Foreword

This document contains a collection of promising ideas and reflections stemming from my collaboration with Peter Lindener1 on Social Choice from 2004 until 2009. Peter and I have decided to make these ideas public with the hope that they will inspire others to push the field of Social Choice forward and perhaps pave the way for real change in the way groups of people make decisions. All the technical content presented here was developed and written up jointly.

For different reasons both Peter and I have decided that we will not be actively pursuing the ideas and theory presented here for the foreseeable future. In my case, as a graduate student in a different field I simply to not have the time and energy to sustain two separate research tracks. Peter is an independent researcher and thinker, but unfortunately he has not been able to find the kind of supportive community necessary to maintain this research effort without my participation. While either or both of us may return to these developments at some future date, we would both be happy if you or others find these ideas intriguing and choose to pursue them. If you do find these ideas useful as a basis for your own work, all we ask is that you properly attribute to us whatever you gleaned from our work.

Peter and I had planned a three part paper with the following basic structure:

Part I: Moderation as a voter-specified hybrid of Condorcet's and Borda's methods
Part II: MaxRep: Cycle resolution through global moderation
Part III: Mathematical analysis of the properties of our cycle resolution method

The paper which follows is Part II of this plan. What we had once thought of as just Part I has been published on its own in the journal Voting Matters (J. W. Durham and P. Lindener, "Moderated Differential Pairwise Tallying: A Voter Specified Hybrid of Ranking by Pairwise Comparisons and Cardinal Utility Sums", Voting Matters, Issue 27, September 2009.). While our cycle resolution method builds off the idea for the moderation span in that paper, the addition of moderation to Condorcet pairwise tallying stands on its own as addressing a fundamental flaw with Condorcet's method. This document assumes that you are familiar with the concept of moderation from that paper.

The idea for the moderation span did not change much after its initial development, but our cycle resolution method (Part II) went through a couple complete re-writes and many revisions. Many early versions were based around an idea we called M-way analysis: if a cycle existed in pairwise analysis we would then try all three-way contests, then all four-way contests, etc looking for the equivalent of a Condorcet-winner. For n candidates, n-way analysis is equivalent to tallying with a real-valued version of Borda's method and thus guaranteed to be cycle free. Dave Daly was also an active participant in the development of these M-way ideas. Most of the technical rationales developed during this time eventually formed the foundation for our subsequent developments, but late in 2005 we dropped the rather ad-hoc M-way approach for a much cleaner way to resolve cycles.

We had intended to submit for publication the paper which follows, however, as we were finalizing the Future Research section Peter became convinced there was a way to do it slightly better and we decided to wait. Since that time I have concluded that one of the reasons we went through so many revisions of our cycle resolution method in the first place is we never formally laid out what exactly we were trying to optimize while resolving cycles. The idea behind the MaxRep method is to use the wider ballot context to resolve ambiguous cycles, but to do so with the minimum amount of context which yields a coherent result. While this context will necessarily bring clarity to a tangled mess of pairwise results, it also opens the door for strategic ballot manipulation which does not exist in isolated pairwise contests. A clearer mathematical definition of the goal of optimally resolving cycles is needed before it would be possible to really say whether MaxRep is optimal or what further revisions there might be (different moderation function shapes, for example). Fortunately, this goal may not be as elusive as it sounds.

Over the four or five years Peter and I worked together on this, we thought of several other ideas in this space which are not represented here. For one, the approach to ballot tallying we adopted for this paper and the previous one can represent many different vote counting methods, not just Condorect and Borda. For example, for a typical vote-for-one plurality voting system, a voter's ballot would be a vector with a 1 for their selected candidate and 0's elsewhere. This can be transformed into a pairwise matrix and tallied by matrix addition just like we do in our moderation span paper, and the winner of the plurality election would then be the Condorect winner of the tally matrix result from these particularly shaped ballots. IRV/STV can also be represented by doing this iteratively, reshaping the voter's ballot each time. There seem to be many interesting insights to be gained from comparing how different voting methods process voter preference information using this common framework. More focused on the content of this paper, we figured out ways to directly compute the optimal value of alpha without a linear search. Other ideas we have had include how to setup a floating representational hierarchy by allowing voters to proxy their vote to a trusted representative, and Peter has also been thinking about what prior information and of what reliability a voter needs to be able to game the system and vote strategically (and thus how to create an automatic system to do this for all voters, so no one voter can gain more sway over the decision).

If you find the ideas in our papers interesting, or are intrigued by some of the other topics I just mentioned, please get in touch with us at the email addresses below.

All the best,
Joey Durham

Author contact info:
Joey Durham: joey.durham@gmail.com
Peter Lindener: lindener.peter@gmail.com

Last Change: March 8, 2009

The MaxRep Choice Function:
Voter-Priority-Based Resolution of Cyclical Majorities

Joseph Durham and Peter Lindener

Abstract: This paper proposes a social choice voting system which can assist democratic decision making by finding a group's genuine consensus. In 1785, Condorcet discovered serious obstacles to this goal; in 1951, Arrow's *impossibility theorem* further clarified their nature. We introduce a vote tallying framework which unifies the methods of Condorcet and Borda, both adapted to tally real-valued preference ballots. The proposed ***"MaxRep" choice function*** maximizes the influence of each voter's priorities while yielding a unique outcome. When it is necessary to resolve ambiguous cyclical majorities, MaxRep employs an adaptive proportional perspective of each voter's relative priorities. This method also minimizes the outcome's dependence on less-relevant alternatives, thereby mitigating the need for speculative voting strategies which can otherwise cloud genuine consensus. The authors hope that the continued development of information theory-based democratic tools will be a step toward more congenial and congruous group decision-making.

2.1 Introduction

Although the voter specified moderation span from Part I reduces the potential for coinciding, cyclical majorities, it does not completely eliminate them. Such ambiguous decision outcomes may still occur across multiple pairwise contests, particularly in highly contentious elections where voters may choose to moderate less. In the second part of this paper we introduce the MaxRep choice function, further unifying the tallying methods of Condorcet and Borda. At the core of our new choice function is an adaptive perspective of each voter relative priorities which can equitably resolve any cycles that occur. This method of cycle resolution permits MaxRep to always yield an unambiguous, equitable result.

In Section 2.2 we will start the development of the MaxRep choice function with the idea that sparked the authors' original collaboration in this work in social choice theory. Analyzing Lindener's initial combination of Condorcet and Borda's methods will suggest the conceptual framework for our subsequent developments. In Section 2.3 we introduce the contest-wide tallying parameter α (alpha) which extends Part I's moderation span concept to form a level playing field for any required cycle resolution. Similar to the per voter span of moderation m_v, α spans the continuum between Condorcet's and Borda's tallying methods. The difference between m_v and α is that α applies uniformly for all voters and the resulting α-parameterized tallying formulation forms a level playing field for adaptive cycle resolution. When Condorcet-style influence maximization causes cyclical ambiguity, raising α will resolve these cycles by introducing the pertinent information on each voter's relative priorities.

Section 2.4 introduces a framework of conceptual tools for understanding the topology of cyclical ambiguity, including edge graphs, cycle sets, and partial set ranking concepts. We also describe a matrix-based method for computing these cycle set relationships. We then leverage these concepts in 2.5 to describe three design constraints that together imply an algorithm for computing the MaxRep choice function. These three constraints are derived from the general design goals of consistently maximizing the influence of each voter's priorities while also minimizing any dependence on less-relevant alternatives. We will show that these goals are actually two interpretations of the same objective (minimizing the required value of α) and thus will be achieved simultaneously. Raising α will resolve any cycles but doing so also introduces some level of necessary dependence on the context of the candidate field under consideration. This additional factor leads to the conservative elimination algorithm which we term the MaxRep choice function.

In 2.6 we present a pair of algorithm flow graphs for the MaxRep choice function, with the second having a significant optimization. Section 2.7 illustrates several example decision computations using these algorithms for the MaxRep choice function. First are a few simplified illustrating cases. These basic examples are followed by two decisions with many alternatives where the ballots are generated randomly. These larger examples show how MaxRep typically resolves real decisions with many options under consideration. We then discuss the properties of the MaxRep choice function in 2.8, including a cursory examination of how MaxRep compares to other social choice methods. In 2.8.3 we reflect on how these developments relate to Arrow's original desired properties. We go on to discuss how MaxRep's elimination can be interpreted as a form of strategically negotiated compromise. Finally, we reflect on how MaxRep's minimized dependence on less-relevant alternatives mitigates the need for voter's to adopt a speculative voting strategy in order to increase their influence in a decision.

2.2 Borda over Condorcet's top cycle

At the beginning of this work, Lindener (second author) proposed the following combination of Condorcet and Borda's methods. He suggested that a choice function could first perform classic Condorcet pairwise analysis and then, if necessary, resolve within any top cycle using Borda's ordinal method. In contrast to Black's method (Black 1998), the subsequent Borda evaluation would be done only with the candidates in the top cycle (the Schwartz set) resulting from the initial pairwise evaluation: for the subsequent Borda assessment, the ordinal ballots would be condensed to contain only the top cycle candidates. Lindener's first proposal appears to be an improvement over straight ordinal Borda in that it removes truly irrelevant alternatives before setting the context for subsequent cycle breaking. After the initial pairwise pass, only the candidates in the top cycle are of any relevance for further evaluation. As Saari has observed, classic pairwise tallying is blind to information about relative preference strengths (Saari 2001). Performing a Borda tally over just the candidates in the top cycle reintroduces something resembling this missing information to resolve the top cycle. The MaxRep choice function introduced here develops on the thinking behind Lindener's initial idea with a pair of enhancements suggested by observations developed from Part I.

First, classic Condorcet pairwise tallying will radically distort relative preference magnitudes, potentially amplifying tiny differential preferences to the same level as the voter's greatest priorities. The moderated differential tallying method introduced in Part I addresses this issue by giving voters the freedom to declare which of their relative preferences do not carry full voting weight. The use of real-valued preference ballots also avoids the artificial and adverse constraints that Borda's classic ordinal ballots would place on the voter's ability to express their true relative preference ratings, as discussed in Section 1.4.

In addition, performing Borda-style cycle resolution over a pairwise tally's top cycle causes a sudden jump in the level of candidate field context in the decision process. Pairwise tally matrix elements (such as Condorcet's and those from moderated differential tallying) have strict candidate pair dependence, while Borda-style, ballot-span normalized tallies are dependent on the whole candidate field under consideration. Jumping directly to full Borda-style perspective will usually introduce a greater dependence on less-relevant alternatives than is strictly necessary to bring some order to the top cycle. The following development will establish a smooth continuum of perspective between the Condorcet and Borda extremes. This continuum will provide a level playing field for any required cycle resolution. In this context, MaxRep will find the minimum span of decision perspective that will yield well-formed progress towards resolving any remaining ambiguous cyclical majorities.

2.3 Contest-wide shared moderation parameter

The moderation span introduced in Part I establishes a per-voter continuum between Condorcet's and Borda's style of delta-tally matrix contribution: voters can specify a span of moderation for weighing relative differential preference priorities. To establish a smooth continuum from Condorcet to Borda for the electorate as a whole, we introduce the shared moderation parameter α. In order to define such a shared parameter, we first require a measure of the relative preference scale of each voter's ballot.

2.3.1 Measure of Ballot Scale

The real-valued preference ballots used in our MaxRep choice function are unconstrained, meaning that any candidate can be placed at any position on the real number line. The relative ratios of differential preference between each pair of candidates on such a ballot establishes the relative priority a voter places on each pairwise relationship, as will be described later. To compare relative priorities between voters in a ballot scale normalized manner, it is necessary to measure the intrinsic scale or size of each ballot. For example, in Eq. 5 we used ballot span to normalize each ballot to span 0 to 1.

The use of the ballot span of all candidates under consideration as the measure of ballot scale in Eq. 5 is significant as it possesses an important property. Any measure of the scale factor of a preference ballot will, by necessity, depend on the placement of candidates and therefore will inherently depend on the context of which candidates are under consideration. However, ballot span (the max-min infinity norm) is the only LP norm that does not involve some form of a summation over the candidate field. It is therefore the only such norm that will not vary with the introduction of candidate clones (additional candidates that are effectively equivalent to an existing candidate, as discussed in Tideman 1987 and Schulze 2003). For our algorithm to be immune to clones, we will therefore employ ballot span as our required measure of ballot scale. In our delta preference matrix formulation, ballot span Δ_v is also equal to the infinity norm of the voter's delta preference matrix.

$$\Delta_v = \max\left(\vec{\mathbf{b}}_v\right) - \min\left(\vec{\mathbf{b}}_v\right) = \max\left(\mathrm{DiffM}\left(\vec{\mathbf{b}}_v\right)\right) \qquad (17)$$

Since we are building towards an elimination algorithm, we clarify that for this paper $\vec{\mathbf{b}}_v$ is the vector of real preference values for only the candidates still under consideration. With this clarification, Δ_v is defined as the ballot span of the still relevant candidates. Using this measure of ballot scale, we can now establish a shared, context-adaptive moderation parameter.

2.3.2 Basic α Tallying Concept

As an introduction to the thinking behind the shared moderation parameter α (alpha), consider replacing the voter specified moderation span m_v with a parameterized fraction of the voter's contextually relevant ballot span Δ_v (which was just defined in Eq. 17).

$$m_v = \alpha \cdot \Delta_v \tag{18}$$

Parameter α is dimensioned as a fraction of the span of the candidates still under consideration on each voter's preference ballot. Assigning all voter moderation spans in Eq. 15 to this expression effectively establishes a level playing field of moderation across all the voters participating in a decision.

$$\mathbf{D}(\alpha) = \sum_{v}^{\text{All voters}} \text{linsgn}\left(\text{DiffM}\left(\vec{\mathbf{b}}_v\right),\ \alpha \cdot \Delta_v\right) \cdot w_v \tag{19}$$

In Eq. 19, α is a *contest-wide, shared moderation parameter*: by setting each voter's moderation span to an equivalent fraction of their relevant ballot span Δ_v, α applies uniformly to all voters. Equation 19 is essentially equivalent to rescaling all ballots to span [0,1] for the candidates under consideration and then setting each voter's moderation span to α. This α-parameterized pairwise delta tally also adapts to the current decision context since Δ_v spans only the candidates currently under consideration.

