An extension and devastating application of Myerson-Weber voting equilibrium. Presentation to the American Political Science Association in San Francisco. September.

Newland, R.A. and Britton, F.S. (1976). How to conduct an election by the single transferable vote, second edition. Electoral Reform Society of Great Britain and Ireland.

Richie, R. and Bouricus, T. (2004). Draft of CVD analysis about IRV vs. Condorcet voting. http://groups.yahoo.com/group/instantrunoff/message/1548

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

Schulze, M. (2003) A new monotonic and clone-independent single-winner election method. Voting Matters 17: page 9-19

Schulze, M. (2004a). Compromising strategy in Condorcet. Discussion list entry. September 19.

http://lists.electorama.com/pipermail/election-methods-electorama.com/2004-September/013908.html

Schulze, M. (2004b). Compromising strategy in Condorcet. Discussion list entry. September 19. http://lists.electorama.com/pipermail/election-methods-electorama.com/2004-September/013906.html

Schwartz, T. (1986). The Logic of Collective Choice. New York: Columbia University Press.

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

Tideman, T.N. (1995). The single transferable vote. Journal of economic perspectives 9(1): 27-38.

Tideman, T.N. and Richardson, D. (2000). Better voting methods through technology: The refinement-manageability trade-off in the single transferable vote. Public Choice 103: 13-34.

Woodall, D.R. (1997). Monotonicity of single-seat preferential election rules. Discrete Applied Mathematics 77: 81-98.

Young, H.P. (1975). Social choice scoring functions. SIAM Journal of Applied Mathematics 28: 824-838.

Zavist, T.M. and Tideman, T.N. (1989). Complete independence of clones in the ranked pairs rule. Social Choice and Welfare 6: 167-73.