

Objective Measures of Preferential Ballot Voting Systems

Barry Wright, III

April 20, 2009

Submitted for Graduation with Distinction:
Duke University Mathematics Department

Duke University
Durham, North Carolina
2009

Advisor: Dr. Hubert Bray

Abstract:

We consider several types of information with which to differentiate preferential ballot voting systems. After establishing a formalism with which to discuss voting methods in a mathematical context, we show that the aggregation of transitive individual preferences does not always result in a unique transitive social preference ordering. Exposition on eleven different preferential ballot voting systems is then given, as possible methods for resolving elections with three or more candidates. To evaluate these methods, we introduce several desirable conditions on voting systems, and then determine which are satisfied by the various methods studied. Extending this, we construct continuous measures of two conditions, to gain more information when methods do not satisfy certain conditions. Finally, we use election simulations (on a uniform vote distribution) to measure how often each pair of election methods provide the same result. We submit this information as suitable for making a reasoned choice of election method in practical application.

Contents

1	Acknowledgements	5
2	Introduction: The Problem of Aggregating Individual Preferences	6
2.1	The Goal of Voting Systems	6
2.2	Reasonable Criteria	6
2.3	Previous Approaches	6
2.4	Summary of Approach	7
3	Notation and Definitions	8
4	The Difficulty of Resolving Elections	10
5	Voting Systems	11
5.1	Single Vote Plurality	11
5.2	Approval Voting	12
5.3	Preferential Ballot Voting Systems	13
5.4	Instant Runoff Voting	13
5.5	Borda Counts	14
5.6	Instant Runoff Borda Count	16
5.7	Least Worst Defeat	17
5.8	Instant Runoff Least Worst Defeat	18
5.9	Kemeny-Young Method	19
5.10	Schulze Method	20
5.11	Ranked Pairs Method	22
5.12	Copeland Method	23
5.13	Breaking Ties	24
6	System Conditions	25
6.1	Basics	25
6.1.1	Voter and Candidate Symmetry	25

6.1.2	Non-Dictatorship	25
6.1.3	Surjectivity	25
6.1.4	Resolution	26
6.2	Practical Conditions	26
6.2.1	Polynomial Running time	26
6.2.2	Margin of Victory Methods	26
6.2.3	Clear Instruction	27
6.2.4	Voter Purpose	27
6.2.5	Transparency	27
6.3	The Majority Condition	28
6.3.1	Proofs	28
6.3.2	Counterexamples	30
6.4	The Condorcet Condition	31
6.4.1	Proofs	32
6.4.2	Counterexamples	33
6.5	The Copeland Condition	35
6.5.1	Proofs	36
6.5.2	Counterexamples	36
6.6	Monotonicity	42
6.6.1	Proofs	42
6.6.2	Counterexamples	44
6.7	Clone Invariance	47
6.7.1	Proofs	48
6.7.2	Counterexamples	51
6.8	Loser Independence	57
6.8.1	Proofs	58
6.8.2	Counterexamples	59
6.9	Summary Table	61
6.10	Conditions Philosophy	62

7	Simulation of Measurable Conditions	63
7.1	Converting Binary Conditions into Continuously Measurable Conditions . . .	63
7.2	Condorcet and Copeland Ratios	63
8	Agreement Simulations	67
8.1	Philosophy	67
8.2	Simulation of Random Elections	67
9	Conclusions	71
9.1	The Importance of the Condorcet Condition	71
9.2	Monotonicity, Clone Invariance, and Loser Independence	72
9.3	Agreement and Practicality	72
9.4	Future Research	72
10	References	74

1 Acknowledgements

First, I would like to thank my research advisor, Dr. Hubert Bray. His guidance, instruction, and support over the last two years have been invaluable. From the inception of the project (which grew out of a teaching assignment for Dr. Bray's seminar, "Game Theory and Democracy" [1]) to the last days of work (debugging Java code), I could not have completed this work without him.

Second, I would like to thank Dr. Mark Huber, Dr. David Kraines, and Dr. Tom Beale for comprising the committee for my thesis. Their input and interest in my work is greatly appreciated, and I thank them for their support through this year. I also note that I appreciate the support of the Duke PRUV program, for financing my studies during Summer 2007.

Third, I would like to thank my friends for their patient help in reading my work, listening to me talk about it, and discussing various presentation ideas with them. I'd especially like to thank Agee Springer, Jared Haftel and Jesse Thorner for their help.