For this preference tallying process as a whole, parameter α establishes a continuum between the tallying methods of Condorcet and Borda. Following from the observations regarding Condorcet and Borda tallying equivalency in Section 1.9, when $\alpha = 0$ this formulation yields classic Condorcet results. Since $\mathbf{D}(0)$ is independent of any voter's ballot span, the results are fully independent of any greater candidate field context. When $\alpha = 1$ this hybrid formulation computes the equivalent results to the differential Borda tally in Eq. 13. At the Borda end of the α continuum, every voter's contribution to $\mathbf{D}(\alpha)$ is their ballot-span normalized delta preference matrix. Since there is no distortion of differential preference ratios with this linear rescaling at $\alpha = 1$, $\mathbf{D}(1)$ will always be free of cyclical majorities.

The introduction of α has provided a mechanism for resolving any cyclical majorities: if a cycle occurs at $\alpha = 0$, there must be a value of $\alpha \le 1$ at which this cycle begins to disentangle. Raising α adaptively introduces the pertinent priority information from each voter required to effectively resolve cycles. By adaptively adding this priority information, α-parameterized tallying can resolve cyclical ambiguity by addressing the loss of information that otherwise would lead to Condorcet's dilemma of coinciding, cyclical majorities.

This initial thought exercise has established the concept of a contest-wide continuum of moderation over the interval $\alpha = [0,1]$ with Condorcet and Borda-style endpoints. Since α applies uniformly across all voters, it provides a level playing field of compromise on which cycles can be resolved. When there is a candidate who wins over all others (a Condorcet winner), no compromise is necessary between voters to make a group decision. However, when cyclical majorities do occur, some form of compromise is necessary to make an effective choice. The shared continuum of moderation defined by α provides a framework that can equitably resolve any cyclical majorities. We assert that the α continuum defines a level playing field for equitably resolving cycles; we term this claim MaxRep's ***level playing field assertion*** and it forms the central tenet of the MaxRep choice function.

If a group chooses to adopt the MaxRep choice function to assist with decision making, they are in essence embracing MaxRep's level playing field for facilitating cycle resolution and reaching compromise. That is, the level playing field assertion represents the social contract that a group has chosen when they employ MaxRep in their quest for more congenial group decisions.

While this shared level playing field for finding compromise is critical for resolving cyclical majorities, the MaxRep choice function should still leave voters the freedom to choose a span of moderation greater than that required by the shared moderation parameter α.

2.3.3 Moderated, α-Parameterized Tallying

To combine the enhancements of Part I's voter specified moderation span and the cycle-resolving α-parameterized tallying concept, we introduce an *effective moderation span* for each voter's ballot. This quantity, $em_v(\alpha)$, is the greater of the voter's specified moderation span and their relevant ballot span multiplied by α.

$$em_v(\alpha) = \max(m_v, \alpha \cdot \Delta_v) \tag{20}$$

Since $em_v(\alpha)$ is set to the greater of the two values, a voter's moderation span will apply unless the α–specified fraction of the voter's ballot span is larger. The introduction of $em_v(\alpha)$ provides the virtues of both voluntary moderation and α-parameterized tallying: a voter can indicate that their smaller preference differentials do not carry full weight while α provides a mechanism for resolving cycles across the whole electorate. Note, however, that while m_v does not introduce any dependence on the candidate field under consideration, if elevated values of α are required to resolve cycles then $em_v(\alpha)$ will depend on the candidate field for some voters. The following equation is an α-parameterized version of the moderated pairwise tallying method in Section 1.8 using $em_v(\alpha)$ for each voter.

$$\mathbf{D}_{Mod}(\alpha) = \sum_{v}^{\text{All voters}} \text{linsgn}\left(\text{DiffM}\left(\vec{\mathbf{b}}_v\right), em_v(\alpha)\right) \cdot w_v \tag{21}$$

When $\alpha = 0$, Eq. 21 is equivalent to moderated differential pairwise tallying as found in Eq. 15 and all voter specified moderation spans will take precedence. At $\alpha = 1$, Eq. 21 is equivalent to the delta Borda tally formulation of Eq. 13 regardless of any voter's chosen moderation span. Between these two extremes, the span of moderation set by α represents a level playing field for voter priority-based cycle resolution, though voters can still choose a higher level of moderation via m_v. To highlight the effect α has on the tally matrix contributions of each voter we will flatten out some of the abstraction in Eq. 21 to produce Eq. 22 below.

$$\mathbf{D}_{Mod}(\alpha)[a,b] = \sum_{v}^{\text{All voters}} \left\{ \begin{array}{ll} \dfrac{\vec{\mathbf{b}}_v[a] - \vec{\mathbf{b}}_v[b]}{\alpha \cdot \Delta_v} & \text{if } \left| \vec{\mathbf{b}}_v[a] - \vec{\mathbf{b}}_v[b] \right| < (\alpha \cdot \Delta_v) \ \& \ (\alpha \cdot \Delta_v) > m_v \\ \dfrac{\vec{\mathbf{b}}_v[a] - \vec{\mathbf{b}}_v[b]}{m_v} & \text{if } \left| \vec{\mathbf{b}}_v[a] - \vec{\mathbf{b}}_v[b] \right| < m_v \quad \& \ (\alpha \cdot \Delta_v) \le m_v \\ \text{sgn}(\vec{\mathbf{b}}_v[a] - \vec{\mathbf{b}}_v[b]) & \text{else} \end{array} \right\} \cdot w_v \tag{22}$$

In this conditional equation, three tallying cases are more clearly delineated. In the bottom “else” case no moderation occurs, the full weight of the voter’s influence is expressed between *a* and *b*. In the middle case, the voter’s expressed delta preference between *a* and *b* falls within the span of m_v (and $\alpha \cdot \Delta_v$ represents a smaller span) and the voter’s moderated contribution remains independent of the candidate field under consideration. However, when $\alpha \cdot \Delta_v > m_v$, the presence of Δ_v in the denominator of the top case introduces a dependence on the voter’s ballot span and thus candidate field context for these tally element contributions. As *α* increases, more voter delta preference matrix elements will fall into this top *α*-moderated case and thus become dependent on the candidate field considered. While increasing *α* increases the level of contextual dependence, it also improves the fidelity of relative preference priority expression. This shift towards the Borda end of the tallying continuum can resolve any cyclical ambiguity.

To analyze the results of the full *α*-parameterized delta tally $\mathbf{D}_{Mod}(\alpha)$ and establish any relative ranking between candidates at a particular value of *α*, we first compute the win Boolean matrix for this value of *α* using an equation similar to Eq. 3.

$$\mathbf{W}_{Mod}(\alpha) = (\mathbf{D}_{Mod}(\alpha) > \varepsilon_{wt}) \tag{23}$$

Win matrix $\mathbf{W}_{Mod}(\alpha)$ contains a 1 (True) in element $[a,b]$ if candidate *a* wins over *b*, and 0 (False) if *a* loses or if they tie at this value of *α*. Note that the diagonal elements of $\mathbf{W}_{Mod}(\alpha)$ are all zeroes where candidates always tie with themselves. The introduction of win threshold ε_{wt} allows flexibility in the definition of the strength of differential supported required for victory.

As tallying parameter *α* increases from the strictly pairwise perspective of relative priorities at *α* = 0 towards the Borda-style full candidate field linear perspective at *α* = 1, some win Booleans in $\mathbf{W}_{Mod}$ will transition. That is, since cycles must vanish as $\alpha \to 1$, for some value of *α* between 0 and 1 there must be a transition in $\mathbf{W}_{Mod}$ that will begin to resolve any ambiguous cyclical majorities. During cycle resolution, our goal remains to consistently maximize the expression of each voter's preference priorities while also minimizing any dependence on less-relevant alternatives. Increasing α will resolve any cycles but also brings with it an inherent increase in dependence on the candidate field under consideration. While this dependence on less relevant alternatives when $\alpha > 0$ is sometimes necessary to resolve cycles, it also suggests that we should proceed cautiously. Before outlining how the MaxRep choice function minimizes this dependence on less relevant alternatives, we will establish a foundation of conceptual tools for understanding win matrix topology.

2.4 Edge graphs and cycle sets

Directed edge graphs are a useful tool for visualizing the combined outcome of pairwise sub-contests, particularly when cycles occur. When one candidate wins over another, we draw a directed edge pointing from the loser to the winner; in our representation, a pairwise tie does not produce an edge between the candidates. The edge graph for a full contest is then the union of the edges from all the associated pairwise sub-contests. If the directed edges for any set of candidates form a loop, then those candidates are said to be in a cycle. Note that a minimum of three candidates are needed to form a cycle since a win-directed edge can only point one way. The simplest example of a three candidate cycle was shown in Fig. 1 which is reproduced below for convenience.

Ballots	Voter 1	Voter 2	Voter 3
First	A	B	C
Second	B	C	A
Third	C	A	B

Win Matrix	A	B	C
A	-	W	L
B	L	-	W
C	W	L	-

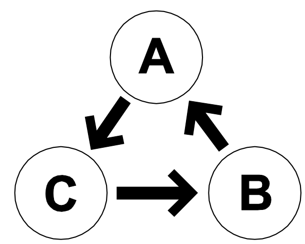

Reproduction of Fig. 1 Pairwise analysis leading to a three-way top cycle. The cyclic nature of the result is clearly shown in the win edge graph at right, with directed edges from sub-contest loser to winner

When there are more than three candidates under consideration, cycles can become quite tangled. As the number of candidates n grows, the number of pairwise edges between them grows as $n(n-1)/2$ and hence the resulting complexity of the edge graph grows of order n^2. This increasing complexity is illustrated by the seven candidate (n = 7) example in Fig. 8.

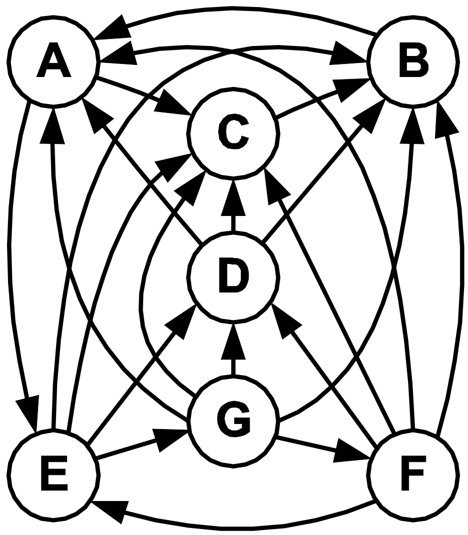

Fig. 8 Seven candidates, with 21 directed edges, all entangled in a single complex cycle. This particular cycle does not have very high redundancy: if candidate A was absent, B would become the Condorcet winner

If edge AE in Fig. 8 were reversed so that A now wins over E, a partial ordering would emerge out of the previously fully entangled cycle, as shown in Fig. 9.

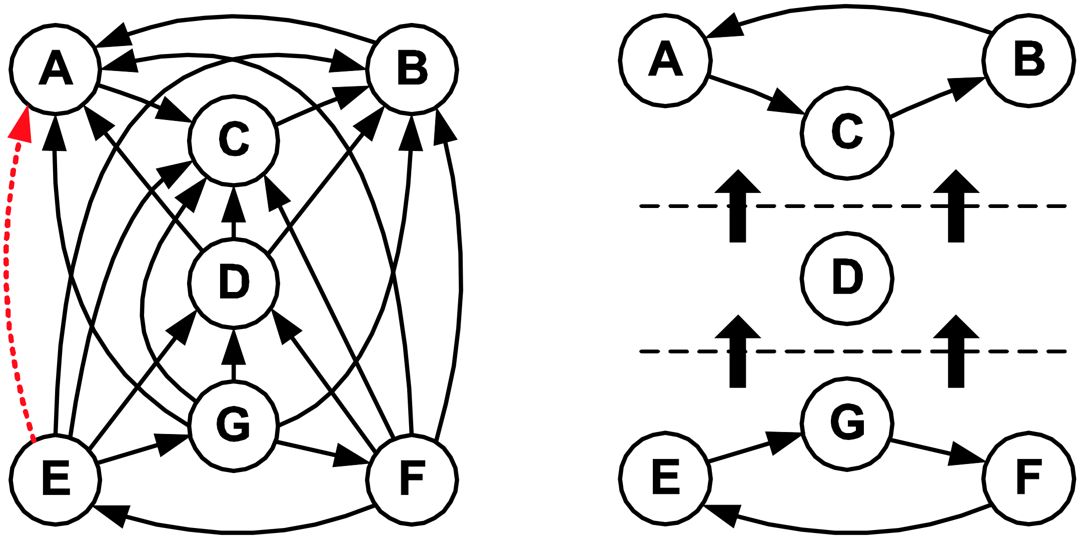

Fig. 9 For this particular example, the reversal of dashed edge AE causes ordering to emerge between subsets of candidates. This emergence of order is more easily seen with the graphical notation on the right

While there remains no single winner, in this example the set {A,B,C} consistently win over all those below them. With this visual representation of cycles and the structure that can emerge with the reversal of even a single pairwise sub-contest, we will now introduce some cycle set terminology that will be used throughout the rest of this paper.

