Moderated Differential Pairwise Tallying: A Voter Specified Hybrid of Ranking by Pairwise Comparisons and Cardinal Utility Sums

Joseph W. Durham and Peter Lindener
joey.durham@gmail.com
lindener.peter@gmail.com

1 Introduction

In our increasingly interactive world, elections and other forms of collaborative group decision-making are becoming ever more important. Many voting systems have been proposed, from common plurality-based methods to the historical approaches of Condorcet and Borda and, more recently, Single Transferable Vote, Range Voting, and Alternative or Instant Runoff Voting ([10], [12], [14], [19]). There is increasing awareness of how the mathematical properties of voting systems affect not just the election outcome but also which options are really considered and the content of pre-election debate. An understanding of the tradeoffs of various approaches appears critical for democratically governed groups of all sizes to realize their full potential.

Democratic decisions made when choosing between only two candidates or options seem straightforward: the option with the most votes for should win. Unfortunately, in voting situations that have more than two alternatives, democratic decisions rapidly become more problematic. This paper presents a new vote-tallying method, *moderated differential pairwise tallying*, that can improve the quality of single-winner elections. This method is a per-voter hybrid of Condorcet's pairwise comparison with a cardinal-weighted revision of the Borda count which gives all voters control over exactly how their preferences are tallied. We also show that the method maintains the virtues which make pairwise comparisons so appealing while reducing the potential for ambiguous cyclical outcomes. Our hope is that the adoption of this and future advances in social choice theory will help build more responsive democracies and encourage greater civic participation.

For this publication, see www.votingmatters.org.uk

2 The Spoiler Effect and Other Challenges

Before describing the developments in this paper, we would like to provide some motivation and introduce some key ideas. Those readers familiar with the spoiler effect and Condorcet's method of pairwise comparison may choose to go on to the next section.

Consider the following common scenario: two candidates, P and Q, each receive the support of about half of the voters. Selecting between just P and Q is simple: whichever candidate receives more support in a head-to-head comparison can be deemed preferred by the group as a whole. Now consider the effect of adding a third option, M, to the set of candidates. In commonly used *plurality* voting systems, each voter is given one vote to cast for either P, Q, or M. With the addition of alternative M each voter must decide whether or not he would like to change his vote from P or Q to M. If a voter still prefers P or Q to M, then logically he should continue to support his prior top candidate. But the voting system has introduced a risk for those voters who might change their vote from P or Q to M, since removing support from either P or Q might cause the other less desired candidate to win.

The voters who are considering a switch to M are in a quandary. Each voter is forced to weigh the potential benefit of switching his vote to M against the risk of causing his least preferred of the three candidates to win. It is even possible that every voter would actually prefer M to either P or Q but, because of the perceived risk associated with voting for a third alternative, the group will remain polarized and stuck choosing between P and Q.

This M, P, and Q scenario is an example of the third-party *spoiler effect*, where the addition or removal of a non-winning candidate can affect which candidate wins. Almost all voting methods in use today are plagued by variations of the spoiler effect, but it is particularly problematic in plurality-based methods where each voter votes for a single candidate. The prevalence of the spoiler effect in a voting method can directly influence how an election unfolds in several ways.

First, in the face of the spoiler effect, voters are often forced to speculate about which candidates are the top contenders to determine how to avoid a detrimental result. The spoiler effect effectively penalizes voters who do not use a voting strategy which looks to vote only for one of the front-running candidates. This unfortunate reality can cause the outcome of an election to depend more on perceptions of popular opinion than the electorate's true preferences.

In addition, once a voter has decided to support one of the perceived front-runners, there is little incentive for him to consider other alternatives. The perceived risk of supporting a third-party candidate limits which options are fully explored. By focusing attention on the perceived front-runners, the spoiler effect contributes to the polarizing divisiveness surrounding some elections and can influence not only which candidate is elected but the very nature of the preceding democratic debate itself.

These distinctly less-than-democratic outcomes are just a few examples of the potential consequences of using a voting system which is subject to the spoiler effect. The frequent worry about "throwing one's vote away" around election time illustrates the awareness of many voters about the spoiler effect, even if they do not use the term. These issues are so common that they are often accepted as being an inherent part of the election process. However, the potential for electoral results to be influenced by the spoiler effect highlights the very real value of a well-formed, truly democratic voting system. The question then becomes: how can a democratic selection be made between several candidates without inviting the spoiler effect?

As we have mentioned, the spoiler effect is particularly troublesome in plurality methods where each voter has a single vote to cast. In an attempt to reduce the severity of the spoiler effect, many political systems using plurality methods hold a series of smaller contests including primaries, runoffs, or both. Although holding a final runoff between the top two candidates does return to the simple two candidate scenario for the last stage of an election, the voting which determines who will be in the runoff is still subject to the spoiler effect. Recently, Single Transferable Vote (STV) and Alternative Vote or Instant Runoff Voting (IRV) have received a lot of attention for bypassing the logistical need for a separate runoff. However, the implementations of these methods still involve a series of plurality votes and so the spoiler effect still influences election results.

To show how a voting system can be designed to avoid the spoiler effect, we will return to the P, Q, and M scenario. When P and Q were the only two choices, selecting a winner was easy and there was no threat of the spoiler effect. The simplicity of the two candidate scenario holds a clue about how the spoiler effect can be eliminated from democratic decision-making. When candidate M is included in the pool of options and voters are asked to pick one of three candidates, then the spoiler effect appears. Consider instead if we asked voters to pick one candidate out of each possible pair of candidates: {P,Q}, {P,M}, or {Q,M}. With this approach, a voter can still support P over Q, but also express his preference for M over either P or Q, for example. By evaluating the candidates in pairs in this way, the risk of the spoiler effect caused by support for newcomer M is removed. This approach to tallying votes is known as *exhaustive pairwise comparison*.

To give voters the freedom to express their preferences over many pairs of candidates, a different kind of ballot is required, something known as a *ranked choice ballot*. Ranked choice ballots, which are also used in STV, IRV, and other methods, allow voters to rank the candidates against each other. Ranked choice ballots allow the concept of pairwise contests to be easily extended to all possible combinations of candidates. For example, if a voter ranked M higher than P on a ranked ballot, then it would be the same as the voter supporting M in the pairwise comparison {P,M}.

The combination of ranked choice ballots and exhaustive pairwise comparison forms the foundation of the voting method first proposed by Condorcet in 1785 [5]. Condorcet's approach effectively holds a simultaneous runoff between every possible pair of candidates. To resolve the P, Q, and M scenario with Condorcet's method, each voter would submit a ballot ranking the three choices. If candidate M won her pairwise comparison against both P and Q, then M would be the *Condorcet winner* and would be selected as the group's most preferred choice.

Pairwise comparisons have the important advantage of *strict candidate pair dependence*: a voter's ranking for candidate M has no effect on the relative standing of P and Q. When exhaustive pairwise comparison produces a Condorcet winner, then the result is *independent from irrelevant alternatives*, meaning that any non-winning candidate can be removed without affecting the outcome. Thus, we have found a method for making a truly democratic decision between multiple candidates without risking the emergence of the spoiler effect.

It is worth noting that plurality-based runoff methods, including STV and IRV, will often eliminate a potential Condorcet winner. The spoiler effect in the early rounds of these plurality-based methods will often lead to a runoff between two less-preferred candidates. A Condorcet winner will always win a runoff election against any of the other candidates and is therefore distinctly the choice of the voters overall.

While exhaustive pairwise comparison has many important virtues when there is a Condorcet winner, unfortunately such a decisive winner does not always occur. When no single candidate wins all of her pairwise sub-contests, a set of more-preferable candidates will often emerge above the rest of the choices. This *top cycle set* wins over all candidates outside the set, but there is no coherent ordering of the candidates within the set [16]. The figure below gives one example of a cycle set, where P wins over M, M wins over Q, but Q wins over P to create a cycle. There exists a class of voting methods known as Condorcet methods which differ only in how they select a winner in the case of a top cycle set. There is little data, however, on how frequently top cycles would occur in real elections, particularly when there are a large number of candidates. Nonetheless, the potential for these ambiguous results with Condorcet's approach means that some further analysis is needed.

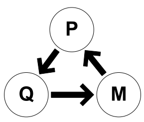

In this paper, we pursue an understanding as to why cycles can occur in pairwise tally results and develop a method to reduce their likelihood.

3 Contents of this Paper

The core contribution of this paper is a new pairwise tallying formulation which unifies Condorcet's method with a linear version of the *Borda count*. As we will show, classic pairwise comparison discards critical relative priority information from voters which can resolve cycles. Our new hybrid method, *moderated differential pairwise tallying* (MDPT), is built on the premise that not all voters will choose to strategically maximize their ballot's influence to the Condorcet-style extreme. A new voting parameter called the *moderation span* will be introduced, which allows voters to express slight preferences between candidates they find similarly preferable. We show that when voters use the new moderation parameter, this new tally formulation decreases the chances that a cycle will occur without introducing any dependence on irrelevant alternatives. This concept of individual moderation also suggests intriguing new approaches for resolving top cycles.