Finally, I thank my family for their continued love and support through this effort. Their love and consideration is particularly important in an extended project of this magnitude, when completion is never a foregone conclusion. Their belief in my ability has helped me persevere in creating this work.

2 Introduction: The Problem of Aggregating Individual Preferences

2.1 The Goal of Voting Systems

The purpose of voting, especially in the context of democratic forms of government, is to aggregate the preferences or opinions of individuals and process them to produce a single opinion, which purportedly will be an accurate reflection of the views of the electorate. Since the inception of democracy, determining the best possible method for accomplishing the task of voting has been an important, open question. The question has many facets; political, philosophical, practical, and mathematical. To unravel what makes a good (or the best) voting system requires philosophy, to measure those things and determine whether they can occur requires mathematics.

2.2 Reasonable Criteria

Intuitively, we can quickly declare a number of criteria for a voting system that would universally be deemed “reasonable.” It’s clear that one particular person’s vote shouldn’t determine the outcome of the election (this is called nondictatorship), it’s clear that the order of the candidates and voters shouldn’t affect the outcome (candidate and voter symmetry), and it’s clear that my voting for a candidate shouldn’t decrease their final ranking (this is called monotonicity). While each of these conditions are intuitive, they require rigorous mathematical formulation to test. Further, there are other, more subtle criteria that are not obviously reasonable, but prevent unreasonable outcomes. We’ll see that in elections involving three or more candidates, we can’t have everything. Guaranteeing one condition often precludes another. Determining when this occurs is mathematics, deciding what to do about it is philosophy.

2.3 Previous Approaches

Throughout the past few centuries, but mostly in the last fifty years, researchers have employed a number of different approaches to the question of determining the best voting system. Several philosophers have developed their own methods for aggregating votes, many of which we’ll consider through this paper. They have built these systems around a number of concepts; symmetry, avoiding particular paradoxes, satisfying a certain set of conditions. None of these things can individually crown the best voting system. Researchers have shown (in numerous different combinations) that some sets of reasonable criteria are, in fact, mutually exclusive. A natural approach is then to determine if certain sets of conditions determine a unique voting system, adding conditions to eliminate systems at each stage. Philosophically, we must be careful to add conditions for the sake of the conditions, not in order to cleanly narrow the field of potential voting systems. Yet, still we lack a good way to compare systems which satisfy different sets of conditions, leading us to the approach of this paper.

2.4 Summary of Approach

After setting up the basic notation and definitions for discussing voting systems, we begin by explaining the major voting methods under consideration. Then, we follow the “standard approach” of defining a number of reasonable conditions and classifying the systems based on which conditions they satisfy. This, however, is insufficient. Because no system can satisfy all conditions, inevitably we’ll need to compare systems which satisfy different sets of conditions. We approach this problem in two ways:

First, the majority of work concerning voting systems and various criteria has been binary, that is, a focus on determining whether a particular system satisfies a given condition or not. Some work has been done in computing the probability of the occurrence of various paradoxes (especially Condorcet’s Paradox), in the abstract setting of some type of random elections. We extend the spirit of this work to consider the probability that given voting systems abide by the criteria. In this way, we intend to provide a more complete profile of information for the philosophers to discuss. A voting system which fails several criteria, but only does so in one out of every million cases may very well be more desirable than another criteria which passes other criteria, but fails another over half the time. Measuring these conditions in a probabilistic sense will allow us to make more informed decisions in comparing voting systems.

Second, we recognize that if two voting systems produce the same social preference given *any* set of individual votes, then they should be considered the same system (even if this is not obvious based on the method of calculation). Further, measuring how often two voting systems agree, and determining in what situations they disagree provides another dimension of information to consider their quality. This information will also be useful in introducing comparisons based on method ease and efficiency, as we’ll be able to determine exactly what is lost by moving to a more complicated method.

3 Notation and Definitions

The study of the aggregation of social preferences is not as universal a topic as say, calculus or algebra. As such, there is a comparably wider range of definitions and notations used in the literature (which have evolved over time, but still remain relatively diverse). Each writer has their own personal bias, so it is important to clearly lay out notation from the beginning. The exposition of notation also serves as a good introduction to the definitions, jargon, and basic concepts of voting theory.

Let's begin with the basics. Any election (which we'll define precisely) must contain at least two **candidates** (also called **alternatives**), with the implication that either a ranking of candidates or a single "winning" candidate is desired. We will denote candidates by capital letters; $A, B, C...$ typically beginning with A . Further, we will denote the number of candidates in an election by m .