A *top cycle* contains the smallest set of candidates who do not lose to any candidate outside (below) this set, which will sometimes be a single candidate. In Fig. 9, {A,B,C} is the top cycle. The *top cycle set* is defined as the union of top cycles from a win edge graph, often referred to as the *Schwartz set* (Schwartz 1990).[1] Only when there are ties between candidates can there be more than one distinct top cycle in this union. When there is a Condorcet winner in pairwise analysis, the Schwartz set will contain only that candidate.

In Section 1.11 we asserted that, when voters are free to moderate, always selecting the Condorcet winner when one exists is a requisite property for a well-formed choice function. When cycles do occur, we similarly assert that a viable choice function must select a member of the Schwartz set. This property of being *moderated Schwartz definite* implies the prior moderated Condorcet winner definite property since a Condorcet winner would be the only member of the Schwartz set. We also note that the Schwartz set for a Borda-style tally will contain only the candidate who wins the Borda contest. The α-parameterized tallying framework we have just introduced can tally either Condorcet or Borda style results; computing the Schwartz set provides a method of determining the top candidate(s) which works for any value of α.

In addition to the top cycle set, we also have use for the *bottom cycle set* which is the union of all bottom cycles ({E,F,G} in Fig. 9). Candidates in the bottom cycle set do not win over any candidate outside the set. We will use the term *middle cycle set(s)* to refer to the sets of candidates neither in a top or bottom cycle ({D} in Fig. 9).

Cycle sets can be computed from win-Boolean matrices (determined via Eq. 23), which contain the edge-graph information in an equivalent matrix from. To compute the top cycle set from this win-Boolean matrix, we first compute an associated Boolean *connectivity matrix*. This matrix will contain information on which pairs of candidates can be traversed between along a chain of directed edges. For example, in Fig. 9 a chain of directed edges exists from A to B but not from A to E.

The first step in this computation is to compute the one-hop connectivity matrix, which contains a 1 (True) where ever a directed edge chain of length 1 or 0 leads from one candidate to another. To produce this matrix, we OR together the transpose of $\mathbf{W}_{Mod}(\alpha)$ and the Boolean identity matrix. ORing in the identity matrix includes the 0-hop connections on the diagonal where candidates are connected to themselves.

$$\mathbf{C}_{\text{one-hop}} = \mathbf{W}_{Mod}^{t} \mid \mathbf{I} \qquad (24)$$

Squaring $\mathbf{C}_{\text{one-hop}}$ yields the zero through two hop matrix. This matrix multiplication must be done as a Boolean operation to yield a matrix of only 0's and 1's. When $\mathbf{C}_{\text{one-hop}}$ is raised to at least the power $(n-1)$, where n is the number of candidates, it produces the ultimate connectivity matrix $\mathbf{C}$. When using Boolean matrix multiplication, raising $\mathbf{C}_{\text{one-hop}}$ to a higher power will produce no further change in the resulting matrix.

[1]The Schwartz set differs from the Smith set only in its handling of ties: ties are considered edges in the Smith set, but not in the Schwartz set.

$$\mathbf{C} = \mathbf{C}_{\text{one-hop}}^{n-1} \tag{25}$$

Element $\mathbf{C}[a,b]$ is 1 (True) wherever there exists a direct edge chain of any length that leads from candidate b to candidate a, or equivalently if a beats someone, who beats someone, who beats someone, etc, who beat b. For a candidate to be a member of the top cycle set, they must either be connected in both directions or not connected at all to every other candidate. From the directed edge graph perspective, to be a member of the top cycle set a candidate must be able to, for all other candidates, either get back from the other candidate or not be able to get to them in the first place. To compute what we term the *Schwartz matrix* $\mathbf{M}_{Schwartz}$, we OR together the Boolean inversion and transpose of the ultimate connectivity matrix $\mathbf{C}$.

$$\mathbf{M}_{Schwartz} = \overline{\mathbf{C}} \,|\, \mathbf{C}^{T} \tag{26}$$

Candidates whose corresponding row in $\mathbf{M}_{Schwartz}$ is all True are in the top cycle (Schwartz) set. There will always be at least one member in this set. The MaxRep choice function will also have use for the bottom cycle set, whose candidates have their corresponding column in $\mathbf{M}_{Schwartz}$ all True. If all candidates are entangled in just a single cycle, they will all be in the top (as well as bottom) cycle set and $\mathbf{M}_{Schwartz}$ will be matrix entirely full of Trues. With these methods for computing top and bottom cycle set membership, we can now describe the design of the MaxRep choice function.

2.5 Algorithm design constraints

The α-parameterized tallying method in Eq. 21 (and its equivalent formulation in Eq. 22) forms the foundation of the MaxRep social choice function. Parameter α, introduced in Section 2.3.2, provides a degree of freedom when tallying ballots which can equitably resolve cyclical majorities. Increasing α introduces pertinent information on each voters delta preference priorities, shifting tallying towards the Borda end of the α continuum. We will now detail our reasoning on how to best control this tallying parameter. MaxRep is designed to resolve cycles while maximizing the influence of each voter's preferences and also minimizing any dependence on less-relevant alternatives. Three design constraints will combine to imply the algorithm for the MaxRep choice function: iterative candidate elimination at the minimum coherent value of α.

2.5.1 Minimize Alpha

Parameter α has two important properties. First, minimizing the value of α at which decisions are made uniformly maximizes the influence of each voter. Since α is in the denominator of the top condition of Eq. 22, larger values of α reduce each voter's influence over their smaller delta preference matrix components by expanding the width of the moderation sigmoid's linear region. Therefore, making a decision at the minimum possible value of α maximizes each of the contributions from every voter to the resulting delta tally matrix. This maximization of influence over the candidate field under consideration reduces any potential gain from speculative voting strategies (such as the dilating and clipping strategy discussed in Section 1.4).

The second desirable property of minimizing α is doing so also reduces the dependence of the resulting delta tally on the candidate field under consideration, as discussed in Section 2.3.3. When $\alpha = 0$, Eqs. 21 and 22 are independent of the span of any voter's ballot since they essentially perform moderated differential tallying (developed in Section 1.9). Because of this, when $\alpha = 0$, the delta tally matrix elements produced by Eqs. 21 and 22 exhibit strict pairwise dependence. However, as the value of α increases, more voter delta preference matrix elements will be tallied using the top condition in Eq. 22 and will therefore become dependent on the voter's ballot span and hence the candidate field under consideration. When a voter's specified moderation span is greater than $\alpha \cdot \Delta_v$ the voter's tally contributions will remain independent of candidate field context. If $\alpha \cdot \Delta_v$ grows beyond m_v, then only the candidate pairs separated by less than $\alpha \cdot \Delta_v$ are in the linear region of the moderation sigmoid and are thus dependent on the greater candidate field. As α increases from 0 to 1, the number of candidate pairs in each voter's tally contribution that are dependent on ballot span increases monotonically. Therefore, making decisions at the minimum possible value of α minimizes any net dependence on candidate field context and thus any less relevant alternatives currently under consideration.

Minimizing the value of α at which our algorithm makes decisions achieves these two goals simultaneously: maximizing voter influence while also minimizing any outcome dependence on less relevant alternatives. The rationale for minimizing the value of α at which our algorithm makes decisions can also be seen through another interpretation. Raising α expands a voter's effective moderation span causing their delta tally contribution for the candidate pair closest together on their ballot to be reduced first. In other words, expanding α first moderates a voter's differential support between candidates they find most similarly preferable. When cycles exist and no clear winner can be initially determined, some amount of voter priority-based cycle resolution is necessary to find an equitable solution. Raising α resolves ambiguous cyclical majorities by selectively moderating over only each voter's lower priorities. By finding the minimum value of α where progress towards an eventual choice can be made, our algorithm will find the coherent decision where each voter has to compromise the least. Building on the observations surrounding the introduction of α in Section 2.3.2, decisions at min-α represent the most equitable compromise for resolving cycles: since all voters moderate over the same fraction of their relevant ballot span, each voter will compromise over their relative priorities no more or less than any other.

Constraint #1: **Minimize alpha**
The MaxRep algorithm should make decisions at the minimum possible value of α that will permit making coherent progress towards a choice.

To satisfy this constraint, the MaxRep algorithm will start at $\alpha = 0$ and search upwards until it finds the minimum value of α where not all of the alternatives remaining under consideration are entangled in a single cycle set. Unless the candidates are truly tied, some ordering will emerge as α is swept between 0 and 1 since there can be no cycles at $\alpha = 1$. At the minimum value of α where cycles begin to resolve the candidate field will separate into two or more distinctly ordered sets: a top cycle set, a bottom cycle set, and sometimes the remaining candidates in between (middle cycle set(s)). We refer to this lowest separation value of α as *min-α*. While there may remain no coherent ranked ordering inside each of these cycle sets at min-α, there will be an unambiguous ranking between them. This partial ordering scenario is demonstrated in Fig. 9. The next constraint will clarify the actions the MaxRep algorithm will take at min-α.

2.5.2 Eliminate least preferable candidate(s)

At min-α, none of the candidates in the bottom cycle set will win a pairwise contest with respect to any candidate in the cycle set(s) above them. For this candidate field, this separation is the candidate set ranking which requires the minimum amount of compromise between voters. While it may not yet be possible to determine the top choice, such partial set ranking means we can reduce the size of the candidate field. By eliminating only the least preferable candidate(s) at min-α we adhere to Constraint #1 while also making coherent progress towards an eventual choice.

When min-α is greater than 0, for some voters $em_v(\alpha)$ will depend on their ballot span. This necessary introduction of sensitivity to ballot span means that the results of the election will depend on the candidate field considered. While this dependency may be required to make a coherent decision, the interdependence between candidates also constrains how our algorithm should handle candidate sets at elevated values of α. We thus observe that to minimize the choice function's dependence on any less relevant alternatives we should drop only the least preferable candidate(s). If only the lowest candidate(s) are dropped then this elimination decision will be framed only by candidates considered more or equally preferable to those that will be removed. All other possible decisions that could reduce the size of the candidate field under consideration would be framed by some candidates deemed less preferable than the candidate(s) affected. We note, however, that when separation occurs at $\alpha = 0$ there is no dependence on the candidate field under consideration. We will take advantage of this independence of context at $\alpha = 0$ to improve the computational efficiency of the MaxRep algorithm subsequent to its initial development below.

Constraint #2: **Eliminate only the least preferable candidate(s)**
When the MaxRep algorithm makes an elimination decision, it should eliminate only the candidates in the bottom cycle set.

Though the elimination of only the bottom cycle set is the most conservative change to the candidate field at this value of min-α, it will sometimes drop multiple candidates at once. While raising α beyond min-α could eventually clarify ordering within the bottom cycle, continuing to elevate α would violate Constraint #1. Removing the whole bottom cycle set at min-α is the unique cycle-resolving decision which both requires minimal compromise and entails minimal dependence on the candidate field under consideration. The alternative — raising α to determine a single candidate to eliminate — would violate both our goals of maximizing voter influence and minimizing dependence on candidate context. For this reason, making a candidate elimination decision at min-α takes precedence over raising α in an attempt to remove fewer candidates.

By conservatively removing just the bottom cycle set candidates from consideration, the MaxRep algorithm minimizes the influence of any less-relevant alternatives while still making well-formed progress towards an eventual choice. After the elimination of these least preferable candidate(s), any subsequent evaluation should only consider those candidates who are still in contention. The third and final design constraint will clarify how to assess the value of α after the elimination of a bottom cycle set.

2.5.3 Reassess min-α for new candidate field context

Once a candidate has been eliminated from consideration, they should not affect any subsequent progress towards an eventual choice. After an elimination, the span of relevant candidates on each voter's ballot must be retabulated since the prior value may have been set by a now eliminated candidate. Furthermore, if the candidate(s) was removed at $\alpha > 0$ then some voters' tally matrix contributions were computed using the top condition in Eq. 22. In this situation the outcome of the elimination decision was dependent on at least one voter's ballot span. Due to this dependence on the prior candidate field, MaxRep must reassess min-α for any subsequent decision. It is entirely possible that the reduced candidate field will resolve on the next iteration at a lower value of α than was necessary to bring partial ordering to the prior cycle. With the removal of a contender and their associated edges from the edge graph, there may now even be a Condorcet winner in this smaller candidate field. Therefore, we must reevaluate min-α after each reduction to the candidate field so that our choice function consistently maximizes voter influence and minimizes any dependence on less-relevant alternatives as we converge towards a final choice.

> Constraints #3: **Re-evaluate min-α after each reduction of the candidate field**
> When the candidate field under consideration changes, the value of min-α must be reassessed in this new context by restarting the search for min-α from $\alpha = 0$.

The minimum value of α where the candidate field separates depends on the candidate field considered and thus min-α must be recomputed when the candidate field shrinks. With all three algorithm design constraints elaborated, we can now describe the algorithm flow for the MaxRep choice function.

2.6 MaxRep choice function algorithm

Combining the three design constraints above, we will now describe the control flow for MaxRep's resolution of Condorcet's dilemma of coinciding cyclical majorities. The MaxRep choice function employs all of the pieces we've described in this paper: differential pairwise tallying, the voter specified moderation span, and α-parameterized tallying. Applying the design constraints described above leads to the following voter priority-based iterative bottom cycle set elimination algorithm. To both maximize the influence of each voter and to minimize dependence on decision context, MaxRep iteratively reduces the candidate field under consideration at the minimum value of α where the field separates into distinctly ranked subsets.

2.6.1 Minimal implementation

The flowchart in Fig. 10 and the associated description of the blocks in Table 1 below describe the simplest implementation of the MaxRep choice function. For conceptual simplicity, we will illustrate both versions of the algorithm using a straightforward linear search for min-α.