This paper is organized as follows. In Section 4 we explore the mathematics of Condorcet's pairwise analysis. The issue of ambiguous cyclical results and their inevitability in the face of Arrow's impossibility theorem is the topic of Sections 5 and 6. We present a fully linear Cardinal utility system in Section 7, but the necessary introduction of a constraint on voter weight in Section 8 leads to a real-valued formulation similar to Borda's method. In Sections 9, 10, and 11 we transform the Condorcet and real-valued Borda tallying methods into a common difference matrix framework. From these transformations we derive our new voter-specified hybrid method, MDPT, in Section 12. We show in Sections 13 and 14 how this new tallying formulation reduces the potential for cyclical results. In Section 15 we discuss some properties and interpretations of moderated tallying. Section 16 focuses on the practical aspects of implementing MDPT, including a suggestion of how voters might cast preference ballots and set the new voting parameter. We conclude in Section 17 and offer some intriguing future directions for this work in Section 18.

For those already versed in social choice theory, we hope that the perspectives on Condorcet, Borda, and strategic voting we present will spark further insights. For those new to some of these concepts or interested in how our proposal might be used, you may want to read the more practical material in Section 16 before diving into the other sections. We have tried to ensure that your effort to grasp these ideas is rewarded with some new and interesting understanding. We would also like to open up these ideas for discussion and look forward to dialog with others interested in improving commonly used democratic group decision methods.

4 Condorcet pairwise tallying

As we described in Section 2, pairwise tallying holds a sub-contest for every pairwise combination of candidates. By considering just the candidate pair in question for each sub-contest, this formulation remains free of the third-party spoiler effect. Based on recently rediscovered manuscript transcriptions, the concept of making group decisions using exhaustive pairwise comparison has a long history. In the 13th century Llull described an iterative procedure for small groups to elect a leader by holding a head-to-head vote for every possible pairwise combination of candidates [8]*. Condorcet proposed the first known pairwise tallying method based on ranked choice ballots in 1785 [5], [11].

We will first present the equations for pairwise tallying and then show an example of the computation. The first step in determining the winner of candidate pair {A,B} with pairwise tallying is to tally how many voters ranked A higher than B,

$$T_{cond}\,[\mathrm{A,B}] = \sum_{v}^{\text{all voters}} \left(\vec{b}_v[A] > \vec{b}_v[B]\right) \qquad (1.1)$$

For this paper, $\vec{b}_v$ will always be a voter's real-valued preference ballot vector and the term $\vec{b}_v[A]$ is the preference rating given to candidate A by voter *v*. The relative standing of $\vec{b}_v[A]$ and $\vec{b}_v[B]$ indicates the voter's relative preference for A or B with the expression $(\vec{b}_v[A] > \vec{b}_v[B])$ yielding 1 when $\vec{b}_v[A]$ is greater than $\vec{b}_v[B]$ and zero otherwise. Iterating (1.1) over all pairs of candidates computes a square *pairwise tally matrix*, T_{cond}, using Condorcet's approach. $T_{cond}\,[\mathrm{A,B}]$ contains the number of voters who ranked candidate A higher than B. While not necessary for (1.1), in this paper all preference ratings will be real values. The use of real values will prove significant in later sections.

4.1 Example: Condorcet Tally of One Ballot

For this example we will compute how the single real-valued preference ballot below is added to an aggregate tally.

Voter 1	
1.0	A
0.9	C
0.0	B

* As Llull was well-read in the writings of Arabic scholars, there is some speculation that Llull's ideas may have their roots in prior Arabic thinking, see [12]

This voter has ranked A > C > B. For all the examples in this paper ballots will span the interval [0,1] for simplicity, but any interval is acceptable. When this ballot is tallied using (1.1), the voter's support is added to $T_{cond}\,[\mathrm{A,B}]$, $T_{cond}\,[\mathrm{A,C}]$, and $T_{cond}\,[\mathrm{C,B}]$, as shown in matrix form below.

Voter 1

	A	B	C
A	0	1	1
B	0	0	0
C	0	1	0

Notice that while the ordering of candidates on the voter's ballot can be determined from this matrix, the original spacing between the candidates cannot.

4.2 Picking a Winner from a Condorcet Tally

To determine a choice from a pairwise tally, we first compute a *delta-tally matrix* where each element compares how candidate A fared relative to candidate B. Subtracting the transposed tally matrix T^T_{cond} from T_{cond} yields the difference between row A, column B and row B, column A,

$$D_{cond} = T_{cond} - T^T_{cond} \qquad (1.2)$$

In other words, (1.2) tallies the number of voters who ranked candidate A higher than candidate B versus B higher than A across all pairs of candidates in matrix form. The resulting delta-tally matrix, D_{cond}, is anti-symmetric and has zeros on the diagonal (where candidates tie with themselves). From D_{cond} we can easily compute a pairwise win-Boolean matrix W_{cond},

$$W_{cond} = (D_{cond} > 0) \qquad (1.3)$$

As before, the $(x > y)$ operator yields 1 (true) when x is greater than y, 0 (false) otherwise. When this operator is applied to a matrix of values it operates on an element-by-element basis. $W_{cond}\,[\mathrm{A,B}]$ will be 1 if $D_{cond}\,[\mathrm{A,B}] > 0$ (ie, when A beats B) and 0 otherwise. Overall contest results are then determined by examining the full contents of this pairwise win matrix.

Together (1.1), (1.2), and (1.3) perform Condorcet's pairwise analysis. If all the win-Booleans in a candidate's row in W_{cond} are 1 (except the 0 diagonal where candidates tie with themselves), this candidate has won a direct comparison with every other candidate in the election and is distinctly the best choice. As mentioned in Section 2, a candidate who wins all of her pairwise comparisons is termed a *Condorcet winner*. Since the delta-tally matrix is anti-symmetric, a Condorcet winner will also have her corresponding column in W_{cond} all 0 since she will not have lost to anyone.

4.3 Example: Condorcet Winner

For this example we will examine a three voter, three candidate election which produces a Condorcet winner.

Voter 1		Voter 2		Voter 3	
1.0	A	1.0	B	1.0	C
0.9	C	0.2	C	0.8	A
0.0	B	0.0	A	0.0	B

As in Example in section 4.1, we can compute what each voter will contribute to the tally matrix.

Voter 1

	A	B	C
A	0	1	1
B	0	0	0
C	0	1	0

Voter 2

	A	B	C
A	0	0	0
B	1	0	1
C	1	0	0

Voter 3

	A	B	C
A	0	1	0
B	0	0	0
C	1	1	0

Summing all of these contributions produces the following pairwise tally matrix.

T_{cond}

	A	B	C
A	0	2	1
B	1	0	1
C	2	2	0

To compute the corresponding delta-tally matrix, we subtract T_{cond}^T from T_{cond}, so $D_{cond}\,[\mathrm{A},\mathrm{B}] = (T_{cond}[\mathrm{A},\mathrm{B}] - T_{cond}\,[\mathrm{B},\mathrm{A}])$. From the delta-tally we then compute the pairwise win-Boolean matrix.

D_{cond}

	A	B	C
A	0	1	-1
B	-1	0	-1
C	1	1	0

W_{cond}

	A	B	C
A	0	1	0
B	0	0	0
C	1	1	0

The winner of this election is candidate C who wins over both A and B. This can also be seen directly from D_{cond} since C's row is all positive except the 0 diagonal.

4.4 Properties of Condorcet Tallying

Since each matrix element $T_{cond}\,[\mathrm{A},\mathrm{B}]$ depends only on relative ballot positioning of the associated pair of candidates, each element of T_{cond} possesses *strict candidate-pair dependence*. This desirable property is a direct result of Condorcet's method of tallying votes. Since each voter's full support is given to his preferred candidate in every pairwise sub-contest, all pairwise results are independent of the presence of any candidate not in the pair. Therefore, when there is a Condorcet winner, pairwise tallying exhibits the very desirable property of *independence from irrelevant alternatives*: any non-winning candidate can be added or removed from the contest without changing the result. As a consequence of this independence, pairwise tallying is free of the spoiler effect. The catch, however, is that a Condorcet winner does not exist for all possible collections of ballots.

5 Coinciding Cyclical Majorities

As Condorcet discovered in 1785, even though each voter submits a strictly ordered list of candidates, the set of pairwise contest results can form an ambiguous cycle [3, pg 193]. One collection of ballots which produces a cyclic outcome is shown in Example below.

5.1 Example: Coinciding Cyclical Majorities

In this example we will modify one of the ballots from Example in 4.2 such that the pairwise tally results in a cycle. We will switch Voter 1's rating of candidates B and C so that the previous winner, C, is now at the bottom of his ballot.

Voter 1		Voter 2		Voter 3	
1.0	A	1.0	B	1.0	C
0.9	B	0.2	C	0.8	A
0.0	C	0.0	A	0.0	B

Notice that each voter is using a different rotation of the same preference order. Tallying these ballots produces the following win-Boolean matrix.

W_{cond}

	A	B	C
A	0	1	0
B	0	0	1
C	1	0	0

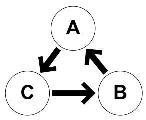

Each candidate wins one of their pairwise comparisons and loses the other one, forming a cycle: A wins over B, B wins over C, but C also wins over A. A cycle can also be shown graphically using a win edge-graph as above, where directed edges point from the loser to the winner of each pairwise sub-contest.