For our purposes, we can consider **votes** and **voters** as equivalent objects (we require no information about who made which vote, anonymity which is typically held sacred in modern electoral processes). The implication is, of course, that each voter makes exactly one vote. Also, because of the anonymity, we will not need to refer to specific voters to distinguish them, as we might need to with candidates. We will denote the number of votes/voters in an election by n .

While we'll see that different voting systems are based on different interpretations of what a vote is, the vast majority (that we'll consider) are based on a set of **pairwise preferences**. Intuitively, a pairwise preference is a single voter's opinion on two candidates. We denote, for example, the preference of A over B by $A \succ B$. We can represent this as a *linear binary relation* on the set of candidates, following the convention of Markus Schulze (though he does not tacitly assume linearity). This provides a few basic, reasonable properties:

The relation \succ is **antisymmetric** and **linear**. This means for any distinct candidates A, B , either $A \succ B$ (exclusive) or $B \succ A$. As a technical note, we say $A \equiv A$, indicating that a candidate can't be preferred over itself.

Further, \succ is transitive, meaning that for distinct candidates A, B, C ;

$$(A \succ B \text{ and } B \succ C) \Rightarrow A \succ C \tag{1}$$

This is a nontrivial assumption about our voters; specifically that their pairwise preferences are **rational**. Consideration of non-transitive pairwise preferences is important, but not something we'll consider deeply in this paper.

Given a pairwise preference for each pair of candidates in the election, there is a unique **preference ordering** of the candidates. For example, if a voter has the following pairwise preferences:

$$A \succ B, A \succ C, C \succ B \tag{2}$$

then their preference ordering is:

$$A \succ C \succ B \tag{3}$$

Typically, voting systems which collect all pairwise preferences as input will have voters provide this information as a preference ranking, because it is faster (though it could be argued that thinking about each individual pairwise matchup is simpler for the voter). We call such systems (technically those that *use* all pairwise information) **preferential ballot voting systems**.

Finally, we are in a position to define an election. An **election** is simply a set of votes (typically preference orderings) on a set of candidates. Generically, we'll denote an election by a capital E (making distinction between a candidate and an election if necessary). We'll see that the most important characteristics of an election are the number of candidates, m and the number of voters n . Therefore, we define spaces of elections $\epsilon(m, n)$ to contain all possible elections of n voters, voting on m candidates. This will be especially important when considering random elections later on.

As we mentioned above, the study of voting systems would not exist if different voting systems didn't give different results on the same election. This idea motivates the precise definition of voting systems.

A **voting system** is a function such that for each reasonable election space (positive numbers of candidates and voters) $\epsilon(m, n)$ the voting system selects on vote on m candidates for each election $E \in \epsilon(m, n)$. This vote (typically a preference ordering) is called a **social preference ordering**. We will denote voting systems as functions, and thus by a lower case letter (or short string of letters).

Finally, we require notation for measuring the aggregation of individual votes.

Given two candidates A, B , we'll denote the number of voters who prefer A to B (that is, the number of voters with pairwise preference $A \succ B$) by $[A, B]$. Similarly, the number of voters with pairwise preference $B \succ A$ is denoted by $[B, A]$. By convention, we set $[A, A] = 0$.

A convenient way to store this information is the **margin of victory matrix** M . This is an antisymmetric matrix with the following entries:

$$M_{AA} = 0 \tag{4}$$

$$M_{AB} = [A, B] - [B, A] \tag{5}$$

This will be a very important tool for the computation of several voting systems. In fact, there are a number of systems which determine the social preference ordering solely from the margin of victory matrix. These methods are called **margin of victory matrix voting methods**.

4 The Difficulty of Resolving Elections

Now, we'll use the formalism developed in the previous section to demonstrate that while two candidate elections are easily determined, elections with three or more candidates do not always have an obvious resolution.

Recall that we required each voter to submit a *transitive* preference ordering. We can encode this transitive ordering as a single-voter margin of victory matrix, which we'll call a **vote matrix**. It is then clear that an equivalent definition of the margin of victory matrix is the sum of the individual (transitive) vote matrices.