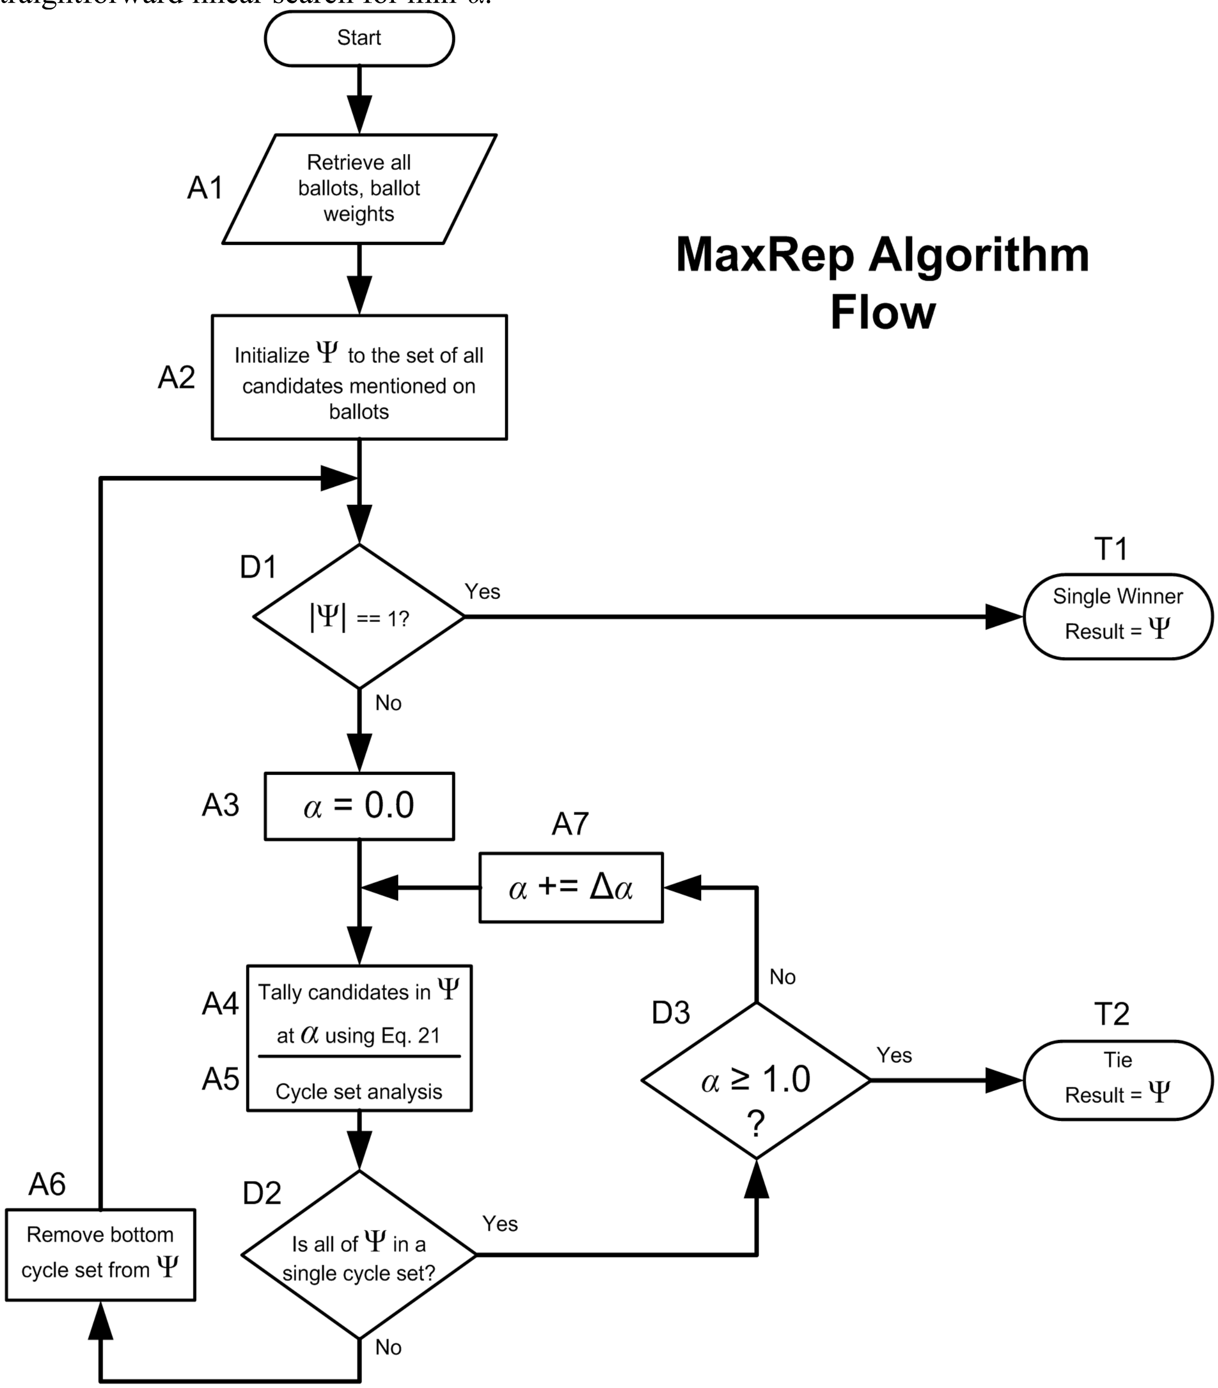

Fig. 10 The simplest implementation of the MaxRep elimination algorithm, with no computational shortcuts. The algorithm iteratively eliminates the bottom cycle set at min-α. The rectangular A blocks are actions, the diamond D blocks are decisions, and the oval T blocks are terminators

Table 1: MaxRep algorithm flow block descriptions.

A1	Input the real-valued preference ballot collection with the associated voter specified moderation spans and ballot weights. This collection will remain fixed over the execution of this algorithm.
A2	Initialize the set of candidates under consideration, Ψ, from which candidates will be eliminated. For an open nomination election, this would be the union of all candidates mentioned on any ballot.
A3	Initialize α to 0 in preparation for the search for min-α.
A4	Tally all ballots at the current value of α using Eq. 21 over the candidate set Ψ currently under consideration.
A5	Analyze the tally results and determine the bottom cycle set, if one exists.
A6	Eliminate the bottom cycle set from Ψ.
A7	Increase α to continue the search for min-α.
D1	Check if a single candidate remains in Ψ. This is the standard single-winner termination condition and will also exit in the case of a trivial, one candidate election.
D2	Check for whether Ψ separated at this value of α. If not, a higher value is necessary.
D3	Check for termination in search for min-α. This occurs whenever the candidates in Ψ are in perfect tie and no value of α on [0,1] will bring any sense of coherent ordering.
T1	Single winner termination, may be either a Condorcet winner or the candidate may have won outright after a round of elimination at elevated α.
T2	Multi-candidate perfect tie termination.

At the beginning of a decision, the MaxRep choice function is supplied with the ballot collection and associated ballot weights (A1). During the computation of a winner, the working ballot collection remains fixed and consistent. At this point, ballot integrity would be verified and ballot collections could be synchronized across all redundant tallying entities that are verifying this computation. After securing the ballot collection, the candidate set under consideration, Ψ, is initialized to the union of all options mentioned on any of the ballots (A2). Allowing open nomination in this manner satisfies Arrow's preference for an unrestricted domain. To handle this open candidate field, we also stipulate that all ballots include a specified default value that applies for any candidate for whom a preference value was not explicitly assigned.

Following initialization, the MaxRep algorithm iteratively eliminates the least preferable candidates from consideration. Before beginning the computation, the algorithm checks whether a single candidate remains in Ψ (D1). On the first pass, D1 will catch a trivial, single candidate election and exit. On each iterative elimination pass, the algorithm begins at $\alpha = 0$ (A3) and searches upwards to find min-α, following Design Constraint #1. Control flow then enters a search loop looking for the minimum value of α at which Ψ separates into at least two distinctly ordered cycle sets. For each value of α starting with 0, the algorithm tallies all the ballots using the α-parameterized moderated differential pairwise tallying defined in Eq. 21 (A4) over just the candidates still in Ψ.

The resulting square delta tally matrix is then analyzed by computing the Schwartz matrix and then performing cycle set analysis (A5), as described in Section 2.4.

Using the results of A5's cycle set analysis, D2 checks if any coherent partial ranking has emerged between the candidates in Ψ at this particular value of α. On the first pass, it is likely that some completely irrelevant candidates will separate out of Ψ at $\alpha = 0$. When separation occurs, A6 will eliminate the candidate(s) in the bottom cycle set from consideration, adhering to Design Constraint #2. Only the bottom cycle set candidates are removed from Ψ while the top and any middle cycle set candidates remain in consideration. After this conservative reduction of candidate field size, Design Constraint #3 requires that we reinitiate the search for min-α starting again from $\alpha = 0$ and consideration only the remaining contenders.

If no ordering has emerged at this value of α and all candidates are determined to be in a single cycle at D2, then a higher α value will be necessary to bring some coherent ranking to Ψ. Before incrementing α, D3 checks whether the entire interval between 0 and 1 has been searched. If no ordering has emerged even at $\alpha = 1$, the candidates in Ψ are in a perfect tie and the algorithm terminates in this tied condition at T2. If further values of α remain to be checked, A7 increments α by $\Delta\alpha$ and the algorithm then re-tallies all the ballots at this new value of α. This simplified algorithm employs a linear search in α, where $\Delta\alpha$ is defined prior to beginning the computation. Smaller values of $\Delta\alpha$, while computationally more demanding, produce a more gradual introduction of proportional decision perspective. Though not the topic of this paper, computationally efficient methods for finding min-α to high levels of precision have been shown to be feasible.

Figure 10 and Table 1 have defined a basic algorithm for the MaxRep choice function. We will now demonstrate how to improve the efficiency of this algorithm while still computing the same result. When Ψ partitions at $\alpha = 0$, each of the tally matrix elements from Eq. 21 has strict candidate pair dependence. Because all the delta tally matrix elements are independent of the greater candidate field context, removing the bottom cycle set will not affect the relative ranking between any of the other candidate sets. In fact, the top cycle set will remain invariant under the removal of any of the candidates below that set. This invariance of the top cycle set means that iterative bottom-cycle elimination at $\alpha = 0$ will proceed predictably until only the top cycle set remains. Therefore, when partial ranking separates the candidate field at $\alpha = 0$, we can take a significant computational shortcut and keep only the top cycle set in contention instead of eliminating just the bottom cycle set.

2.6.2 Optimized MaxRep Algorithm

The flowchart in Fig. 11 and annotation in Table 2 below describe a more efficient algorithm for the MaxRep choice function with the above mentioned significant computational shortcut. The prior candidate elimination loop has been partly unwound to allow the context independent results at $\alpha = 0$ to be more efficiently handled. When rounds of elevated α are necessary, this algorithm will proceed into the loop that evaluates these context dependent results using the same mechanics as before.