5.2 Cause of Cycles

There are several names for the cycle phenomenon, including "majority rule cycle" and "cyclical majority." In his writings on the topic, Condorcet often used the word "contradictoire" when referring to cycles [5]. The potential for cyclical results is also often referred to as Condorcet's "paradox": although each voter submits a ranked ballot, it is possible for the group tally to have no coherent ranking. We prefer the term *coinciding cyclical majorities* which emphasizes that the pairwise majority-rule victories which produce the cyclical result all occur from the same collection of ballots. Regardless of the name, the potential for ambiguity in the ranking produced by pairwise tallying can seem paradoxical, contradictory, or at least disconcerting. Under some situations a result with no decisive winner may be acceptable, but in general we require a method which can resolve any possible set of input ballots. This paper shows why coinciding cyclical majorities are actually an expected result of the mathematics of pairwise tallying and what can be done to reduce their likelihood.

Contrary to potential misconception, the potential for coinciding cyclical majorities is not so much due to underlying irrational preferences of the voters as it is a product of how votes are tallied. As noted earlier, Condorcet's method of maximizing a voter's influence over each pairwise sub-contest independently means that pairwise tallying is immune to the spoiler effect. However, since each pairwise sub-contest is tallied based only on the relative rank order of the two candidates involved, information regarding relative preference magnitudes is entirely lost. The sign of the preference difference between the candidates is conserved, but the strictly pairwise perspective in Condorcet's evaluation cannot distinguish between a voter's significant, modest, or trivial preference differentials.

As noted by Saari [14], the emergence of cycles can be seen as a result of this information loss: coinciding cyclical majorities occur because of the distortion of all voter priorities to the same weight by Condorcet's style of pairwise voter influence maximization. On occasion, the distortion caused by the non-linearity of Condorcet's binary pairwise comparison can overwhelm a weaker consensus and this incoherence of pairwise victories may manifest as a cycle. As we will show in Section 10, relative voter priority information loss and the resulting potential for cycles become even more problematic as the number of candidates increases.

Condorcet discovered cycles in the late 18th century, but it would be another 150 years before a now famous theorem would more clearly explain the obstacles to designing an optimal social choice method.

6 Arrow's Impossibility Theorem

In his 1951 book *Social Choice and Individual Values* [1], economist Dr. Kenneth Arrow proposed a list of properties that an ideal social choice method or voting system would possess:

1. unrestricted domain (or universality);

2. positive association of values (monotonicity);

3. independence of irrelevant alternatives (binary independence);

4. non-imposition (or citizen sovereignty);

5. non-dictatorship.* †

Arrow's famous *impossibility theorem* concludes that a democratic voting system with three or more options cannot achieve all of these desirable properties for all possible collections of ballots. In his 1972 Nobel Prize lecture, Arrow concluded: "*Certainly, there is no simple way out. I hope that others will take this paradox as a challenge rather than as a discouraging barrier*" [2].

While Arrow's theorem may seem discouraging at first, Condorcet's pairwise analysis provides some hope. When there is a Condorcet winner, pairwise analysis achieves all of the desirable properties Arrow listed. For many collections of ballots a winner can thus be found without relaxing any of these properties. The challenge then becomes two-fold: reducing the occurrence of cycles while maintaining Arrow's properties and resolving cycles when they occur with minimal relaxation. There have been several proposals for how to resolve cycles since Condorcet discovered them, including ones by Condorcet [19], Tideman [17], Schulze [15], and Green-Armytage [7]. We will instead directly address reducing the prevalence of cyclical majorities by pursuing a deeper understanding of information loss in pairwise analysis which causes cycles.

* There exist many variations of this famous impossibility theorem, including a 1963 version which replaces the monotonicity and non-imposition criteria with the Pareto efficiency. Monotonicity, however, is in its own right a frequently discussed property of social choice functions so we have chosen to use the original version of the theorem.

† Expanded description of Arrow's five properties:

a) Unrestricted domain means that (1) each voter must have the freedom to rank all of the choices available in any order, (2) the voting mechanism must be able to process all possible sets of voter preferences, and (3) it must consistently give the same result for the same profile of votes — no randomness is allowed.

b) Monotonicity, also termed "positive association of social and individual values", means that a change in a candidate's placement on a ballot (either higher or lower), if it causes a change in the candidate's ranking, can only result in a change in final ranking in the same direction.

c) Indepedence of/from irrelevant alternatives means that if A is preferred to B out of the choice set {A,B} by the electorate as a whole, then introducing a third alternative X, thus expanding the choice set to {A,B,X}, must not make B preferred to A. This property is also referred to in the literature as *binary independence*.

d) Non-imposition means citizens must be free to vote for the candidate(s) of their choice.

e) Non-dictatorship means no single voter determines the entire contest outcome.

7 An idealized, real-valued linear system perspective

Since the potential for cycles results from the distortion of differential preference magnitudes, an idealized choice function that will avoid this pitfall seems straightforward. We can instead simply sum every voter's real-valued preference for each candidate,

$$\vec{\tau}_{benefit-cost} = \sum_{v}^{\text{all voters}} \vec{b}_v \qquad (1.4)$$

With this approach, all relative preference information is aggregated into the tally vector $\vec{\tau}_{benefit-cost}$ and the winner is then the candidate with the highest component. Since it is a completely linear system, (1.4) does not distort preferences at all.

In this idealized framework, each voter's preference ballot can be interpreted as a vector of *expected cardinal utilities* or *von Neumann-Morgenstern utilities* [20]. The preference value assigned to each possible outcome would be the voter's expected benefit from that outcome minus any associated cost. Under such a system it would be considered the social responsibility of every voter to understand how to map his personal utility into the group's summation. The optimal solution for the group as a whole is then the alternative with the highest tally in the benefit-cost sum. This choice function is also known as the Bentham-Edgeworth sum of individual utilities. It can also be thought of as a Range Voting variant where each voter can pick his own range.

If all voters submit appropriately scaled ballots, then (1.4) represents a social choice function that effectively achieves all of Arrow's five stated properties. Unconstrained ballots mapped linearly by a sum into a global tally vector represent a truly unrestricted domain. Full linearity implies positive association of values, as all changes to a voter's preference ballot are positively conducted directly into the tally. Irrelevant alternatives never affect the ranking of other candidates because ballots are unconstrained and each candidate can be given any value completely independent from the others. With this function, voters are also free to vote as they wish for each candidate and no voter determines the entire election results (unless all other voters agree that they should).

However, even though each voter should constrain the envelope of his preference schedule to only exert his appropriate level of influence, it is in the voter's interest to maximize the influence of his preference schedule. This reality causes this function to be unusable. Contentious decisions and differing concepts of individual social responsibility could easily cause an endless escalation of ballot vector magnitudes as competing factions wrestle for control of a decision outcome.

8 Real-world constraints

Rescaling ballots to standardize the minimum and maximum preference values that a voter can express appears to be an easy solution to the issue of voter self-interest which renders (1.4) effectively unusable. Subtracting the minimum preference value and then dividing by the ballot span $\Delta_v = \max(\vec{b}_v) - \min(\vec{b}_v)$ normalizes each voter's ballot to a standard interval between 0 and 1,

$$\tau_{\text{rvborda}} = \sum_{v}^{\text{all voters}} \frac{\vec{b}_v - \min(\vec{b}_v)}{\Delta_v} \qquad (1.5)$$

Similar to (1.4), (1.5) retains all information regarding the relative priorities of a voter but also limits the maximum magnitude of support a voter can express for a candidate. We will refer to this approach as a real-valued Borda method but note that it is also equivalent to Range Voting when all voters mention every candidate and rescale their ballot.*

8.1 Example: Real-valued Borda

To show how real-valued Borda differs from pairwise tallying, we will tally the same set of ballots as in Example in 5.1. With Condorcet-style tallying these ballots produce a cyclic result because of the loss of relative preference magnitude information.

Voter 1		Voter 2		Voter 3	
1.0	A	1.0	B	1.0	C
0.9	B	0.2	C	0.8	A
0.0	C	0.0	A	0.0	B

Since all voters' ballot spans are 1.0 and all ballots start at 0.0, computing the real-valued Borda tally vector for this example is a simple sum of the candidate components. Note, however, that the results would be equivalent if any ballot was scaled to span [0, 100], shifted to span [13], [14], or some combination of the two.

$\vec{\tau}_{\text{rvborda}}$	
1.9	B
1.8	A
1.2	C

* More information about Range Voting, visit the Center for Range Voting at `http://www.rangevoting.org/`

For these ballots B gets the highest total and is thus the winner.

8.2 Properties of Real-valued Borda

A real-valued Borda tally is distinguished from the classic ordinal Borda tally, where preference metrics are constrained to an evenly spaced, strictly ordered ranking. By forcing preferences onto a constrained ordinal number line, classic ordinal Borda has no means to express either that one candidate stands apart from the others or that two candidates are equally preferable. For example, with ordinal tallying the ballots in Example in 8.1 would produce a tie since all middle candidates would be effectively forced to 0.5. We use unconstrained cardinal-weighted ballots because we believe ballots should be instruments for representing all of a voter's relative preference information for any viable alternatives (an extension of Arrow's unrestricted domain).

In addition, when candidates are added to or removed from the middle of an ordinally constrained ballot, the voter's expressed preference value will change for all other candidates except the top and bottom choices. In contrast, inserting or removing a candidate from the middle of a real-valued ballot causes no change in the preference expression for the other candidates. The use of ordinally constrained ballots causes unnecessary and detrimental dependence on all the other alternatives under consideration to be introduced into a Borda-style tally. The use of real-valued preference ratings will also prove crucial in the development of our moderated differential pairwise tallying method.