Now, while each transitive preference order corresponds to a unique vote matrix, a margin of victory matrix does not necessarily correspond to a transitive preference ordering. Put another way, the sum of transitive preferences is not necessarily transitive. Consider the following example:

The author and two friends are ordering a pizza, but can only afford a single topping; sausage (S), pepperoni (P), or tomato (T). The preference orderings are as follows;

Voter	Preference Order
1	$S \succ P \succ T$
2	$P \succ T \succ S$
3	$T \succ S \succ P$

The resultant margin of victory matrix is then;

$$\begin{pmatrix} & S & P & T \\ S & 0 & 1 & -1 \\ P & -1 & 0 & 1 \\ T & 1 & -1 & 0 \end{pmatrix}$$

We can see that this seems to imply an intransitive cycle of preferences $S \succ P \succ T \succ S \succ \dots$, which is not a reasonable social preference result. While in some sense this result (often called Condorcet's Paradox) [7] is a generic tie, we must develop means for selecting a winner when the magnitudes of victory are non-identical.

We can show by exhaustion that any 2×2 margin of victory corresponds to a transitive preference (either $A \succ B$ or $B \succ A$) unless there is an exact tie. Thus, all two candidate elections are easily resolvable. On the other hand, as the above example shows, this is not the case for elections with more than two candidates. Thus, we are forced to consider exactly how to approach such election scenarios. We should also take care to choose a voting system which selects the implied transitive ordering, when one exists.

5 Voting Systems

We'll continue with an exposition of the most common voting methods (in the community of people who study voting methods), providing salient examples of how they calculate a social preference ordering given an election. The following sections will explain various important conditions for voting systems, and demonstrate why these systems satisfy or fail those conditions.

5.1 Single Vote Plurality

The most familiar voting system in use today is the single vote plurality system. In this system, each voter selects a single candidate to vote for. The votes are tallied, and the candidate with the largest number of total votes is then the winner. If a ranking of candidates is desired, we simply order them based on number of votes received, from highest to lowest. We'll denote the plurality vote by p , and in this case the individual votes are simply one candidate (the voter's preferred candidate). Consider the following example election:

Candidate	Votes
A	30
B	55
C	21
D	3

Because candidate B receives the most votes, B is the winner of the election. We denote this as $p(E) = B$. The social preference ordering, if desired, is $B \succ A \succ C \succ D$.

Three reasons for the relative ubiquity of this system are:

- The system is extremely simple, both to vote in and to compute the result.
- The system satisfies all reasonable conditions in elections with two candidates, in fact, it is the de facto choice for two candidate elections.
- The system has been entrenched in the political landscape, which influences our perception of how elections should be run.

Unfortunately, many problems can occur in plurality elections involving three or more candidates, which we'll detail in the following section. One commonly understood problem is that voting candidates with small chances of winning (for example, third-party candidates in the United States) often feels like "wasting" ones vote. This can cause voters to represent their true preferences dishonestly, by voting for their second-choice candidate.

5.2 Approval Voting

In Approval Voting, the voter is instructed to cast a vote for every candidate he or she would “approve” of winning the election. For example, in determining what type of ice cream to buy for a class, a teacher may ask each student to vote for every flavor that they would be able to eat. No distinction is made based on preference, each candidate is either approved or denied by the voter. As in Single Vote Plurality, the winner is the candidate which receives the most votes (“approvals”), and a ranking of candidates can be had by ordering based on total number of votes. We’ll denote this voting method by a . Consider the following example election:

Candidates Approved	Not Approved	Number of Votes
A, B, C	none	8
A, B	C	14
A, C	B	20
B, C	A	22
A	B, C	9
B	A, C	10
C	A, B	14
none	A, B, C	5

These represent all the votes of the election. To compute $a(E)$ we need to tabulate the total number of approvals for each candidate:

Candidate	Total Approvals
A	51
B	54
C	64

Thus, we have that $a(E) = C$, and that the social preference ordering is $C \succ B \succ A$.

Note that the approval voting method does not favor polarizing candidates. One can easily imagine the following election given two very polarizing candidates A, B and one moderate candidate C . The first column of the table, **true preferences** represent the “internal” knowledge each voter calls on when voting (this is only relevant when the voting system does not ask for the full preference ordering).

True Preferences	Candidates Approved	Not Approved	Number of Votes
$A \succ C \succ B$	A, C	B	81
$B \succ C \succ A$	B, C	A	75

Candidate	Total Approvals
A	81
B	75
C	156

Despite not being the favorite candidate of any voter, candidate C is the approval winner in a landslide! Again, this is not a philosophical statement, the merits of this characteristic can be debated. We can only use the mathematics to uncover these various characteristics of voting systems.