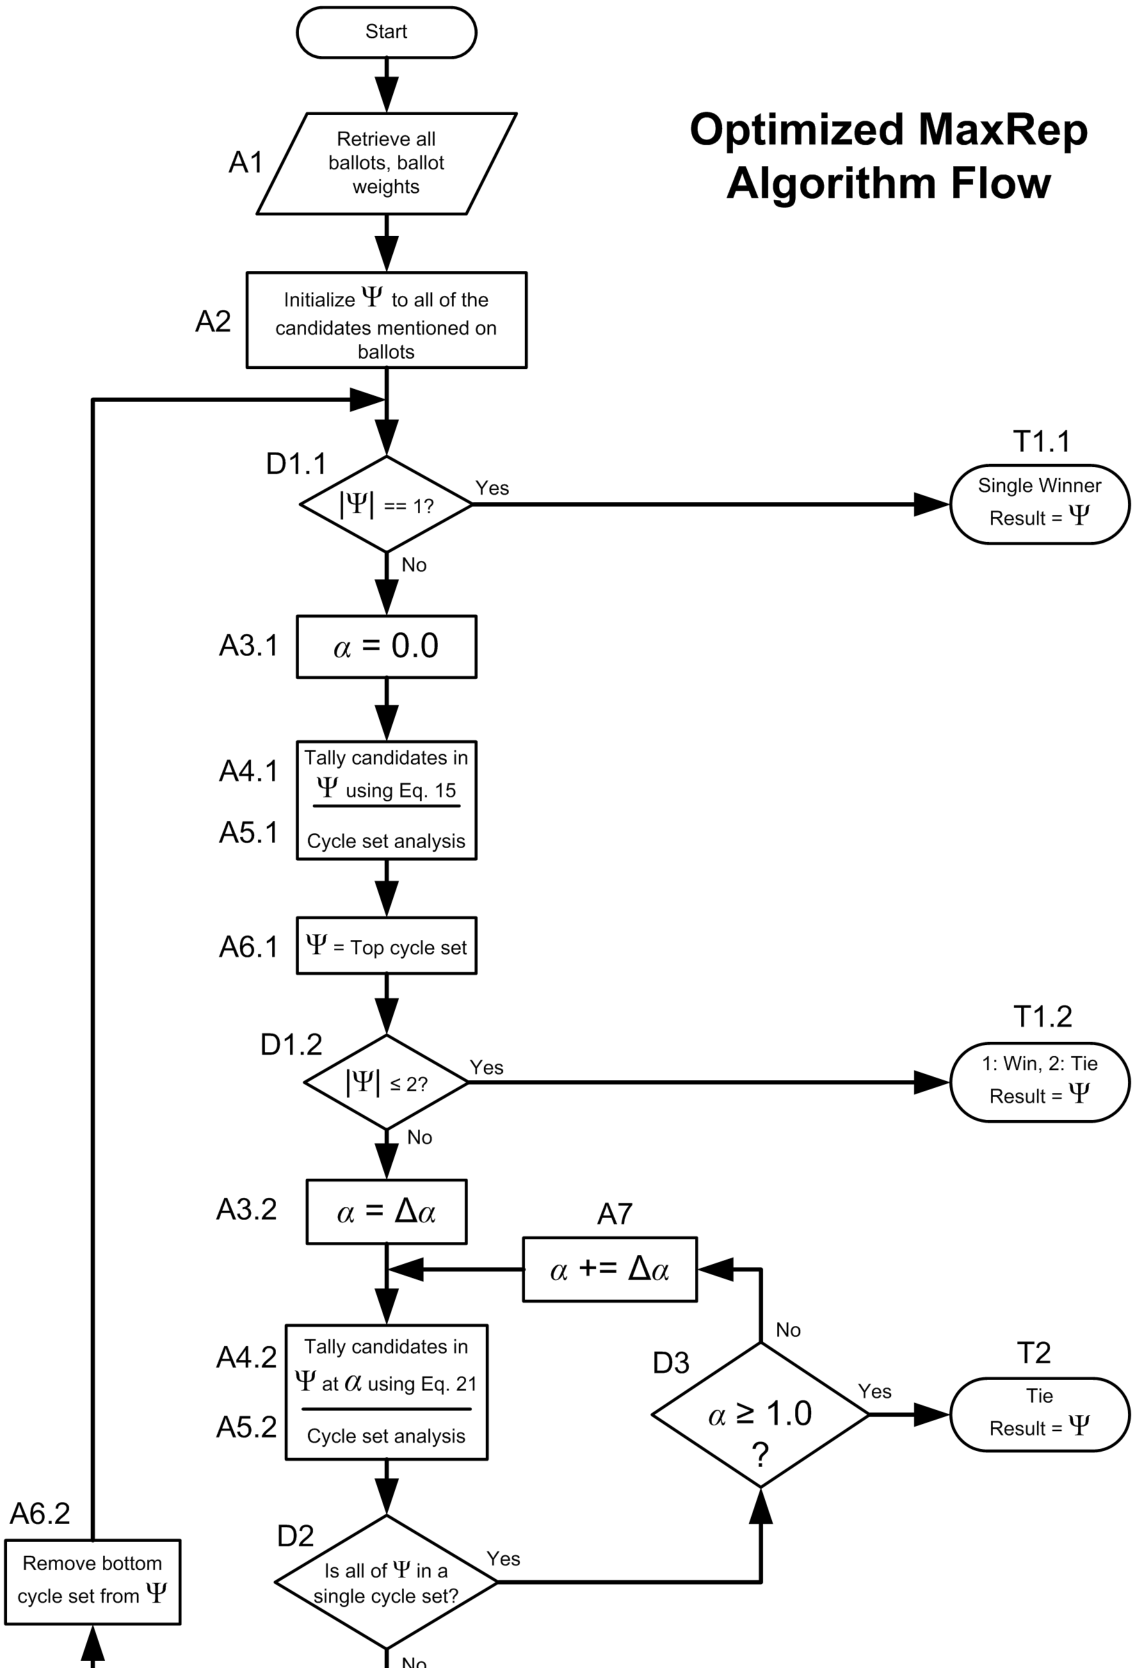

Fig. 11 The MaxRep choice function algorithm with an optimization for separation at α = 0

Table 2: Optimized MaxRep Algorithm Flow block descriptions.

A1	Consolidate the real-valued preference ballot collection with the associated voter specified moderation spans and ballot weights. This collection will remain fixed over the execution of the algorithm.
A2	Initialize the set of candidates under consideration, Ψ, from which candidates will be eliminated. For an open nomination election, this will be the union of all candidates mentioned on any ballot.
A3.1	Set α to 0 for initial, binary independent analysis.
A3.2	Set α to first step above 0 in preparation for search for min-α.
A4.1	Tally all ballots using Eq. 15 (since $\alpha = 0$), over all candidates in Ψ. The tally matrix is independent of candidate field context since the tally matrix elements have strict candidate pair dependence.
A4.2	Tally all ballots at the current value of α using Eq. 21 over the candidate set Ψ.
A5.1	Analyze the tally results and determine the top cycle set (which will include all candidates if there is a single cycle set). The top cycle set can be used because the strict candidate pair dependence in A4.1.
A5.2	Analyze the tally results and determine the bottom cycle set, if one exists.
A6.1	Set Ψ to only candidates in the top cycle set from A5.1.
A6.2	Eliminate the bottom cycle set from Ψ.
A7	Increase α to continue the search for min-α.
D1.1	Check if a single candidate remains in Ψ after elevated α. This will also exit in the case of a trivial, one candidate election.
D1.2	Check for single winner or a two candidate tie after moderated pairwise analysis. Two way cycles cannot exist in pairwise analysis, so if two remain they are in a true tie that will not be broken at elevated α. This will also find a Condorcet winner on first pass.
D2	Check for whether Ψ separated at this value of α. If not, a higher value of α is necessary.
D3	Check for termination in search for min-α. This occurs whenever the candidates in Ψ are in perfect tie and no value of α on [0,1] will bring any coherent ranking.
T1.1	Single winner termination after rounds of elevated α (or trivial one candidate election).
T1.2	Single winner or two-way tie after moderated pairwise analysis.
T2	More than two candidates in a perfect tie termination.

This optimized algorithm for the MaxRep choice function begins with the same initialization steps (A1, A2) as previously described. Blocks A3.1, A4.1, A5.1, A6.1 and D1.2 are the explicit unwinding of the calculations at $\alpha = 0$ from the initial algorithm description in Fig. 10. Since $\alpha = 0$, ballots can be tallied using the more computationally efficient moderated differential tallying in Eq. 15 (A4.1). The elements of the tally matrix produced by A4.1 will be independent of the greater candidate field context. Cycle set analysis of these pairwise results in A5.1 remains unchanged but A6.1 takes advantage of the tally's independence from candidate field context to remove all but the top cycle set candidates from consideration. For the initial pass of a large election, keeping only the top cycle in contention will most likely cause a dramatic decrease in the candidate field before the computationally expensive min-α passes. The termination condition in D1.2 is slightly modified. A cycle with only two candidates cannot exist in pairwise analysis since a directed edge can only point in one direction. As a consequence of this, a two candidate top cycle set after $\alpha = 0$ analysis indicates the candidates are truly tied.

Following these modified $\alpha = 0$ mechanics, A3.2 sets α to its first non-zero value. The linear search for min-α then proceeds exactly as described before. If multiple candidates remain after A6.2, this algorithm will again take advantage of independence from greater candidate field context when α is reset to 0. This optimized version of the algorithm computes exactly the same results as the basic version, but the computational efficiency improvement is generally substantial. This enhancement is especially significant during the first pass of the algorithm when there may be large numbers of fully irrelevant candidates who will all be removed during the initial $\alpha = 0$ evaluation.

2.6.3 MaxRep choice function summary

In summary, the MaxRep choice function we have just presented extends Part I's moderation span concept via a shared parameter of proportional perspective, α. Like the moderation span, α specifies a continuum between Condorcet ($\alpha = 0$, strict pairwise perspective) and Borda ($\alpha = 1$, fully linear perspective) style tallying methods. The level playing field defined by the α-parameterized moderated pairwise differential tallying method in Eq. 21 is the foundation of MaxRep's resolution of coinciding cyclical majorities. Raising the scope of decision perspective (α) will most equitably resolve any cycles by introducing pertinent information regarding each voter's relative priorities. The MaxRep algorithm for computing a group's top choice iteratively removes the least preferable candidates from consideration at the minimum possible value of α. By adhering to the three design constraints developed in Section 2.5, MaxRep maximizes the influence of each voter's preferences while also minimizing any dependence on less relevant alternatives.

2.6.4 Extending MaxRep to a full choice ranking

Since the MaxRep choice function employs an elimination algorithm, an extension of the basic choice function to produce a full candidate ranking is fairly straightforward. Candidates are eliminated in reverse preferability order, although several candidates can be eliminated together in a bottom cycle. Candidates dropped together can either be considered tied or, if strict ranking is required, α could be raised further in this candidate field context to determine relative ranking within just this bottom cycle set.

2.7 MaxRep algorithm results

To demonstrate the behavior of the MaxRep choice function, we will first present four hand-crafted illustrative cases. After these initial examples, we will discuss two 15 candidate decisions created from random ballots. All the results presented here were produced using the optimized MaxRep algorithm in 2.6.2 implemented in the Python programming language.

2.7.1 Hand-crafted examples

Example #1: Three candidate election with a Condorcet winner

Our first example of the MaxRep choice function in action is the set of ballots in Table 3 below which produce a Condorcet winner. In all of the first four examples with three voters and three candidates, we'll set each voter's moderation span to 0 for simplicity.

Table 3: Three ballots that produce a Condorcet winner.

Voter 1		Voter 2		Voter 3	
1.0	B	1.0	A	1.0	C
0.5	C	0.8	C	0.4	B
0.0	A	0.0	B	0.0	A

With these ballots no candidate receives a majority of first place votes, but candidate C receives some support from each voter. In head-to-head comparisons with each other candidate, C has the support of two of the three voters and is therefore the Condorcet winner. The MaxRep choice function selects C as the winner on the first iteration at $\alpha = 0$, as shown in Fig. 12.

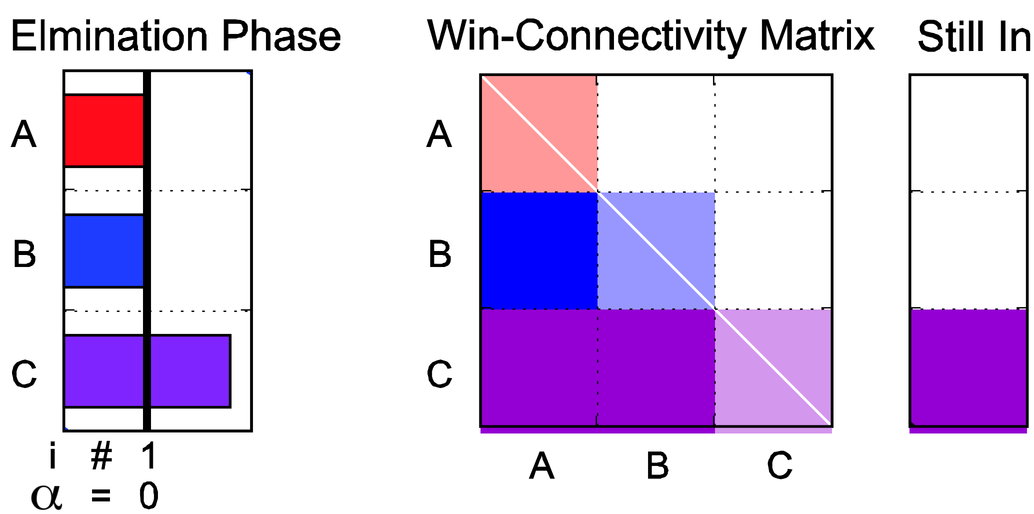

Fig. 12 The optimized algorithm for the MaxRep choice function picks C on the first iteration. On the left of the figure is the iteration-by-iteration survival of the candidates, with candidates along the vertical axis and the iterations (i) and min-α's along the horizontal. On the right is the connectivity matrix for the selected iteration, where direct victories are shown in full colors while the diagonal (where candidates tie with

themselves) are shown in half-tones. In this example, candidate C beats A and B head-to-head, and so is the only member of the top cycle, as shown by the "Still In" indicator at the far right

If this election was run using the non-optimized algorithm in 2.6.1, candidate A (the bottom cycle set) would be eliminated in the first round and then B (the middle cycle set) in a second round, still picking C as the overall winner. Regardless of which version of the algorithm is used, whenever a Condorcet winner exists in the initial moderated tally, the MaxRep choice function will select it.

Example #2: Three candidate cycle resolution

Finding a Condorcet winner (as in Example #1 above) is the simplest case and requires only a single iteration. For the second example, we will craft a top cycle set by switching Voter 1's middle candidate from C to A, as in Table 4 below.

Table 4: Three ballots that produce a pairwise cycle.

Voter 1		Voter 2		Voter 3	
1.0	B	1.0	A	1.0	C
0.5	A	0.8	C	0.4	B
0.0	C	0.0	B	0.0	A

This change produces an ambiguous cyclical edge graph at $\alpha = 0$, as each candidate has one win and one loss, as shown in Fig. 13. The placement of the middle candidate on each ballot is then vital in resolving an overall winner from this cycle.

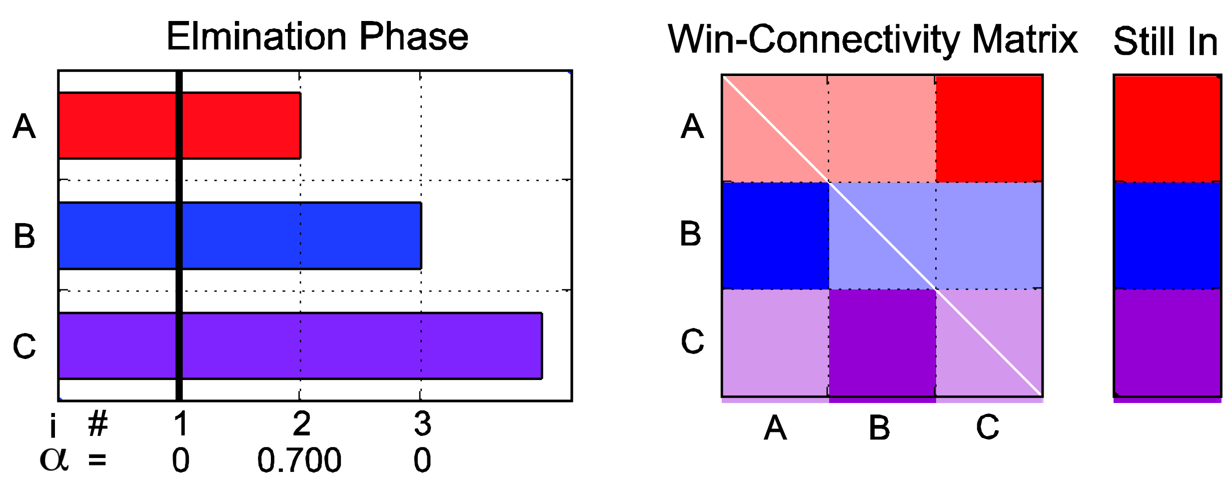

Fig. 13 The three candidate cycle does not resolve on the initial $\alpha = 0$ pass. For the first iteration, the connectivity matrix shows all candidates are connected in a cycle, with direct wins in full colors and multi-hop connections in half-tones

After the initial $\alpha = 0$ pass, all three candidates are entangled in a single cycle. The algorithm then proceeds to introduce information about each voter's relative priorities to resolve this cycle. Glancing at the ballots in Table 4, we might first presume that candidate B would be eliminated first since the sum of B's row is smallest. If α elevated all the way to 1 and performed a Borda-style contest, this is indeed what would occur.