The introduction of the normalization by ballot span for the real-valued Borda method in (5) has addressed the primary issue with (4). However, this introduction of a limit on the strength of a voter's authority renders Arrow's five properties mutually unachievable. Although (5) achieves Arrow's four other properties, it sacrifices independence from irrelevant alternatives. This loss of independence from irrelevant alternatives opens the door for strategic voting: as Borda himself observed, choice functions like (5) which are sensitive to less relevant alternatives work "only for honest men" [3, p. 215].

Ballot span normalization limits the ability of an individual voter to set her own weight of influence but, as we will now illustrate, encourages the adoption of a voting strategy. When there are larger numbers of candidates under consideration, a voter's ballot may be significantly stretched by irrelevant alternatives. This ballot stretching decreases a voter's authority over the true contenders, increasing the po-

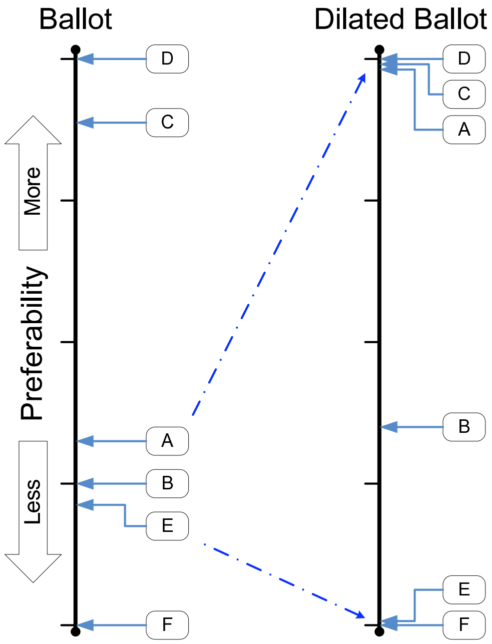

Figure 1.1: In this hypothetical example, a voter's true preferences are shown at left, while a strategically dilated version is shown at right. If the voter believes D, C and F have no chance in the election, it is in her best interest to strategically dilate and clip her ballot so that the contending candidates A, B and E define the span. While this distorts preference information for the candidates the voter believes are non-contenders, it maximizes the weight of influence between those considered to be of real importance.

tential gains from ballot manipulation. If a voter can predict the top contending candidates, she can increase her influence in the end decision by dilating and clipping her ballot to span just these top contenders; an example of this strategy is shown in Fig. 1.1. In the face of large-scale speculative gaming of this kind, the decision outcome becomes predominantly dependent on perceptions of popular opinion instead of true voter preference. We term this adverse situation of outcome dependence on perceived top contenders a form of *speculative indeterminacy*.

It is worth noting that while ordinal preference ballots preclude dilating and clipping, they encourage other strategic possibilities which are even worse. Instead of dilating their ballots, voters are encouraged to stuff the middle of their ballot with irrelevant candidates to increase their authority between the top contenders. This strategic reordering of candidates on the ballot further obscures the true wishes of the voters, yielding close to meaningless results. When there are large numbers of candidates, the potential gain from strategic voting with ordinal Borda increases as there are more irrelevant alternatives that can be stuffed in the middle of an ordinal ballot.

The only reasonable goal of any voting strategy is to elect the highest possible candidate from the voter's sincere preference schedule. As we have just described, with Borda's method voters can manipulate the placement of perceived non-contenders to increase the influence of their ballot over the front-runners. All vote tallying systems which do not exhibit strict candidate pair dependence will effectively encourage some kind of similar speculative voting strategy. When voters no longer express their true preferences to a social choice function, the election result cannot reflect the true desires of the electorate.

One way of interpreting Condorcet's tallying method is that it automatically maximizes each voter's influence since the voter's full influence is expressed between each candidate pair. Therefore, a strategic voter cannot do anything to change a pairwise sub-contest where her sincerely preferred candidate falls on the losing side. The property of strict candidate pair dependence limits a voter to trying much riskier indirect strategies. A voter can only attempt to flip a pairwise contest to a candidate she finds less preferable in the hope of creating a cycle which might end up resolving in her favor.

Unfortunately, as shown in Example in 5.1, when there are more than two candidates Condorcet's influence maximization also invites cyclical outcomes. In contrast, Borda-style methods with ballot span normalization always yield a distinct outcome but encourage several forms of strategic voting. In the following sections, we will transform these two classic tallying methods into a common delta-preference framework. This new framework will clarify their similarities and differences, and suggest a hybrid method which can exhibit the desirable properties of both methods.

9 Pairwise difference matrix of a vector

In preparation for the use of concise matrix notation throughout the rest this paper, we introduce an operator that computes the pairwise *difference matrix* of a vector. This operator will be used to express both

Condorcet's and Borda's methods into a unified matrix formulation. First, we build a square matrix M from a vector $\vec{v}$ by replicating the vector in each matrix column,

$$M = [\,\vec{v}\,\ldots\,\vec{v}] \tag{1.6}$$

The difference matrix of $\vec{v}$ is then defined as the column-replicated matrix M minus its transpose,

$$\mathrm{DiffM}\,(\vec{v}) = M - M^T \tag{1.7}$$

The subtraction of the transpose yields an anti-symmetric matrix which contains the pairwise difference between every combination of components in the incoming vector. The element [A,B] of the resulting matrix is equal to the difference in value between components [A] and [B] of the vector: $\mathrm{DiffM}\,(\vec{v})\,[\mathrm{A},\mathrm{B}] = \vec{v}\,[A] - \vec{v}\,[B] = -\mathrm{DiffM}\,(\vec{v})\,[\mathrm{B},\mathrm{A}]$ All the diagonal elements of $\mathrm{DiffM}\,(\vec{v})$ are 0. Note that this subtraction of the transpose is the same operation used in (1.2) to find the Condorcet delta-tally. We also note that the DiffM operator is linear, a property we will use to reorder operations in the development of the re-factored formulations that follow. When this operator is applied to a voter's preference ballot vector, we will refer to the resulting matrix as the voter's *delta-preference matrix* $\mathrm{DiffM}(\vec{b}_v)$.

9.1 Example: A Delta-Preference Matrix

For this example we will consider the same ballot as in Example in 4.1.

Voter 1	
1.0	A
0.9	C
0.0	B

As described, the [A,B] element of $\mathrm{DiffM}(\vec{b}_v)$ is given by $\vec{b}_v[A] - \vec{b}_v[B]$.

Voter 1

	A	B	C
A	0	1.0	0.1
B	-1.0	0	-0.9
C	-0.1	0.9	0

10 Condorcet pairwise tallying in difference matrix form

In preparation for forming a hybrid method that unifies the underlying approaches of Borda and Condorcet, we will reformulate Condorcet pairwise tallying employing the above difference matrix operator. The voter's delta-preference matrix $\mathrm{DiffM}(b_v)$ is used to compute the pairwise tally matrix across all pairs of candidates in a single element-by-element matrix operation,

$$T_{\mathrm{Cond}} = \sum_{v}^{\text{all voters}} \left(\mathrm{DiffM}\left(\vec{b}_v\right) > 0\right) \tag{1.8}$$

This equation is the equivalent matrix form of the classic pairwise comparison in (1.1).

In (1.2) we computed the Condorcet delta-tally D_{Cond} from pairwise tally T_{Cond}. We will now derive an equation for determining the Condorcet delta-tally directly from ballots by combining (1.2) and (1.8). We can reorder the differencing of T_{Cond} and its transpose from (1.2) to be inside the ballot tallying summation by again employing the delta-preference matrix $\mathrm{DiffM}(\vec{b}_v)$, which then gives D_{Cond} as:

$$\sum_{v}^{\text{all voters}} \left(\left(\mathrm{DiffM}(\vec{b}_v) > 0\right) - \left(\mathrm{DiffM}(\vec{b}_v) > 0\right)^T\right) \tag{1.9}$$

This delta-tally formulation is equivalent to performing both (1.1) and (1.2) but is tabulated directly from the voters' ballots without the need of the intermediate pairwise tally.

To condense (1.9) we will use the signum function which is defined as

$$\mathrm{sgn}(x) = \begin{cases} 1 & \text{if } x > 0 \\ 0 & \text{if } x = 0 \\ -1 & \text{if } x < 0 \end{cases} \tag{1.10}$$

Signum can also be written as the difference of two inequalities $\mathrm{sgn}(x) = (x > 0) - (x < 0)$ or equivalently $\mathrm{sgn}(x) = (x > 0) - (-x > 0)$. Since the delta-preference matrix $\mathrm{DiffM}\,(b_v)$ is always anti-symmetric, its transpose is equal to its negation: $(\mathrm{DiffM}\,(b_v) > 0)^T = (-\mathrm{DiffM}\,(b_v) > 0)$ $= (\mathrm{DiffM}\,(b_v) < 0)$. Using this, (1.9) simplifies to

$$D_{\mathrm{Cond}} = \sum_{v}^{\text{all voters}} \mathrm{sgn}\left(\mathrm{DiffM}(\vec{b}_v)\right) \tag{1.11}$$

This equation is the direct difference matrix expression of the Condorcet delta-tally in (1.2).

10.1 Example: Direct Delta-tally of Condorcet Winner

For this example we will use the set of ballots from Example in 4.2 which produced a Condorcet winner.