5.3 Preferential Ballot Voting Systems

As we defined above, any system which takes into account each voter's entire preference ordering is a **preferential ballot voting system**. It is important, now, to make a distinction between these first two voting methods and those that will follow. Single Vote Plurality and Approval Voting collect no knowledge about the various pairwise preferences the voters hold. A single vote system only learns what each voters first choice is, rather than a full profile, first through last. Approval Voting, while allowing multiple votes, only breaks the candidates into two sets, approved and denied, and collects no ranking information within these sets. We can anticipate the inadequacy of these voting methods simply based upon the fact that they take in and use less information than the preferential ballot systems which follow.

5.4 Instant Runoff Voting

The most popular preferential voting system is Instant Runoff Voting, also known as Alternative Voting, and it is an iterative process. In each round, we first check if any candidate has a majority of the first place votes. If so, that candidate is selected as the winner. Otherwise, the last place candidate (that is, the candidate with the fewest number of first place votes) is eliminated, and removed from all ballots. Thus, any voter who selected the losing candidate as their first preference now has a different first preference. This process is iterated until some candidate gains a majority of first place votes. If a full ranking of candidates is desired, the process can be continued with second place votes, third place votes, and so on. This voting system will be denoted by *irv*. Let's walk through a multi-stage example election:

Preference Ordering	Number of Votes
$A \succ B \succ C \succ D$	12
$A \succ C \succ B \succ D$	10
$A \succ D \succ B \succ C$	6
$B \succ A \succ D \succ C$	10
$C \succ B \succ A \succ D$	4
$C \succ B \succ D \succ A$	5
$D \succ A \succ C \succ B$	3
$D \succ B \succ A \succ C$	14

First place votes at the first pass:

Candidate	First Place Votes
<i>A</i>	28
<i>B</i>	10
<i>C</i>	9
<i>D</i>	17

Thus, candidate *C* is eliminated from the elections and the ballots. There are two ways to represent this; to simply remove *C* from all ballots, essentially moving to an election with three candidates, or to move *C* to the last position on each preference ordering. We'll use the latter method, since it will be useful when programming random elections under Instant Runoff Voting later on. Moving to the second stage:

Preference Ordering	Number of Votes
$A \succ B \succ D \succ C$	22
$A \succ D \succ B \succ C$	6
$B \succ A \succ D \succ C$	14
$B \succ D \succ A \succ C$	5
$D \succ A \succ B \succ C$	3
$D \succ B \succ A \succ C$	14

First place votes at the second pass:

Candidate	First Place Votes
<i>A</i>	28
<i>B</i>	19
<i>C</i>	eliminated
<i>D</i>	17

Since all of the voters who voted for *C* had *B* as their second choice, *B* gets by on the skin of his/her teeth, forcing candidate *D* to be eliminated. This sets up the final round in which we see a majority winner:

Preference Ordering	Number of Votes
$A \succ B \succ D \succ C$	31
$B \succ A \succ D \succ C$	33

Thus, we have that $irv(E) = B$. Again, this may be an unexpected result given how few first place votes the winning candidate had at the onset. Such is one facet of the nature of the Instant Runoff Voting system. We do note that Instant Runoff will agree with Plurality voting when there is a majority winner (one receiving at least half of the first-place votes), in other cases (like this one), however, the results can be very different.

5.5 Borda Counts

There are several different versions of the Borda Count, a method often attributed to Jean-Charles de Borda (1770) [2], though there have been many independent developments. We'll

adopt the following version; given a voter's preferential order, we award n points to the voter's first preference, $n - 1$ points to the voter's second preference, and so on, until the voter's least preferred candidate receives 1 point. The winner is then the candidate who scores the most points, after adding the points from all voters. The social preference ordering is simply a descending rank order of points scored.

Typically, any point-based system is classified as a Borda Count, even if there is a different weighting scheme (some systems are based on having the fewest number of points). For systems ranking candidates from most points to least, we require that an m th-place vote score at least as many points as an $(m + 1)$ th-place vote, though typically this inequality will be made strict to ensure distinction between places. It is important to note that Borda Counts with different weight schemes are **not** necessarily equivalent (in fact, if the schemes are not scalar multiples, there will always exist a set of votes which produces different election outcomes).