However, the MaxRep algorithm finds a different consensus at a lower value of α, as show in Fig. 14.

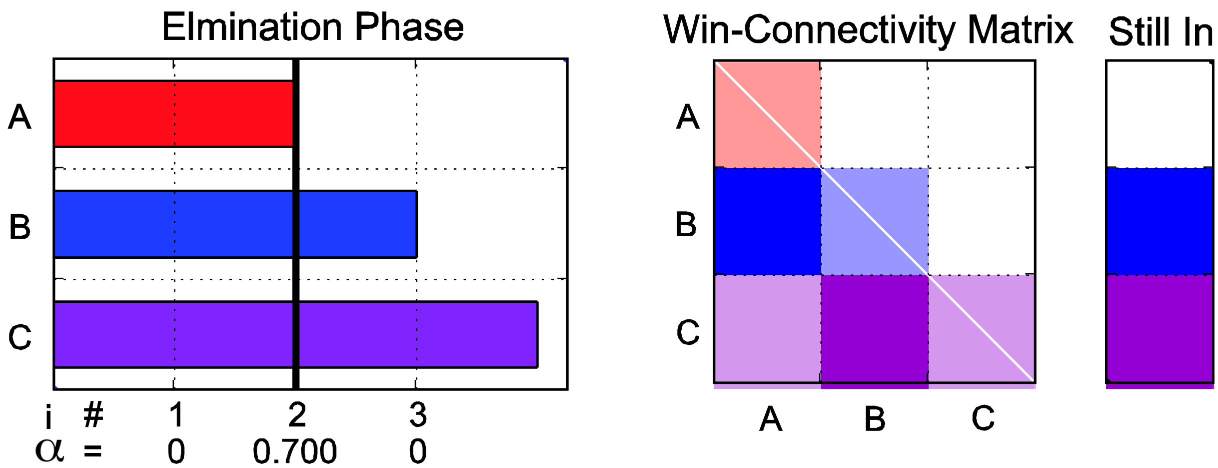

Fig. 14 At α = 0.700 on the second iteration, the algorithm determines that candidate A is the least preferable. In particular, A's pairwise victory over C disappears because the difference between these candidates is small on ballots 1 and 2. As α increases, these small differences are moderated first while Voter 3's insistence that C is better than A does not. At α = 0.700, these effects balance exactly, producing a tie between C and A, causing A to drop from the cycle. The cycle separates with C on top, B in the middle, and A at the bottom. Since this occurs at elevated α, only the bottom candidate is removed on this iteration. In iteration 3, C wins over B head-to-head as can be seen in the connectivity matrix of Fig. 13, so C is chosen as the overall winner

To clarify why the MaxRep choice function eliminated candidate A on the second iteration, we will examine each voter's differential support for the various candidates. The delta preference matrix form of a voter's ballot was first introduced in section 1.5, and the matrices for this example are shown in Fig. 15 below.

	Voter 1			Voter 2			Voter 3		
	A	B	C	A	B	C	A	B	C
A	0	-0.5	0.5	0	1.0	0.2	0	-0.4	-1.0
B	0.5	0	1.0	-1.0	0	-0.8	0.4	0	-0.6
C	-0.5	-1.0	0	-0.2	0.8	0	1.0	0.6	0

Fig. 15: The delta preference matrices for each ballot from Table 4.

For pairwise sub-contest AC, Voter 1 and Voter 2 support A by 0.5 and 0.2 respectively. Voter 3 supports C over A by the full span of their ballot, 1.0. Since all moderation spans are 0, at α = 0 the tally is two votes for A and 1 voter for C. At α = 0.700, Voter 1's and Voter 2's differential support for A are both moderated and since ballot spans are 1.0, they're simply divided by α. The sum of 0.5/0.7 and 0.2/0.7 is one, so at α = 0.700 the tally is now a net 1 vote for A and 1 vote for C, producing a tie between these candidates. Since this victory over C was A's only victory, A is now in the bottom cycle set and is removed from contention. If α had continued to climb, the edge between A and C would have completed its transition and C would beat A at these higher values of α.

The edge transition we might have initially expected, where B would switch to losing to A, does not occur until $\alpha = 0.900$. At this higher value of α, the differential support of A over B by voters 1 and 2 by 0.5 and 0.4, respectively, balances Voter 3's support of B by their full ballot span. When votes are expressed in delta preference matrix form, it is easier to see that on net the voters considered B beating A a higher priority that A beating C. The additional benefits of making elimination decisions at lower values of α were also discussed in Section 2.5.1.

Example #3: Three candidate perfect tie

Although it would be a very rare situation in real-world voting, it is possible for a set of candidates to be in a perfect cyclical tie. When all three submitted real-valued ballots are exact rotations over three candidates, there is no net relative priority information to be found. To produce this from the cycle above, change all the middle preference values to the same fraction of ballot span, as shown in Table 5.

Table 5: Three ballots that produce a perfect tie.

Voter 1		Voter 2		Voter 3	
1.0	B	1.0	A	1.0	C
0.5	A	0.5	C	0.5	B
0.0	C	0.0	B	0.0	A

Since all moderation spans are 0 for these initial cases, at $\alpha = 0$ the three candidates will emerge in a cycle of the same topology as the previous example, as shown in Fig. 16.

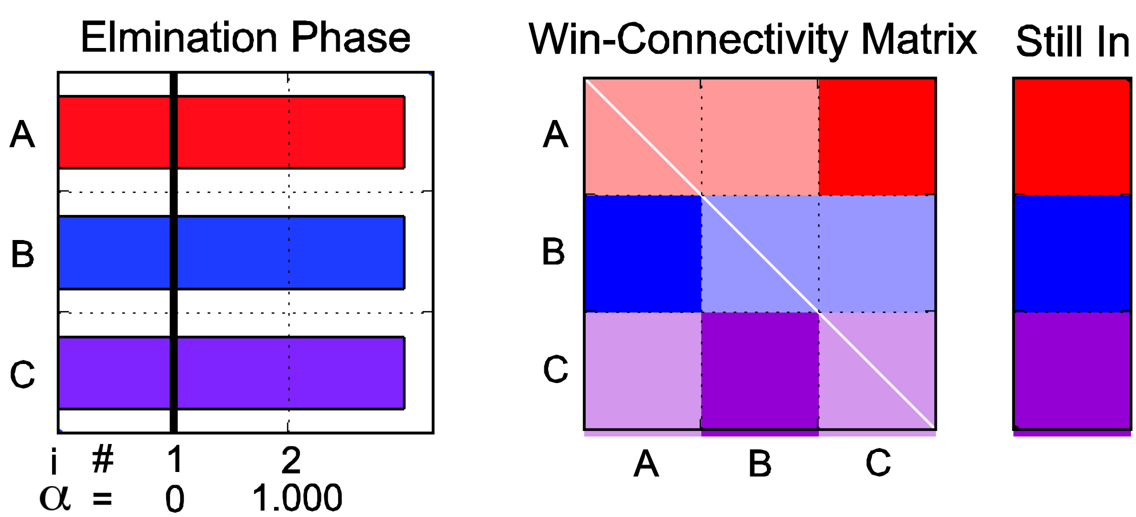

Fig. 16 On the initial $\alpha = 0$ pass, all candidates are in the top cycle set as shown by the connectivity matrix

Unlike the previous example, in this case introducing additional relative priority information from each voter cannot break the perfect tie because of the perfect symmetry of this set of ballots. The candidates remain in a cycle all the way through $\alpha = 1.0$, at which point all the edges simultaneously transition from wins to ties as shown in Fig. 17.

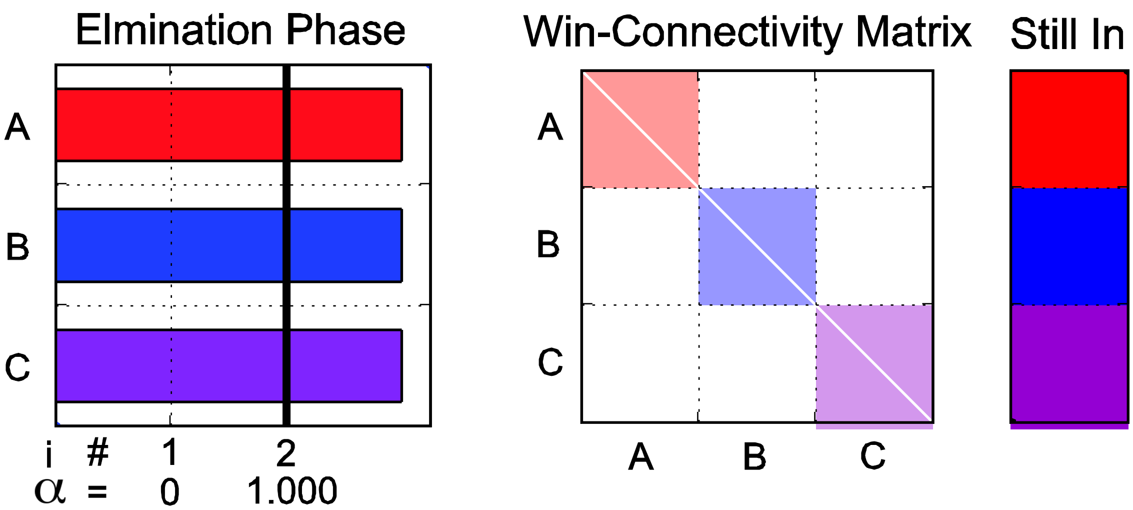

Fig. 17 At $\alpha = 1.0$, all ballots go fully linear and the results are equivalent to a real-valued Borda tally which essentially sums each row in Table 5. Since each candidate receives a 1.5, the three are perfectly tied. The connectivity matrix in this case shows no edges (victories) between any candidates, so all are in the top cycle set

For these exactly balanced ballots, the algorithm cannot find a single winner and all three are determined to be in a tie. In essence, this is also what would occur if the voters in Example 2 were required to vote with ordinal ballots.

Example #4: Three candidate cycle above lower candidates

For this final three voter example we will add several "irrelevant" alternatives to the cycle creating ballots also used in Example #2. To make room for these new candidates, we'll expand the voters' ballot spans but keep the same delta preferences between pairs of candidates A, B, and C, as shown in Table 6.

Table 6: Three ballots from Table 4 with extra candidates

Voter 1		Voter 2		Voter 3	
2.0	B	2.0	A	2.0	C
1.5	A	1.8	C	1.7	D
1.0	C	1.0	B	1.4	B
0.8	D	0.5	G	1.0	A
0.5	E	0.3	D	0.5	F
0.2	F	0.2	E	0.2	G
0.0	G	0.0	F	0.0	E

Each of the added candidates D through G receives the support of at most one voter when compared with any of the original three candidates. Because of this, all four of these newly introduced candidates will be removed in the initial $\alpha = 0$ pass, leaving the original cycle to be resolved, as shown in Fig. 18.

Fig. 18 On the initial $\alpha = 0$ pass, only candidates A, B and C are in the top cycle set, while D is in a middle cycle and E, F, and G are the bottom cycle. Once the irrelevant alternatives have been removed, iterations 2 and 3 proceed exactly as in Figs. 13 and 14

After the irrelevant candidates are removed in the initial $\alpha = 0$ pass, subsequent iterations proceed exactly as in Example 2. After D, E, F, and G are removed, the ballot spans for all three voters are back to 1.0. In addition, each voter's differential preferences over the set {A,B,C} have not changed, so their delta preference matrices will be identical to Fig. 15. The remaining elevated α calculations follow exactly the same logic as previously described. Even after the removal of A on the second iteration, the third iteration proceeds as if {A,D,E,F,G} had never been considered.

This independence has highlighted an important consequence of Design Constraints 2 and 3 which we term *contest regularity*: if an election was started with any of the intermediate values for the candidate set Ψ, subsequent elimination will proceed in an identical manner, yielding the same outcome as the full elimination sequence.

2.7.2 Examples with random ballots

In the following examples, between 3 and 11 random ballots were used with each mentioning 15 candidates. Each ballot also had a random moderation span that could be up to 10% of their ballot span. We chose a small number of random ballots since any sense of consensus amongst these synthetic ballots statistically diminishes into a randomly distributed noise floor as the number of ballots grows larger. When elevated values of α are necessary to resolve cyclical majorities, these examples compute min-α to an accuracy of 10^{-3}.

Example #5: Large number of candidates with random ballots

For this first example with random ballots, we selected a decision that required a single iteration of elevated α. When displaying these large candidate fields, we display the candidates in a particular order to highlight structure in the connectivity matrix. Two criteria are used, the first being the order of removal from contention. When a block of candidates are dropped on the same iteration they are ordered from least to most pairwise wins in the initial $\alpha = 0$ pass. The added clarity from this display ordering can be seen in the connectivity matrix in Fig. 19.

.