Voter 1		Voter 2		Voter 3	
1.0	A	1.0	B	1.0	C
0.9	C	0.2	C	0.8	A
0.0	B	0.0	A	0.0	B

Using (1.11), we can directly compute what each voter will add to the delta-tally.

Voter 1

	A	B	C
A	0	1	1
B	-1	0	-1
C	-1	1	0

Voter 2

	A	B	C
A	0	-1	-1
B	1	0	1
C	1	-1	0

Voter 3

	A	B	C
A	0	1	-1
B	-1	0	-1
C	1	1	0

As in Example in 4.1, the voter's original spacing between candidates cannot be recovered from his contribution to the Condorcet delta-tally. As expected, summing all voter contributions produces the same aggregate delta-tally as we found in Example in 4.1.

D_{Cond}

	A	B	C
A	0	1	-1
B	-1	0	-1
C	1	1	0

As before, C wins both of its pairwise comparisons and is therefore the winner.

It is also worth noting that since (1.11) is equivalent to (1.2), (1.11) will yield an ambiguous cyclical outcome for the same ballot collections as (1.2).

11 Real-valued Borda Tallying in Difference Matrix Form

We will now transform the vector Borda tallying method from Section 8 into an equivalent difference matrix form. This new form will have a similar structure to the Condorcet delta-tally presented in (1.11). To start, we can compute the delta-Borda tally matrix for a Borda tally vector produced by (1.5),

$$D_{\text{rvborda}} = \text{DiffM}(\vec{\tau}_{\text{rvborda}}) \tag{1.12}$$

The D_{rvborda} matrix contains the same information on the relative standing of candidates as the τ_{rvborda} vector, but it possesses a similar anti-symmetric structure to the Condorcet delta-tally in (1.2) and (1.11). We can commute the difference matrix operator in (1.12) to inside the summation across voters, creating an equivalent delta-Borda matrix formulation that implements the preference differencing operation on a per ballot basis similar to (1.11),

$$D_{\text{rvborda}} = \sum_{v}^{\text{all voters}} \frac{\text{DiffM}(\vec{b}_v)}{\Delta_v} \tag{1.13}$$

As in (1.5), Δ_v is the span of the voter's ballot, i.e. $\Delta_v = \max(\vec{b}_v) - \min(\vec{b}_v)$. $\text{DiffM}(\vec{b}_v)$ is the delta-preference matrix from the voter's ballot as described in Section 5. Normalizing $\text{DiffM}(\vec{b}_v)$ by ballot span limits the voter's contribution to a given delta-tally component to ± 1 , since $\max\left(\text{DiffM}(\vec{b}_v)\right) = \Delta_v$.

In the absence of exact ties, the tally vector from (1.5) always yields a distinctly ordered candidate ranking. The delta-Borda matrix derived from this tally vector contains the same ranking information and, therefore, the win matrix produced by this delta-Borda tally will also yield the same unique, cycle-free candidate ranking. In other words, since Borda tallying is a linear process that does not distort relative preference magnitude information, it will never yield an ambiguous cyclical result. This statement is true for the vector-based tabulation from (1.5) as well as the equivalent difference matrix form in (1.13).

11.1 Example: Real-valued Borda Delta-tally

To demonstrate that (1.13) is equivalent to (1.5), we will perform a real-valued Borda delta-tally using the ballots from Example in 8.1.

Voter 1		Voter 2		Voter 3	
1.0	A	1.0	B	1.0	C
0.9	B	0.2	C	0.8	A
0.0	C	0.0	A	0.0	B

Each voter's contribution to the delta-tally is the voter's delta-preference matrix divided by the span of his ballot. All ballot spans are 1.0 for this example.

Voter 1

	A	B	C
A	0	0.1	1.0
B	-0.1	0	0.9
C	-1.0	-0.9	0

Voter 2

	A	B	C
A	0	-1.0	-0.2
B	1.0	0	0.8
C	0.2	-0.8	0

Voter 3

	A	B	C
A	0	0.8	-0.2
B	-0.8	0	-1.0
C	0.2	-1.0	0

For the Condorcet case in Example in 10.1, only a voter's ordering of the candidates could be determined from his contribution matrix. For the real-valued Borda contribution matrices above, the relative magnitudes of delta-preference are conserved. The aggregate delta-tally is the sum of all voter contributions. As in the Condorcet case, we determine the winner by computing a win-Boolean matrix.

D_{rvborda}

	A	B	C
A	0	-0.1	0.6
B	0.1	0	0.7
C	-0.6	-0.7	0

W_{rvborda}

	A	B	C
A	0	0	1
B	1	0	1
C	0	0	0

As in Example in 8.1, candidate B is the winner. This difference-matrix form for real-valued Borda also demonstrates that determining a winner from a Borda tally is very similar to finding a Condorcet winner. In both circumstances, the winning candidate will have her row in the win-Boolean matrix all 1 (except to diagonal). The key difference is that the linearity in Borda tallying guarantees that it will never yield a cycle.

In the next section, we will combine the transformed Condorcet and Borda tallying formulations in (1.11) and (1.13) to reduce the prevalence of cyclical outcomes.

12 Moderated Differential Pairwise Tallying: A Hybrid

We will now bring together the desirable properties of Condorcet and real-valued Borda tallying, balancing Condorcet's strict candidate pair dependence with real-valued Borda's undistorted transmission of relative priority information. A significant form of information loss in Condorcet's pairwise comparison is the removal of degree for each voter's smaller differential preferences. Cycles are more prevalent due to information loss when no distinction is made between candidates far apart versus close together on a voter's ballot. Making use of the similar structure of (1.11) and (1.13), we can address this shortcoming in Condorcet's method by forming a parameterized hybrid of Condorcet and Borda tallying using a linear sigmoid. A linear sigmoid introduces a proportional, sloped linear region around near equal preference to the classic signum from (1.11). This sloped region will address issues caused by the hypersensitive step transition in a signum function. We define this linear sigmoid function as

$$\text{linsgn}(x,h) = \begin{cases} \dfrac{x}{h} & \text{if } \mid x \mid < h \\ sgn(x) & \text{if } \mid x \mid \geq h \end{cases} \quad (1.14)$$

The parameter h is the half-width of the linear region of the sigmoid, with the equation's conditional written in terms of the magnitude (absolute value) of x. As $h \to 0$, the linear region vanishes towards the signum's step discontinuity.

Using this linear sigmoid we can insert a parameterized Borda-like proportional region into the saturated, binary comparison of Condorcet's method. To control the width of this linear region we introduce the *moderation span*, m_v. This voter specified parameter allows each voter to choose where on the tallying continuum between the Condorcet and Borda methods her ballot will be tallied. Moderated differential pairwise tallying (MDPT) can be written in matrix form as

$$D_{\text{Mod}} = \sum_{v}^{\text{all voters}} \text{linsgn}\left(\text{DiffM}(\vec{b}_v), m_v\right) \quad (1.15)$$

We can also compute the same moderated delta-tally element-by-element and with less abstraction,

$$D_{\text{Mod}}[\text{A},\text{B}] = \sum_{v}^{\text{all voters}} \begin{cases} \dfrac{\vec{b}_v[A]-\vec{b}_v[B]}{m_v} & \text{if } \mid \vec{b}_v[A]-\vec{b}_v[B] \mid < m_v \\ \operatorname{sgn}\left(\vec{b}_v[A]-\vec{b}_v[B]\right) & \text{if } \mid \vec{b}_v[A]-\vec{b}_v[B] \mid \geq m_v \end{cases} \tag{1.16}$$

As they are equivalent, both (1.15) and (1.16) divide the difference in candidate preference values by the voter's moderation span, m_v, reverting to the previous Condorcet formulation when the difference between the candidates is greater than m_v. Regardless of any voter's moderation span or candidate placement, each matrix element in D_{mod} depends only on the pair of candidates in question. Because of this property, MDPT possesses the same property of strict candidate pair dependence as Condorcet's classic pairwise tallying.

12.1 Example: Moderated Differential Pairwise Tallying

For our example of MDPT we will use the collection of ballots from Examples 5.1 and 8.1. With pairwise tallying these ballots produced an ambiguous cycle, while a winner was found using real-valued Borda. For this example we will set each voter's moderation span to half of their ballot span, 0.5.

Voter 1		Voter 2		Voter 3	
1.0	A	1.0	B	1.0	C
0.9	B	0.2	C	0.8	A
0.0	C	0.0	A	0.0	B

We next compute the moderated delta-preference matrix for each voter. All preference differentials smaller than 0.5 are moderated.

Voter 1

	A	B	C
A	0	0.2	1.0
B	-0.2	0	1.0
C	-1.0	-1.0	0

Voter 2

	A	B	C
A	0	-1.0	-0.4
B	1.0	0	1.0
C	0.4	-1.0	0

Voter 3

	A	B	C
A	0	-1.0	-0.4
B	1.0	0	-1.0
C	0.4	1.0	0

The sum of these voter contributions produces the moderated delta-tally from which we can also determine a win-Boolean matrix.