We'll denote this voting system by bc . Consider the following example election:

Preference Ordering	Number of Votes
$A \succ B \succ C \succ D$	5
$D \succ B \succ C \succ A$	2
$C \succ B \succ D \succ A$	1
$C \succ D \succ B \succ A$	1

Recall that candidates will receive 4 points for a first-place vote, 3 points for a second-place vote, 2 points for a third-place vote, and 1 point for a last-place vote. Thus, the Borda Count scores are:

Candidate	Score
A	24
B	26
C	22
D	18

Thus, we have that $bc(E) = B$. Notice that the plurality winner, candidate A does not win, while a candidate with zero first-place votes, candidate B does win. To see the effects of a different weighting system, consider one which rewards first-place votes, giving 5 points for a first-place vote instead of 4:

Candidate	Score
A	29
B	24
C	24
D	20

Now, the election winner is given to be $bc_2(E) = A$. As we might expect, the candidate with the most first-place votes was rewarded with the victory. Notice how even a one-point change can alter the election structure (the effect is, of course, increased if we increase the number

of voters). This means we must be very careful in determining a Borda Count weighting system, since that decision will affect what types of results we value (polarizing candidates, widely accepted candidates, etc.).

We can also compute the Borda Count social preference order by summing the rows of the margin of victory matrix. To see why, consider this deconstruction of the Borda Count score. Since even a last place candidate gets 1 point, each candidate automatically gets n points, where n is the number of voters. Then for each pairwise victory, the candidate must be ranked one slot above another candidate on a particular ballot. Thus, the remaining points are exactly equal to the number of pairwise victories the candidate has. Since there is a clear bijection between the total number of pairwise victories and the sum of the entries in a candidate's row of the margin of victory matrix, we can simply use this value (which is easier to compute when programming elections).

5.6 Instant Runoff Borda Count

Instant Runoff Borda Count behaves in a slightly different manner than Instant Runoff Voting. Instead of checking for a majority winner in each round, the process always iterates until a single candidate remains. In each round, the candidate with the worst (least or most, depending on the point allocation system) Borda Count score is eliminated, and removed from the ballots, prompting a recalculation of the scores. This proceeds until a single candidate remains, though formulas can be calculated which will indicate a guaranteed winner given a particular point allocation system.

This method will be denoted by *irbc*. We'll use the same preferences as in the previous example, and the standard Borda Count weighting to conduct an example election. Recall the initial scores:

Candidate	Score
A	24
B	26
C	22
D	18

This means that candidate D is eliminated. We then recalculate the preference orderings and then calculate the new Borda scores:

Preference Ordering	Number of Votes
$A \succ B \succ C$	5
$B \succ C \succ A$	2
$C \succ B \succ A$	2

Now that there are only three candidates in consideration, the point values will be 1, 2, and 3 (worst-to-best).

Candidate	Score
<i>A</i>	19
<i>B</i>	20
<i>C</i>	15

In the second pass, candidate *C* is eliminated, leading to one final recalculation (which amounts to a two-person plurality election under the standard weighting):

Preference Ordering	Number of Votes
$A \succ B$	5
$B \succ A$	4

Thus, clearly we have that $irbc(E) = A$. (For reference, the Borda Count scores are 14 for candidate *A* and 13 for candidate *B*). Interestingly, the candidate with the highest current score in the initial rounds did not win the election.

5.7 Least Worst Defeat

The Least Worst Defeat method is a margin of victory matrix method. In this method, we determine the worst defeat of each candidate (the minimum value in a candidates row of the margin of victory matrix). The candidate with the greatest such value (that is, with the least severe defeat, or no defeat if the value is zero) is the winner. While the method can be extended to provide a ranking of candidates (by second-least worst defeat, third-least worst defeat, etc.), this is typically not done. We'll denote this method by *lwd*. Consider the following example election, as represented by the margin of victory matrix (we'll label the rows and columns for each candidate, the letters signifying them are, of course, not part of the matrix):

$$\begin{pmatrix} & A & B & C & D & E & F \\ A & 0 & -5 & -3 & 1 & -7 & -5 \\ B & 5 & 0 & 7 & 1 & -3 & 5 \\ C & 3 & -7 & 0 & -5 & 7 & 1 \\ D & -1 & -1 & 5 & 0 & 11 & 3 \\ E & 7 & 3 & -7 & -11 & 0 & 1 \\ F & 5 & -5 & -1 & -3 & -1 & 0 \end{pmatrix}$$

We compile the worst defeat of each candidate (minimum number in the candidate's row of the margin of victory matrix):

Candidate	Worst Defeat
<i>A</i>	-7
<i>B</i>	-3
<i>C</i>	-7
<i>D</i>	-1
<i>E</i>	-11
<i>F</i>	-5

Thus, since candidate D has the greatest worst defeat (greatest numerically, least in terms of severity), we have that $lwd(E) = D$. Notice that the candidate with the fewest number of defeats (candidate B) does not necessarily win.