Fig. 19 Example election using random ballots and 15 candidates. Half of the candidate field is eliminated on the first $\alpha = 0$ iteration, including a 4 candidate middle cycle. The win-connectivity matrix at right shows the connections between candidates on this first iteration. At $\alpha = 0.409$, candidate G separates from the other 7 remaining candidates. G is the only candidate who beat N in the initial $\alpha = 0$ connectivity matrix at right. Once G is eliminated and α reset to 0, N is then chosen as the winner

One interesting feature of the connectivity matrix in Fig. 19 is the number of ties. Ties can be located by examining elements of the connectivity matrix that are reflections across the diagonal. When neither element is drawn with a darker, full color tone, the candidates tied head-to-head. In this case, L ties with B as neither cell has a full tone, but since L beats M who beats B there is a connection from B to L as indicated by the half-tone in the ultimate connectivity matrix.

Example #6: Multiple rounds of elevated α

The final example we will present here contains multiple iterations of elevated α, as shown in Fig. 20.

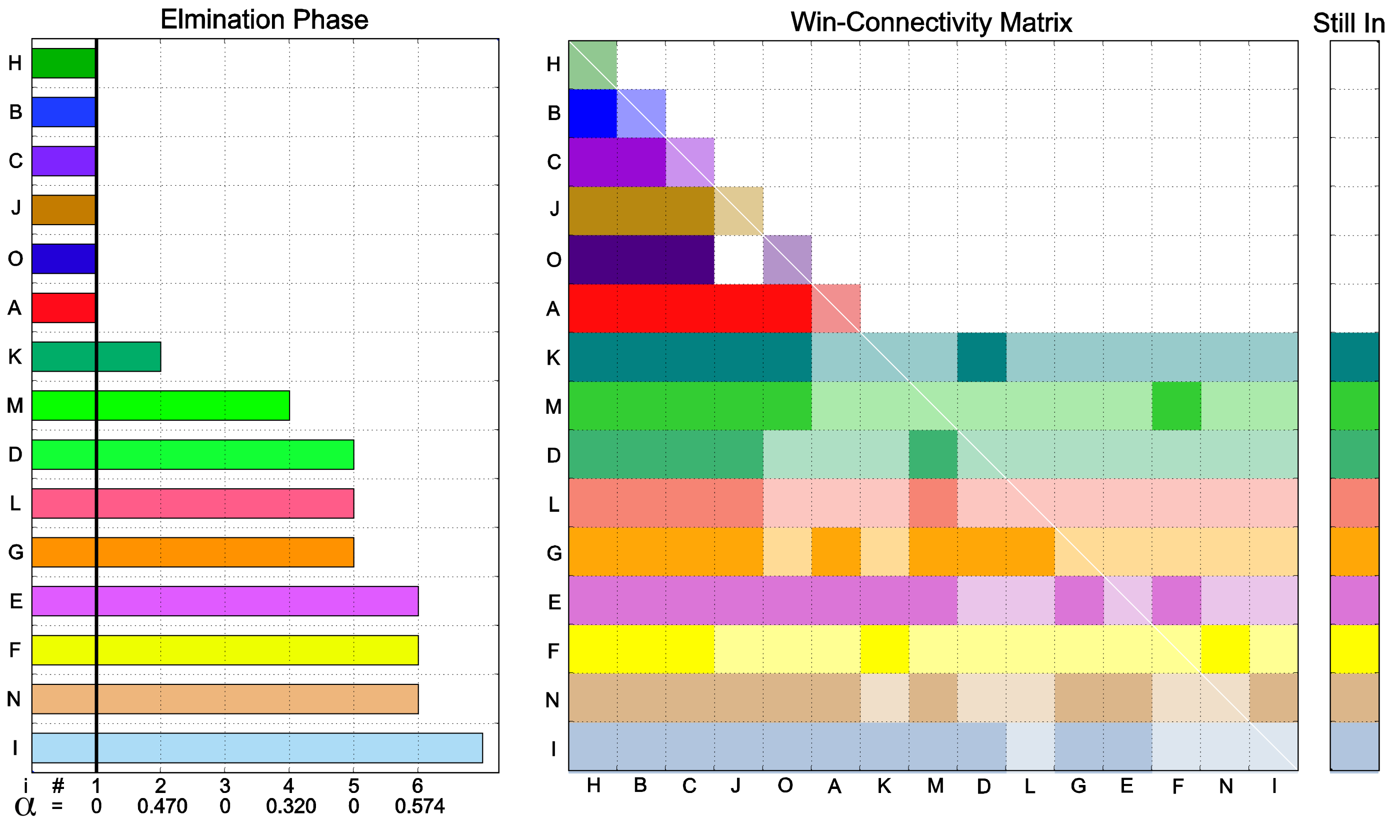

Fig. 20 An example decision between 15 candidates chosen by random ballots that requires multiple rounds of elevated α to resolve. On the initial $\alpha = 0$ pass half of the candidates drop based on the win-connectivity matrix on the right. After iteration 2 where candidate K drops at $\alpha = 0.470$, no candidates are removed at $\alpha = 0$. Another iteration at elevated α is required to determine that M is the least preferable candidate remaining. In this case, M was the key to holding D, L, and G in the running, as can be seen in the initial connectivity matrix at right. M is the only candidate of the four to have a victory over one of the bottom four candidates. Finally, on iteration 6 at $\alpha = 0.574$, candidate I emerges above the cycle set {E,F,N}. These bottom cycle candidates are all removed at once, leaving I as the winner

Notice that a lower value of α brings ordering on iteration 4 than is required on iteration 2. After candidate K is removed, some voter ballot spans must have changed. As explained in 2.5.3, when the candidate field context of a decision changes so may the minimum value of α required to separate the candidates.

2.8 Discussion

2.8.1 Properties of the MaxRep Choice Function

On the first pass of the optimized MaxRep algorithm in 2.6.2, only the top cycle set or Schwartz set candidates are kept in consideration for subsequent analysis. As a result, the eventual winner will always be a member of this initial top cycle set, a property we termed *moderated Schwartz definite* and first introduced in Section 2.4. This property is an extension of the standard Schwartz criterion since MaxRep uses moderated pairwise tallying. As a consequence of this, MaxRep also satisfies the *moderated Condorcet winner criterion* from Section 1.11 since it will always select a Condorcet winner when one exists.

When cycle resolution is necessary, both implementations of the MaxRep choice function presented follow the three design constraints from 2.5. As a consequence of these three design constraints, the choice function also has the *contest regularity* property discussed at the end of Section 2.7.1: if a new contest was started using the candidate field considered in any iteration of a prior run of MaxRep, the new contest will always choose the same candidate as the winner.

2.8.2 Monotonicity vs. Maximal Influence

On each iteration of MaxRep's elimination algorithm, the choice function makes progress towards a decision while maximizing the expression of each voter's preferences. While searching for min-α candidate field separation, ballot scale factors remain constant and the width of the linear sigmoid monotonically increases. Each round of elimination, therefore, provides positive association of voter preference and the resulting candidate standing: if you raise candidate A on a ballot and rerun a given round of elimination, this cannot decrease A's chance of remaining. Monotonic results are guaranteed on each iteration so long as the candidate field considered for the iteration does not change. However, any change in the elimination order of the least-preferable candidate(s) at elevated α causes a shift in context for future iterations. There are two potential effects of this perturbation in the candidate field considered which can lead to non-monotonic results: change in ballot scale factors and change in edge-graph topology.

To illustrate these two potential effects, consider the following scenario: a decision has been computed which involved at least one round of elevated α. One voter in this decision opts to move candidate A up their ballot, while keeping all other candidates at their original values. This move need not cause A to pass any other candidates, at elevated α just changing the delta preferences between candidates is sufficient. With this single change the results are then re-computed. This new decision also requires cycle resolution and on an elevated α iteration of the algorithm the change in A's position on the ballot causes the cycle to resolve in a different manner than before, resulting in a different set of candidates being considered on subsequent iterations.

One effect of eliminating different candidates is that one or more voters' ballot spans may change for subsequent evaluation. If further rounds of elevated α are required, then these voters' tally contributions will change. These changes can, in rare circumstances, cause candidate A to be removed earlier than they would have been without the modification to the original voter's ballot. That said, MaxRep's conservative bottom-cycle elimination results in the smallest possible change to the candidate field between iterations, minimizing the changes in ballot scale factors between iterations. In addition, making elimination decisions at the minimum possible value of α also minimizes the effect of a change in context after elimination. While further analytic treatment seems necessary, minimizing both of these factors appears to minimize the probability of non-monotonic behavior resulting from the change in ballot scale associated with influence maximization. In addition, when voters choose to moderate not only are cycles rare but any change in ballot scale factor between iterations is reduced and non-monotonic behavior is less prevalent.

Another effect of removing a different candidate(s) is that the topology of the edge-graph will be different in subsequent iterations. By changing which candidates remain in consideration, the resolution of subsequent iterations should also change. It is possible that candidate A will not fare as well in these subsequent iterations because of the altered structure of the edge graph. Since the results of one iteration must depend only on the voters' preferences between the candidates considered in that iteration, non-monotonic effects from changing edge graph topology are inescapable. However, we conjecture that by eliminating the smallest number of candidates possible and by incorporating voter priority information at lowest value of α, MaxRep should reduce the potential for non-monotonic results.

It is worth noting that there is a simple modification to MaxRep's elimination method that can guarantee positive association of preference values and final ranking: instead of raising α until a bottom cycle separates, raise α until a single winner emerges. Unfortunately, this switch sacrifices consistently maximizing the expression of each voter's preferences on round of elimination and thereby expands the potential gain from strategic voting. In effect, not eliminating candidates at min-α and then adjusting ballot scale factors is akin to unnecessary and involuntary moderation. Uniform maximization of voter influence by minimizing α remains the higher priority, as discussed in 2.5.

The relaxation of the hard constraint of positive association of preference values with the final ranking is necessary to achieve the more important goal of uniform maximization of preference expression. This result is a sobering effect of balance between influence maximization and the desire for monotonicity. However, MaxRep still achieves positive association of values on an iteration-by-iteration basis. Voluntary moderation also gives voters the ability to reduce the prevalence of cycles (as shown in Section 1.10) and any of this chance of non-monotonic behavior.

2.8.3 MaxRep's Relationship to Arrow's constraints

In 1.2 and 1.3, we examined two methods which nearly achieved all five of Arrow's original desired properties. However, as Arrow's impossibility theorem shows, when a social choice method must always yield a result and limit voter authority, no choice function can achieve all five properties. When a constrained design problem definition returns a null solution set, a reassessment of the original constraints is necessary. With this perspective, we now outline the set of prioritized goals that have been incorporated into the design of the MaxRep choice function. We believe that this prioritized set of objectives, when realized on a foundation of neutrality, yields a well-formed framework for the solution of the generalized social choice problem.

First, we clarify four hard constraints which define a foundation for a neutral choice function: *anonymity*, *neutrality*, *non-imposition*, and *unrestricted domain*. Anonymity and neutrality, as defined in Schulze 2003, insist that the choice function be unbiased with respect to both voters and candidates. Non-imposition and unrestricted domain are two of Arrow's properties which require a choice function to allow the submission of all possible preference ballots and consistently return a deterministic result for any set of input ballots.

Beyond this foundation of neutrality and consistency, the MaxRep choice function uses the following prioritized list of criteria to define how to best make a choice:

1. resolvability
2. voluntary moderation
3. uniform maximization of preference expression
4. maximized probability of positive association of values

Resolvability is the property that the choice function must consistently return a result: as the number of voters grows, the chances that there is more than one potential winner must diminish towards zero. With a high number of voters, the chance of a true tie between candidates drops to zero because it requires a perfect balance in ballots as in Example #3. Ambiguous cyclical majorities, however, increase as the number of voters increases (as discussed in Section 1.10) and so a viable choice function must resolve these cycles.

Voluntary moderation signifies that each voter has primary control over their ballot's balance between influence maximization and fidelity of preference priority expression. The voter specified moderation span introduced in 1.8 properly formulates this concept. Moderation control effectively provides voters with an additional degree of freedom when casting their ballot, but because it is voluntary does not encourage strategic voting. This property expands on the idea of non-imposition by allowing voters to choose how they would like their preferences to be expressed. The choice function can require a higher level of moderation when necessary to resolve a decision.

Uniform maximization of preference expression means that the expression of each voter's preference priorities outside of their span of moderation will be expressed with the voter's full weight. This criterion, as the third priority in the list, is subject to both the constraint of resolvability (priority #1) and allowing voters to choose a level moderation higher than that needed to find a solution (priority #2). Subject to an allowance for voluntary moderation, the desirability of uniform maximization of voter preference expression cannot be understated. When a choice function maximizes influence automatically, voters do not need to strategically distort their preference ballot in an attempt to increase their influence with respect to any top contending candidates. The MaxRep choice function was designed to maximize the influence of every voter's ballot and, when necessary to resolve cyclical ambiguity, find compromise while keeping all voters on a level playing field. This new optimality criterion of influence maximization profoundly improves Arrow's constraint of non-dictatorship, instead stating that each voter should be given as much of a voice in the decision possible.

As we have noted previously in Section 2.5.1, maximizing voter influence by minimizing α also yields a result with maximum independence from less relevant alternatives. When a Condorcet winner exists, the result of the MaxRep choice function is completely independent from less relevant alternatives. It is only when the electorate's consensus is clouded by distortion from a lack of moderation that the additional relative priority information introduced by raising α is needed to resolve cyclical ambiguity. The MaxRep choice function exhibits full independence from irrelevant alternatives when possible, otherwise yielding a result with maximal independence from less relevant alternatives.