D_{mod}

	A	B	C
A	0	0.2	0.2
B	-0.2	0	1.0
C	-0.2	-1.0	0

W_{mod}

	A	B	C
A	0	1	1
B	0	0	1
C	0	0	0

The winner for this moderated example is candidate A. This hybrid method produced a different result than the Borda and Condorcet methods. With full ballot span normalization, the winner for these ballots is B as shown in Example in 11.1. With Condorcet's pairwise tallying in Example in 5.1, these ballots produced a cycle. To resolve the cycle, the two voters on the winning side of one pairwise contest in the Condorcet tally need to moderate sufficiently to change the pairwise result. In this example, Voter 2 and Voter 3 both indicated that they found A and C relatively similar, allowing A to win. If the voters' moderation spans were expanded to 1.0, then the result of the sub-contest between A and B would also flip and B would be the winner. This result would be the same as the Borda case since moderated tallying with all $m_v = 1.0$ is equivalent to real-valued Borda.

We will discuss some details of how this new vote tallying mechanism works and why a voter might choose to moderate in the following sections.

13 Moderation as a transfer function

Moderation in MDPT can be understood as follows. If option A is much more preferable than B, the signum function limits the support of A with respect to B to $+1$. Conversely, if B is much more preferable than A, then the differential opposition of A to B is similarly limited to -1. If the difference in preference between A and B is less than the voter's moderation span, then the expression of relative support or opposition between the candidates will be

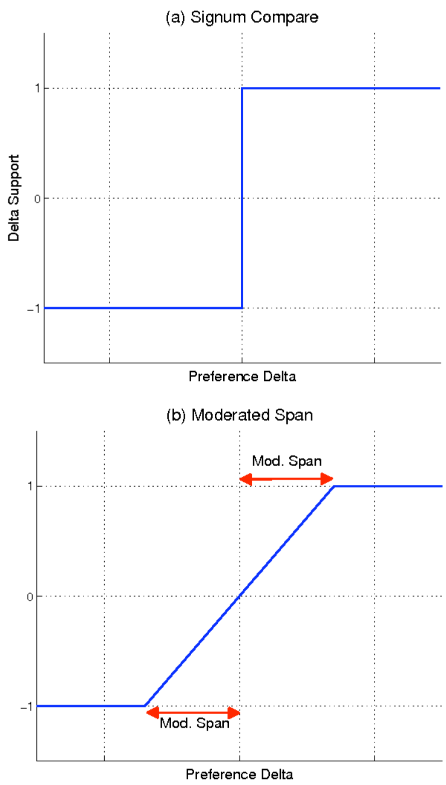

Figure 1.2: Voter support of one candidate over another based on the preference delta of those candidates on the voter's ballot. The left graph shows this transfer function for Condorcet pairwise comparison, the right for moderated pairwise comparison.

less than the voter's full weight. The moderation span introduces a linear region into a voter's contribution to the delta-tally matrix. With this addition to pairwise tallying, the voter can choose to moderate his expression of relative support/opposition for alternatives he finds nearly equally preferable. Fig. 1.2 shows a voter's differential support versus delta-preference for pairwise tallying using both Condorcet's classic quantization and MDPT.

When m_v is shorter than the smallest distance between candidates on a voter's ballot, the voter's tally contribution is equivalent to Condorcet's comparison in (1.1). If a voter's moderation span is equal to the whole span of his ballot, the voter's delta-preferences will all fall within the moderated linear region as the magnitudes in the delta-preference matrix are bounded by ballot span. In this circumstance, the voter's tally contribution is equivalent to

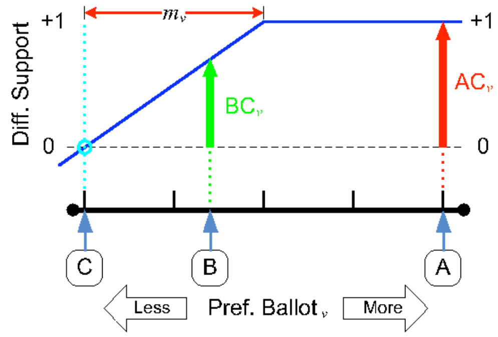

Figure 1.3: Assessing candidates A and B with respect to C, where the voter's ballot is shown below the sigmoid. The voter's contribution into the tally matrix for each sub-contest is represented by the height of the BC and AC arrows. Note that B receives moderated support while the voter exerts full support for A.

that from the linear delta-Borda tally in (1.5). When a voter's span of moderation is set between these two extremes, this partial linearity allows for a hybrid of strong and moderate opinion. The effect of the moderation span on tally contributions is demonstrated in Figs. 1.3 and 1.4.

We note that the transfer function of the linear sigmoid bears some resemblance to the dilating and clipping voter strategy shown in Fig. 1.1. With this strategy a voter effectively created a linear region over the perceived front-runners and compressed the ends of his ballot. In (1.15), however, the $\text{linsgn}(x, h)$ function is applied to a voter's delta-preference matrix $\text{DiffM}(\vec{b}_v)$ instead of directly to the expressed preference ballot vector $\vec{b}_v$. This procedure mitigates the need for one of the two degrees of freedom in such a speculative strategic interval voting strategy: the position of a strategic interval. In the same way that Condorcet's pairwise tallying assesses the relative ranking of all possible candidate combinations independently, MDPT evaluates all transpositions of the voter's desired moderation span centered around each candidate on the ballot. Essentially, m_v defines the extent of a moderation interval around each candidate. The width of the voter's moderation span is the remaining degree of freedom in voting, allowing the voter to choose the level of influence for his smaller differential preferences.

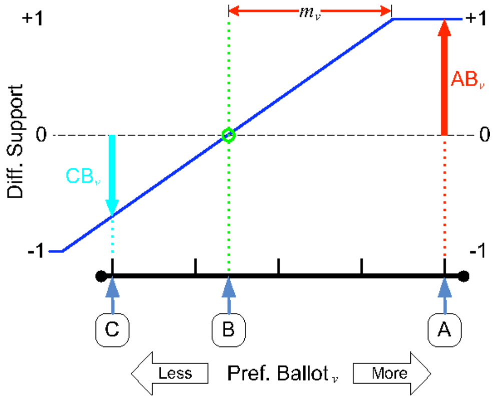

Figure 1.4: To assess candidates with respect to B, the linear sigmoid is transposed to center around candidate B. In sub-contest CB, the voter's opposition to C is moderated to the same degree as in BC from Fig. 1.4 since the delta-tally matrix is anti-symmetric (BC = −CB). In sub-contest AB, A is considered significantly more preferable to B and receives full support.

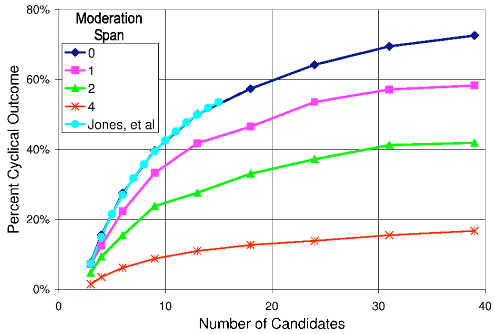

Figure 1.5: This graph shows how the percentage of random elections producing a cyclical result grows with the number of candidates considered for different moderation levels. Each of our data points represents 7,500 elections using seven random ballots. For each election, seven random ballots with uniform candidate distribution were created and then rescaled to span [0,10]. When all moderation spans are 0 (blue diamonds), the results are found to be equivalent to data gathered by Jones et al [9] for classic Condorcet pairwise analysis (light blue circles). When the moderation span for these random ballots is set 40% of ballot span, the data shows a dramatic decrease in the probability of cycles.

14 The cycle reducing effect of moderation

To demonstrate the cycle reducing effect of moderation, we simulated elections using uniform random ballots. Many other voting models for simulating perhaps more realistic elections exist, including issue-space methods [4]. We have chosen instead to use uniform randomness because it readily produces cycles and allows us to show the effect of moderation in the most general way.

For each data point in Fig. 1.5 and 1.6, we computed 7,250 elections. Each simulated election used seven random ballots possessing a uniform candidate distribution. After randomly selecting a preference value for each candidate, we rescaled every ballot to span 0 to 10 so that a particular moderation span had the same meaning for each ballot. Fig. 1.5 shows how the moderated span extension to classic Condorcet pairwise tallying can reduce the occurrence of cyclical results. When voters choose to moderate the expression of their differential preference over candidates they find similarly preferable, cycles are less likely to occur.

For the case when no voters choose to moderate (all $m_v = 0$), Fig. 1.5 shows a sharp growth in the percentage of elections resulting in cycles as the number of candidates increases. One conclusion from this result is that, in classic pairwise analysis, elections with a large number of candidates are much more likely to produce cycles. This result agrees with the data for classic pairwise tallying of random ordinal ballots in Jones et al [9], which is also shown in Fig. 5. For electorates which choose to moderate, however, cycles are significantly less likely to occur even with more than 30 candidates.

Fig. 1.6 presents another view of this same cycle probability simulation. This view shows that as moderation increases the probability of a cycle goes to 0 even for large candidate fields.

15 Discussion

The addition of the moderation span addresses an important shortcoming of Condorcet's pairwise tallying. Although classic pairwise tallying is hardened against manipulation, it does not allow voters to express any difference between their various priorities. All delta-preference magnitudes are treated as the same. We would like to suggest three ways of interpreting this new concept of moderation.