5.8 Instant Runoff Least Worst Defeat

Similar to the Instant Runoff Borda Count, we use a full iterative process to determine the winner of Instant Runoff Least Worst Defeat. In each round, the candidate with the “worst” (the minimum number) worst defeat is eliminated, and their row and column is removed from the margin of victory matrix. The process continues until we are left with a single candidate, the winner, and a full ranking can be determined by listing the candidates in reverse order of elimination. We’ll denote this method by *irlwd*, and use the same starting matrix as the *lwd* example for this example election. We’ll show the successive matrices, with eliminated candidate and worst defeat in boldface:

$$\begin{pmatrix} & A & B & C & D & E & F \\ A & 0 & -5 & -3 & 1 & -7 & -5 \\ B & 5 & 0 & 7 & 1 & -3 & 5 \\ C & 3 & -7 & 0 & -5 & 7 & 1 \\ D & -1 & -1 & 5 & 0 & 11 & 3 \\ \mathbf{E} & 7 & 3 & -7 & \mathbf{-11} & 0 & 1 \\ F & 5 & -5 & -1 & -3 & -1 & 0 \end{pmatrix}$$

$$\begin{pmatrix} & A & B & C & D & F \\ A & 0 & -5 & -3 & 1 & -5 \\ B & 5 & 0 & 7 & 1 & 5 \\ \mathbf{C} & 3 & \mathbf{-7} & 0 & -5 & 1 \\ D & -1 & -1 & 5 & 0 & 3 \\ F & 5 & -5 & -1 & -3 & 0 \end{pmatrix}$$

The next iteration encounters a tie. There are two ways to deal with this, to look at the second-worst defeat, or to utilize a tiebreaking vote. Later, we’re going to explain the methodology of a tiebreaking vote, so for now we’ll just use the second-worst defeat.

$$\begin{pmatrix} & A & B & D & F \\ \mathbf{A} & 0 & \mathbf{-5} & 1 & \mathbf{-5} \\ B & 5 & 0 & 1 & 5 \\ D & -1 & -1 & 0 & 3 \\ F & 5 & -5 & -3 & 0 \end{pmatrix}$$

$$\begin{pmatrix} & B & D & F \\ B & 0 & 1 & 5 \\ D & -1 & 0 & 3 \\ \mathbf{F} & \mathbf{-5} & -3 & 0 \end{pmatrix}$$

$$\begin{pmatrix} & B & D \\ B & 0 & 1 \\ D & -1 & 0 \end{pmatrix}$$

Thus, we have that candidate B is the winner ($irlwd(E) = B$).

5.9 Kemeny-Young Method

The Kemeny-Young method is another system based off of the margin of victory matrix. This method evaluates all possible social preference orderings, and determines the one that matches the margin of victory matrix based on a scoring procedure. It can be computed either using the number of votes for a particular candidate in each pairwise matchup, or using the margins of victory of each pairwise matchup. Since we already have the margin of victory matrix available, we'll use the latter formulation.

Given a preference ordering of m candidates, for example $B \succ A \succ C$, there are $\frac{m(m-1)}{2}$ distinct pairwise matchups that are implied by rationality. In this case, we have that $B \succ A$, $A \succ C$, and $B \succ C$. The Kemeny-Young method then assigns a score to this social preference ordering by adding up the values of the margin of victory matrix corresponding to these pairwise matchups. Consider the following example margin of victory matrix:

$$\begin{pmatrix} & A & B & C \\ A & 0 & -11 & -3 \\ B & 11 & 0 & 13 \\ C & 3 & -13 & 0 \end{pmatrix}$$

Then, the Kemeny-Young Score for the preference ordering would then be $11 + (-3) + 13 = 21$. We'll use this as an example election and calculate the Kemeny-Young scores for all possible preference orderings (there are $3! = 6$ of them):

Preference Ordering	Kemeny-Young Score
$A \succ B \succ C$	-1
$A \succ C \succ B$	-27
$B \succ A \succ C$	21
$B \succ C \succ A$	27
$C \succ A \succ B$	-21
$C \succ B \succ A$	1

Thus, we see that the social preference ordering is $B \succ C \succ A$, and thus the election winner is B .