Maximized probability of positive association of values means that, subject to the top three priorities above, a choice method must be designed to produce a direct relationship between changes in candidate standing on ballots and the final choice. A consequence of MaxRep's success in the consistent maximization of voter influence is that positive association of preference values and final ranking (Arrow's second desirable property, also known as monotonicity) cannot be guaranteed when cycles exist. When there are no cycles, the MaxRep algorithm terminates in a single iteration and end-to-end positive association is assured. However, changing the elimination order of candidates between iterations will alter the candidate field considered in subsequent iterations and thereby introduces some chance for non-monotonic behavior. The tradeoff between influence maximization and monotonicity in MaxRep is discussed at length in the next section.

2.8.4 Game aspects of the MaxRep choice function

The MaxRep choice function was designed to maximize the authority of each voter's priorities in the context of relevant candidates while also providing a level playing field for all voters on which cycles can be equitably resolved. This maximum influence for non-strategic voters reduces the impetus for strategically distorting preference ballots, particularly in the absence of precise information on the way in which any cycles will resolve. As the Gibbard-Satterthwaite theorem states, voting systems can never be completely strategy proof (Gibbard 1973, Saatherthwaite 1975). However, by consistently maximizing the influence of each voter's priorities, the MaxRep choice function reduces the motivation and potential benefit of strategic voting. If voters know the vote tallying method will maximally represent their preference schedule, they are then motivated to express their true preferences. We assert that only when voters submit their true preferences can a social choice function assess a group's genuine consensus.

Nonetheless, when rounds of elevated α tallying are necessary to resolve cycles, there may exist new voter strategies that could take advantage of MaxRep's adaptivity to the candidate field. Among the possible new strategies is evenly spacing the top contending candidates on the ballot so that the voter begins to moderate after other voters with more closely spaced candidates. While this may decrease the chances that a voter has to compromise, it also hides information about which delta preference difference are less important to them. By not compromising over differences they care less about, the voter risks losing a more important difference since they have not given the algorithm any indication of this relative importance.

Another potential precarious strategy is to place candidates ordered on a ballot with exponential spacing. This strategy attempts to take advantage of MaxRep's ballot span adaptivity such that the voter's favorite remaining candidate always receives unmoderated support until α is very close to 1. Even at small values of α, this strategy sacrifices information regarding the standing of all of a voter's other priorities in quest of their most preferred candidate. By only supporting their favorite candidate at elevated α, the voter may cause their second or third choices to lose unnecessarily. If their top candidate is eventually eliminated, they would then have lost the opportunity to select their second choice candidate. This potential loss of influence over the rest of the candidate field makes exponential candidate spacing a risky strategy.

There remains a clear need for further analytic treatment on the susceptibility of MaxRep to voting strategy, particularly under circumstances of variations in available information. However, the MaxRep choice function provides non-strategic voters with maximized fair influence over their expressed priorities so the impetus for strategic ballot distortion has been mitigated. We conjecture that without highly reliable, detailed information on the voting priorities of others the risks of strategically misrepresented preferences will most often outweigh any potential gain.

2.8.5 The MaxRep Choice Function as Strategically Negotiated Compromise

In a generalized negotiation, at any point in the negotiation there are three categories of items under consideration:

- What you can clearly get
- What you realize you will not be able to get
- What will require compromise in further negotiation

One difficulty often encountered in negotiation is effectively and efficiently discriminating which items fall into which of these categories. The MaxRep choice function follows this same pattern. During the initial $\alpha = 0$ pass, MaxRep determines what items should remain under consideration and what should be eliminated. In the Condorcet winner case, there is a conclusive decision between the parties and no compromise is necessary. When cycles exist, however, some compromise will be required to resolve them and make progress towards a final decision. When there are more than two alternatives remaining in consideration, the additional degree(s) of freedom in negotiation allow the necessary room for this compromise.

The equivalent categories to those above using the MaxRep perspective are:

- Options that are low on your preference list that will be eliminated
- Options that are high on your preference list that will be eliminated due to the preferences of others

- Options where your expression of relative priorities and those of the others will determine the most equitable compromise (the top cycle items)

The initial $\alpha = 0$ iteration finds the top cycle which contains the items over which further negotiation is necessary to find the best compromise. After an item is eliminated from consideration at elevated α, resetting α to zero is akin to each party strategically re-maximizing their priorities over what remains on the table. In the MaxRep framework, compromise is transitory and is limited to the decision context of the current iteration. Negotiation then proceeds iteratively, starting by reevaluating the items still under consideration.

We suggest that MaxRep is essentially a practical definition of strategically negotiated compromise in the absence of any a priori information regarding contest outcome. Groups choosing to make decisions using MaxRep will need to assess whether the level playing field assertion from Section 2.3.2 is a fair framework for resolving cycles through strategically negotiated compromise.

2.8.6 Reflections on the possibilities of moderation

As a social decision system, the MaxRep choice function still allows for the idealized linear voting scenario which we presented in Section 1.3. Even after designing MaxRep on the premise that each voter may try to maximize their own influence, the voter specified moderation span still leaves open the possibility that whole communities will be able to see their individual preferences from a broader perspective. While lack of uniform influence maximization encourages strategic voting, voluntary moderation can improve the fidelity of the expression of each voter's preference priorities. When most voters choose to vote moderately the group can make decisions with more of a Borda-like, shared benefit-cost perspective. Some voters may even recognize they do not have as much at stake in a particular decision as others and set their moderation spans larger than the full span of their ballot. Although the reality of contentious and consequential elections requires our algorithm be hardened to strategic voting, MaxRep's voluntary moderation mechanisms do no preclude its use by more moderate groups of people.

2.9 Comparisons with other choice functions

IRV and STV:

Instant Runoff Voting and Single Transferable Vote both involve an iterative elimination of the candidate with the fewest first place votes. This elimination has some similarity to MaxRep's bottom cycle set elimination. However, the MaxRep method has a substantially improved technical foundation of making an elimination decision based on the priorities of the voters' whole preference ballot instead of just the top position. This distinction mitigates the spoiler effect that plagues the plurality-based elimination method in the implementations of both IRV and STV. Without this paper's improvements, both methods still encourage voters to speculate which candidates are top contenders and distort their preferences accordingly.

Range Voting and other Borda-like methods:

Range Voting has the positive aspect of expressing a voter's relative preference information with full fidelity. However, because the ballots are span-normalized as in Eq. 5, Borda-like methods have no sense of independence from less relevant alternatives. The unnecessary reduction of voter influence due to less relevant candidates encourages speculative voting strategies, such as that portrayed in Fig. 2. In contrast, MaxRep maximizes voter influence, only reducing voter influence and introducing information regarding lesser priorities when necessary to resolve cyclical ambiguity.

Condorcet's tallying method with a cycle breaking scheme:

Classic pairwise analysis suffers from an issue opposite that of Borda methods. Pairwise analysis, while independent from less-relevant alternatives, is blind to relative preference magnitudes, as observed in 1.2. The distortion of small differential preference to full voting weight in quest of consistent influence maximization can be seen as the root cause of Condorcet's dilemma of cyclical ambiguity. Cycle breaking schemes such as Beatpath (Schulze) and Ranked Pairs (Tideman) use the relative weakness of some pairwise victories without voter priority weighting to ignore the edges that produce a cycle. MaxRep reintroduces the voter priority information which was inherently lost in pairwise analysis to resolve cycles; breaking cycles based strictly on properties of the pairwise tally is a misplaced attempt to find some substitute for voter priority information.

Binary elimination tournaments:

In light of Condorcet's cyclical outcome dilemma, decision structures like two-at-a-time trees or tournament-style brackets are indeterminate: that is, based on the ordering of comparisons, a different ultimate winner will emerge for the same set of ballots. While it is not practical for sports, a direct side-by-side comparison of all alternatives can allow for a consistent, well-formed decision.

2.10 Conclusion

2.10.1 Summary

This paper formulates the MaxRep choice function, a vote tallying system for group decision making which is an adaptive hybrid of the methods of Condorcet and Borda. Part I presents moderated differential pairwise tallying as the beginning of this hybridization. The introduction of the voluntary moderation span gives each voter the freedom to specify an initial level of proportional priority perspective for their real-valued preference ballot. When voters moderate over options they find similarly preferable, this augmentation to Condorcet's pairwise analysis is shown to reduce the probability of an ambiguous cyclical outcome.

Part II extends Part I's moderation span concept with the tally-wide decision perspective parameter α which similarly specifies a continuum between the Condorcet ($\alpha = 0$) and Borda ($\alpha = 1$) tallying methods. This α-parameterized pairwise differential tallying formulation is the foundation of our MaxRep algorithm. When ambiguous cyclical outcomes occur, increasing parameter α introduces additional pertinent voter priority information which provides for the equitable resolution of these cycles. The requirement of some dependence on candidate field context in the formulation of α necessitates that cycles be resolved conservatively. We show that iterative elimination of the least relevant (bottom cycle set) candidate(s) at min-α maximizes the expression of each voter's preferences while also minimizing any dependence on less relevant alternatives. We assert that this adaptive proportional perspective of voter priorities as a basis for cycle resolution, from a priority information transmission perspective, forms the best resolution of Condorcet's dilemma.

With the introduction of the MaxRep choice function, there may now be a more viable process for making congenial group decisions over larger alternative spaces. Given that expressed preferences are well informed, the authors conjecture that the quality of decisions made using this paper's MaxRep choice function should only improve with the number of options considered.

2.10.2 A New Voting Experience

Be it for voting in small group decisions or large elections, the placement of multiple alternatives on a real-valued preference number line is an intuitive expression of relative ranking. With these developments, voters are able to do more than just select one candidate and hope that their vote has some desired impact. Real-valued preference ballots enable voters to express their relative preferences across all alternatives. The MaxRep choice function consistently represents each voter's ballot in the decision process, irrespective of which alternatives are actually in top contention. In contrast, with other voting systems voters have to guess when casting their ballot which alternatives will be amongst the top contenders to exert some desired influence. Systems without maximum independence from less-relevant alternatives (even those employing preference ballots) encourage voting strategies based on tenuous speculation. With the MaxRep solution presented here, voters do they need to completely understand the underlying vote tallying mechanisms for their votes to be fully heard, nor do not need to speculate on who are the top contest contenders.

2.10.3 Socio-political Implications

Participation in all kinds of group decisions will improve with the convenience of use and the responsive nature of more advanced, network-based social choice mechanisms. The MaxRep choice function's adaptive vote tallying will, when necessary to resolve cycles, take the pertinent relative priorities of each voter into account, irrespective of personal styles of debate and which candidates are top contenders. This choice system can increase the efficiency of the group decision-making process and improve the net equitability of the resulting decision. In the socio-political sphere, preference voting in decisions over many viable alternatives will often improve the morale of the electorate, energize democracy, vastly increase the options available for voters, reflect the electorate's preferences more accurately, and possibly mitigate some of the current divisive nature of "us vs. them" politics. When voters know that all of their preferences will be fully represented, they are more likely to offer well-reasoned and thoughtful input, playing their part in making more effective group decisions.

2.10.4 Closing Perspective

Ever since Condorcet highlighted the dilemma of coinciding, cyclical majorities, people have been contemplating how to deal with this ambiguous outcome. The moderation span we introduce in Part I is perhaps this paper's most significant contribution to this discussion as moderation will reduce the prevalence of cycles. 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 Condorcet's dilemma requires some form of compromise. When cycles do occur, the proposed MaxRep choice function will always select a member of the initial Schwartz top cycle set, picking a winner by reintroducing each voter's pertinent relative priority information.

We anticipate that the adoption of these advances in social choice theory will lead to more congenial and congruous group decisions made across a broad spectrum of scales. At the societal scale, the assessment of the electorate's genuine consensus is a fundamental building block for a truly responsive democracy. Of perhaps more broad reaching effect is the applicability of this choice function to the decisions of smaller social groups. Regardless of scale, a group's agreement on how it will make decisions is fundamentally important. A group's selection of the MaxRep choice function for making decisions is effectively an embrace of MaxRep's level playing field assertion as a mechanism facilitating strategically negotiated compromise.

2.11 Future Research

Further developments in the relationship between social choice and information theory show signs of yielding fruit.[2] Social choice functions represent a challenge in such

[2] "Probabilistic electoral methods, representative probability, and maximum entropy" by Roger Sewell, David MacKay and Iain McLean is an example of a paper using information theory to address gameability

an approach due to the asymmetric nature of information flow and the potentially high order of the internal process: far more information flows into the decision process than emerges in the result. Defining formal metrics for the fidelity of the aggregation of preference information and measures for the effective evaluation of consensus will be a significant step forward. These metrics would allow for direct comparison of the efficacy of proposed social choice functions beyond the current approach of comparing property lists.

The ideas behind the MaxRep choice function could be adapted to other decision scenarios such as proportional representation. Further development of the win-threshold concept from Eq. 23 could also lead progress towards an effective definition of consensus margin. In addition, as discussed in Section 2.8.4, there is also room for research into the susceptibility of MaxRep to strategic gaming, particularly with respect to variation in available information.

We are grateful for the encouraging style at the end of Arrow's Nobel Prize lecture. We would like to similarly encourage others in their quest to further understanding in this challenging field and its potential for significant social impact.

---

issues in voting systems (available at http://www.nuffield.ox.ac.uk/Politics/papers/2004/Sewell%20M-cLean.pdf). MacKay has also authored a book on information theory: Information Theory, Inference, and Learning Algorithms, 2002.

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.

Acknowledgements

Families
The authors would like to thank their families for their consistent support during this project, especially Peter's mother who has made it possible to dedicate the focused time required to make continuing inroads in tackling this challenging problem. Joey would also like to thank his wife, Danielle, for her unwavering love and support.

Editors and reviewers
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.

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." Available http://nobelprize.org/nobel_prizes/economics/laureates/1972/arrow-lecture.html, accessed 15 Aug 2007

[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", available http://fc.antioch.edu/~james_green-armytage/vm/proxy_representation.htm, accessed 30 Nov 2008

[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.