First, voters may simply wish to express slight

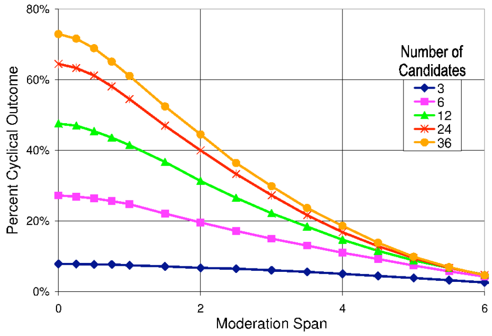

Figure 1.6: This graph shows how the probability of a cyclical outcome decreases with increasing moderation span for several different numbers of candidates. The results were produced using the same uniform random ballot method as Fig. 1.5 with seven ballots. As the moderation spans grow, the prevalence of cycles drops off towards 0. When all moderation spans cover the voters full ballot ($m_v = 10$in this case), the computation is equivalent to a real-valued Borda tally and will never result in a cycle.

preferences. Some Condorcet methods allow voters to express that they consider a pair of candidates equally preferable. Moderation extends this idea to create a smooth continuum between considering two candidates equally preferable and expressing full support for one over the other. For example, a voter might strongly prefer A over C, but have only a small preference for A over B. The moderation span permits the voter to express these differences in priority. Such slight preferences could represent some form of uncertainty a voter has for whether A or B is actually better.

A second perspective on the moderation span is that it gives voters the freedom to choose not to strategically maximize their voting influence over candidate pairs they find similarly preferable. At $m_v = 0$, MDPT tallies a voter's ballot using Condorcet's pairwise analysis, meaning all delta-preference magnitudes are maximized. When m_v is equal to ballot span, only the single largest delta-preference on the voter's ballot is maximized. Moderation gives voters control over the level of strategic maximization for their ballot. In fact, a voter could choose to expand m_v beyond the size of their ballot, which brings us to our final comment.

The third interpretation of moderation we suggest is that it allows a form of voluntary *interpersonal comparison of utilities*. In Section 7 we described an idealized, cardinal utility method for making social choices. The issue which makes this method unusable is well known in economic theory: while individuals can determine the relative costs and benefits of different potential outcomes for themselves, there is no general, well-defined way of comparing the utilities of individuals. However, as discussed by Sen [17], the common background and experiences of members of a society do allow at least a limited form of interpersonal comparison of utilities. While pairwise analysis operates on the premise that all voters will try to maximize their own influence over a decision, the voter-specified moderation span leaves open the possibility that whole communities will be able to see their individual preferences in a broader perspective. When many voters choose to vote moderately, the group can make decisions with more of a Borda-like, shared benefit-cost perspective. Some voters may recognize they do not have as much at stake in a particular decision as others and perhaps then set their moderation spans greater than the span of their ballot. Although the reality of contentious and consequential elections requires that any well-formed vote tallying method be hardened against manipulation, voluntary moderation enables more moderate groups of people to also use pairwise analysis.

The concept of moderation also seems applicable even in the case where there are only two candidates on the ballot. In Sections 2 and 4 we discussed how picking between just two candidates avoided complications from the spoiler effect and was therefore fairly straightforward. However, as we have just described, there are circumstances where a voter may wish to express a slight preference for one candidate over another. The moderation enhancement to classic pairwise tallying in MDPT gives voters this flexibility, even in the two-candidate scenario.

When voters are provided the freedom to express moderate opinion, we assert that a candidate who emerges on top of all head-to-head comparisons with every other candidate under consideration is distinctly the best choice. We term such a candidate a *moderate Condorcet winner*. If no voters choose to moderate, then the moderate Condorcet winner is equivalent to the classic Condorcet winner. When voters do choose to moderate, then the moderate and classic Condorcet winners will sometimes differ. In particular, as we showed in Section 14, a moderate Condorcet winner occurs more often than a classic Condorcet winner.

We refer to a voting method as *moderate Condorcet winner definite* if it always selects the moderate Condorcet winner when one exists. MDPT is intrinsically moderate Condorcet winner definite. Since the moderation span gives voters control over the strategic expression of their ballot, we claim that this proposed criterion is an improvement over the classic Condorcet winner criterion and thus should be a requisite property of any well-formed social choice function.

That said, MDPT does not on its own constitute a complete social choice function. Cyclical majorities can still occur, particularly if voters choose not to moderate in contentious decisions. It would certainly be possible to use a cycle breaking scheme like those proposed by Llull [8], Condorcet [21], Tideman [18], or Schulze [15] on top of MDPT. Since our approach uses real-valued preference ballots, the cardinal-weighted approach proposed by Green-Armytage [7] could also easily be employed. However, as we alluded to in Section 6, we believe that the emphasis in designing a cycle resolution scheme should be on minimally relaxing independence from irrelevant alternatives. Most existing cycle resolution methods are defined in terms of properties of the aggregate tally. We suggest instead that an approach could be built using the concept of moderation. Such a method would resolve cycles by removing edges from the win-edge graph based on which candidate pairs individual voters find most similarly preferable. A voter priority driven method of this type could more directly minimize the amount of compromise individual voters would need to make for the group to reach a coherent decision.

16 Practical Aspects

In the thick of the mathematical detail of our method, it can be difficult to keep track of the bigger picture question: how would it work in actual elections? In this section, we take a step back to address this more practical question. To begin, we will describe how a moderated preference ballot would be cast. We will then describe how elections could be set up to work with our system. Finally, we will discuss some properties of MDPT which should make a transition to this new system easier.

16.1 Casting a moderated ballot

We will now describe how a voter might cast a real-valued preference ballot with a moderation span. This new form of voting will require a new user interface, but with a good design we believe it will be quite intuitive. Our vision for this new interface involves the use of sliders to move candidates up or down on the ballot. These sliders could either be mechanical sliders like those in the figure below, or graphical sliders on a computer screen controlled with a mouse or a touchscreen. Forming a ballot would then be a simple matter of pushing the sliders up or down until the voter is happy with the positions of the various candidates.

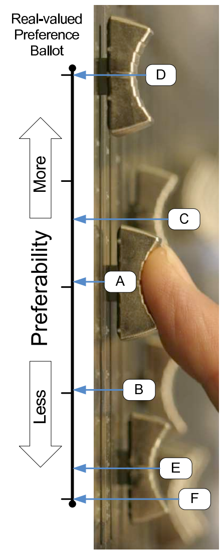

As an example of how a voter could shape her real-valued preference ballot, consider the following scenario:

16.2 Example: Casting a moderated ballot

After careful consideration, a voter has determined her preferences for six candidates in an election. She might start creating her ballot by simply placing the six candidates in order, with her most preferred at the top of her ballot and an equal spacing between the rest, as shown in Fig. 1.7 (a). Suppose she has a strong preference for D, C, or A over any of the other candidates, but does not have strong preferences between those three. She would then separate D, C, and A from the rest and shrink the space between them, as shown in Fig. 1.7(b). Next, if she particularly dislikes candidate F, then she would move F further down (Fig. 1.7(c)).

After placing the six candidates, the voter has two further decisions to make: (1) how to set her moderation span to indicate which candidate pairs she finds similarly preferable, and (2) where to place a default value which would be given to any unrated candidates.

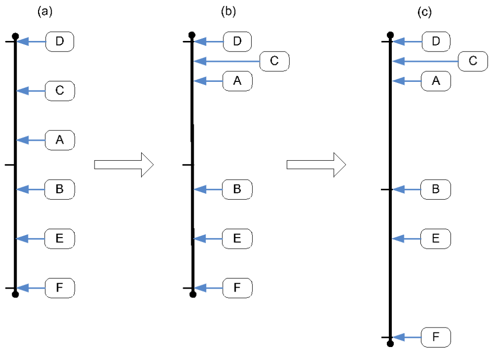

Figure 1.7: Shows the incremental construction of the ballot for Example in 16.2.

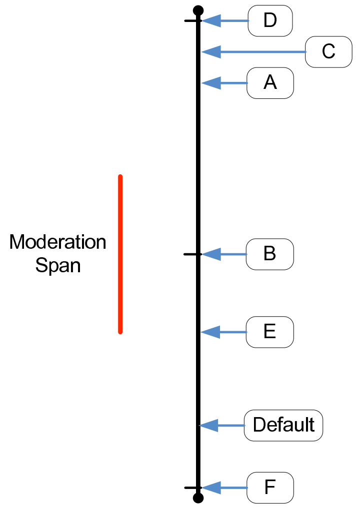

	A	B	C	D	E	F	Def
A	-	1.0	-0.2	-0.4	1.0	1.0	1.0
B	-1.0	-	-1.0	-1.0	0.5	1.0	1.0
C	0.2	1.0	-	-0.2	1.0	1.0	1.0
D	0.4	1.0	0.2	-	1.0	1.0	1.0
E	-1.0	-0.5	-1.0	-1.0	-	1.0	0.6
F	-1.0	-1.0	-1.0	-1.0	-1.0	-	-0.4
Def	-1.0	-1.0	-1.0	-1.0	-0.6	0.4	-

Figure 1.8: Shows the final ballot and moderation span for Example in 16.2, including the contribution matrix for the ballot. The moderation span scales all shorter preference differences on the voter's ballot.

Based on her final ballot rankings, Fig. 1.7(c), the voter would set her moderation span to a distance approximately equal to the space between candidates A and B. This span would indicate her strong preference for the higher of two candidates on her ballot separated by that distance, like A compared to B, or E compared to F. It would also indicate her smaller preference for candidates placed more closely on her ballot, like D compared to C or A.

The final step is that the voter should also specify a default value which will be given to any candidates she has not rated. This default value will usually be near the bottom of the ballot, with only candidates that the voter strongly dislikes below the default. Since this voter finds F particularly unpreferable, she would place her default value between E and F. If there are other candidates in the election, G and H, for example, the system would register that this voter prefers them both over F, even though she has not explicitly ranked them on her ballot. Our voter's ballot is now complete. Fig. 1.8 shows the final ballot at left, as well as how the ballot would be tallied, at right.

16.3 Setup of an election

One of the main goals of this method is to enable voting over a wide range of options. As discussed in Section 2, methods built on pairwise analysis eliminate the need for primaries or other methods of limiting available options. By reducing the potential for cyclical majorities, MDPT allows voters the widest possible range of alternatives. However, when many options are available, voters have to put in greater effort to determine their preferences. It is therefore fundamentally important that voters have easy access to clear information about all alternatives. In circumstances where many voters do not have the time or expertise to form considered opinions, we suggest the use of a *delegable proxy representation system*, similar to that proposed by Green-Armytage [6]. Such a system would allow voters to proxy their voting weight to the representative of their choosing for a given issue, achieving a free-form proportional representation structure and a more responsive democratic process.

In addition, when there are many candidates on the ballot, there must be a way to handle the candidates that a voter does not wish to rate. As mentioned in the description of how to cast a moderated ballot, our suggestion is to allow voters to specify a default value. This default value would be assigned to any unrated candidate. Again, we anticipate that this default value would typically be placed at or near the bottom of the voter's ballot.

16.4 Transition to MDPT

A transition to this new system would require an adjustment for voters who are used to picking a single candidate. Significant voter and poll-worker education campaigns would be necessary for all voters to feel comfortable with this or any new voting system. However, MDPT has a couple features which should make the transition easier.

The first attractive feature of the new system is the removal of the need to vote strategically based on which candidates are top contenders. MDPT would allow a voter to support any candidate of his choice - regardless of that candidate's popularity - without feeling that his vote might be "thrown away." This innovation means that voting can be a simple expression of preference. With our method, voters can feel they have greater choice in an election and that their vote is, therefore, more meaningful.

Flexibility is the second feature of this voting system which should ease any transition. In addition to the moderated preference ballot described above, voters may use our system to cast other styles of ballots with which they are more familiar. If a voter wants to vote only for his favorite candidate, then he can simply put that candidate at the top of his ballot and leave all others at the bottom. A voter who prefers to do a basic ranking of candidates could list them in his preferred order and set the moderation span to 0. We believe voters will appreciate this freedom to express their preferences, and that this appreciation will translate into a better experience while voting and a greater value being placed on the democratic process.

17 Conclusion

In this paper we have presented moderated differential pairwise tallying (MDPT), a per-voter hybrid of the methods of Condorcet and Borda. The foundation for this method is based on Condorcet's pairwise tallying, which has the important property of strict candidate pair dependence. As we described, however, the classic formulation of pairwise tallying discards all voter priority information. This information loss can cause ambiguous, cyclical results for some collections of ballots. At the other end of the spectrum, we examined a real-valued Borda method which always returns a coherent result. The necessary division by ballot span in Borda methods introduces dependence on irrelevant alternatives and encourages speculative voting strategies. To add some Borda-style linearity to pairwise tallying, we developed the voter-specified moderation span. As we have shown, for electorates that choose to widen their moderation spans, cycles will occur less frequently and group consensus will be easier to find.

The introduction of the voter-specified moderation span, in conjunction with real-valued preference ballots, is an important enhancement to Condorcet's pairwise tallying method. Providing voters the freedom to express moderate differential preferences partially addresses the critical information loss issue with classic pairwise tallying. We also proposed to replace the classic Condorcet winner definite criterion for voting methods with a new moderate Condorcet winner definite criterion. Often, cycle resolution will not even be necessary when voters choose to moderate over their diverse opinions. It is only when voters choose not to moderate in contentious decisions that the remaining potential for cycles requires some additional resolution. In the next section we will discuss how the tools we have presented in the paper could be extended to resolve cycles or provide a framework for more directly comparing social choice functions.

18 Future Research

We believe there are some intriguing directions for further research based on the material in this paper. Through the use of real-valued preference ballots and pairwise delta-tallying, we have expressed Condorcet and Borda's tallying methods in a unified framework. This perspective highlighted the similarities and differences of these methods. It appears that other common social choice methods can also be expressed using real-valued preference ballots and pairwise delta-tallying. This delta-preference framework is a potential foundation for a generalized approach for comparing social choice methods.

The concept of individual moderation introduced in this paper provides a new foundation for constructing a complete democratic group decision system. As discussed in Section 15, the additional needed component is a cycle resolution method which minimizes dependence on less-relevant alternatives. With such a system, the spoiler effect would be a thing of the past. Primaries and other methods of artificially pruning the scope of alternatives under consideration would also no longer be necessary.

While there remains work to be done, the potential for significant positive social impact from advances in social choice theory cannot be overstated. We are grateful for the encouraging style at the close of Arrow's Nobel Prize lecture and we would likewise encourage others in this quest to further understanding in this vital and challenging field.

19 Acknowledgements

The authors would like to thank the many people who reviewed drafts of the paper for their ideas and helpful feedback, including Tim Dirks, Jay Hanks, Bruce Carol, Karl Obermeyer, Per Danzl, Jeff Moehlis, and William Durham. The authors would like to especially thank Roger Sewell and James Green-Armytage for their insightful theoretical perspectives on these developments, as well as the anonymous reviewer for his/her insightful comments and focus on the practical side. Finally, both authors would like to thank Dave Daly for his writing expertise, help in forming the framework for this paper, and bringing focus to the relation to Arrow's properties.

20 References

[1] Arrow KJ (1963) *Social Choice and Individual Values*, 2nd edn. Yale University Press, New Haven.

[2] Arrow KJ (1972) "General Economic Equilibrium: Purpose, Analytic Techniques, Collective Choice." For URL, see McDougall site resources page.

[3] Black D (1998) *The theory of committees and elections*, 2nd edn. Kluwer Academic Publishing, Boston.

[4] Chamberlin J, Cohen M (1978) "Toward Applicable Social Choice Theory: A Comparison of Social Choice Functions under Spatial Model Assumptions", *Am Political Sci Rev* 72:1241-1256.

[5] Condorcet MJ (1785) *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix: Rendues à la pluralité des voix*, republished by AMS Bookstore (1972)

[6] Green-Armytage J (unpublished) "Representation by Delegable Proxy", For URL, see McDougall site resources page.

[7] Green-Armytage J (2004) "Cardinal-weighted pairwise comparison", *Voting matters*, 19:6-13

[8] Hägele G, Pukelsheim F (2001) "Llull's Writings On Electoral Systems." *Studia Lulliana*, 41:3-38

[9] Jones B, Radcliff B, Taber C, Timpone R (1995) "Condorcet Winners and the Paradox of Voting: Probability Calculations for Weak Preference Orders." *Am Political Sci Rev* 89:137-144

[10] McLean I (1989) *Democracy and new technology*. Polity Press, Cambridge, UK

[11] McLean I, Hewitt F (1994) *Condorcet: Foundation of Social Choice and Political Theory*. Edward Elgar, Brookfield.

[12] McLean I, McMillan A (2003) *The Concise Oxford Dictionary of Politics*. Oxford University Press, New York.

[13] McLean I, Lorrey H (2007) "Voting in the Medieval Papacy and Religious Orders", *Lecture Notes in Artificial Intelligence*, Vol. 4617. Springer Berlin Heidelberg, p 30-44.

[14] Saari D (2001) *Decisions and Elections: Explaining the Unexpected*. Cambridge University Press, New York.

[15] Schulze M (2003) "A New Monotonic and Clone-Independent Single-Winner Election Method." *Voting matters*, 17:9-19.

[16] Schwartz T (1990) "Cyclic tournaments and cooperative majority voting: A solution." *Soc Choice Welf* 7:19-29.

[17] Sen AK (1970) *Collective Choice and Social Welfare*. Holden-Day.

[18] Tideman TN (1987) "Independence of Clones as a Criterion for Voting Rules." *Soc Choice and Welf* 4:185-206.

[19] Tideman TN (2006) *Collective Decisions and Voting: The Potential for Public Choice*. Ashgate Publishing.

[20] von Neumann J, Morgenstern O (1944) *Theory of Games and Economic Behavior*. Princeton University Press.

[21] Young HP (1988) "Condorcet's Theory of Voting", *American Political Science Review* 82, no. 2 pp. 1231-1244.

About the Authors

Peter Lindener

Peter is an independent researcher who has chosen to dedicate his life to the advancement of truly democratic social decision systems. After the tenuous nature of the 2000 U.S. presidential election, Peter started thinking on the deeper theoretical issues of larger scale group decision processes. His reflections on the virtues and flaws of Condorcet and Borda's tallying methods provided the initial foundation for the rest of this meaningful collaboration. In addition, Peter's style of always debating the other side of an issue really aided the volleyball nature of this joint work.

Joseph "Joey" Durham

Joey is a PhD student in Control, Dynamical Systems, and Robotics at the University of California, Santa Barbara. While always interested in politics and government, Joey's interest in Social Choice Theory was sparked by conversations with Peter in 2004. Joey's continued insights about the nature of voting's contextual interdependency ultimately led to our current focus on the strategic game aspects of Social Choice. Joey's effort in writing up these shared developments has helped bring a needed clarity to this work.