We also mention, briefly, a practical condition (which we'll consider further in the next section). This method requires that we compare the scores of every possible preference ordering, and there are $m!$ such orderings in an election with m candidates. This poses a

computational problem, and while some algorithms have been developed to reduce computation, no polynomial running time algorithms have been found. In elections with a large number of candidates, this may be a prohibitive factor.

5.10 Schulze Method

The Schulze Method is another preferential ballot method which utilizes the margin of victory matrix, in a way. Instead of considering the direct pairwise matchups, we consider a sort of indirect defeat, based on the concept of a **path**.

Instead of using the direct margin of victory matrix, we construct a pairwise vote matrix, in which the entry M_{AB} is equal to the number of votes which have $A \succ B$ (the quantity earlier defined as $[A, B]$). Given a margin of victory matrix and the total number of votes, it's easy to compute the pairwise vote matrix, as in the following example (with 99 voters) and the following margin of victory matrix:

$$\begin{pmatrix} & A & B & C \\ A & 0 & 17 & -23 \\ B & -17 & 0 & -31 \\ C & 23 & 31 & 0 \end{pmatrix}$$

The associated pairwise vote matrix is then:

$$\begin{pmatrix} & A & B & C \\ A & 0 & 58 & 38 \\ B & 41 & 0 & 34 \\ C & 61 & 65 & 0 \end{pmatrix}$$

Now, given this matrix, we can define a path from candidate A to candidate B as follows.

A **path** from candidate A to candidate B is a sequence of candidates $C(1), \dots, C(n)$ such that the $C(i)$ are distinct, $C(1) = A$, $C(n) = B$, and for all $i < n - 1$, $[C(i), C(i + 1)] > [C(i + 1), C(i)]$ (that is, a chain of wins connecting A to B).

We can then define the **strength** of a path as the minimum value of all $[C(i), C(i + 1)]$, that is, the weakest victory along the path.

The Schulze Method then proceeds in the following manner. Given two candidates, we define $p[A, B]$ to be the strength of the strongest path from candidate A to candidate B . Then, we demand that $A \succ B$ if and only if $p[A, B] > p[B, A]$ (if no path from A to B exists, we take $p[A, B] = 0$). Further, we define a candidate C as a **potential winner** if and only if $p[C, D] \geq p[D, C]$ for every other candidate D .

In the vast majority of cases (including if we restrict ourselves to cases in which every nonzero element of the margin of victory matrix is unique), there will be as single potential

winner. However, if there are multiple potential winners, we must either use a tiebreaking vote or a shared victory concept.

Consider the following example election (as represented by a pairwise vote matrix) [3]:

$$\begin{pmatrix} & A & B & C & D & E \\ A & 0 & 20 & 26 & 30 & 22 \\ B & 25 & 0 & 16 & 33 & 18 \\ C & 19 & 29 & 0 & 17 & 24 \\ D & 15 & 12 & 28 & 0 & 14 \\ E & 23 & 27 & 21 & 31 & 0 \end{pmatrix}$$

First, let's look at an example of how to determine a strongest path, in this case $p[A, B]$.

Since $[A, B] < [B, A]$ ($20 < 25$), we can't directly pass from candidate A to candidate B in our path. Similarly, we can't first pass to E . One possible path is A, C, B . This would have a strength of 26, since $[A, C] < [C, B]$. But there is a stronger path; if we proceed A, D, C, B , the weakest link is $[C, B] = 29$. Examination of the other possible paths shows that this is the strongest possible path, so that $p[A, B] = 29$:

Path	Strength	Weakest Link
A, C, B	26	$[A, C]$
A, C, E, B	24	$[C, E]$
A, C, E, D, B	24	$[C, E]$
A, D, C, B	29	$[C, B]$

We then calculate the strength of the path between each pair of candidates, arranging this information in a matrix:

$$\begin{pmatrix} & A & B & C & D & E \\ A & 0 & 28 & 28 & 30 & 24 \\ B & 25 & 0 & 28 & 33 & 24 \\ C & 25 & 29 & 0 & 29 & 24 \\ D & 25 & 28 & 28 & 0 & 24 \\ E & 25 & 28 & 28 & 31 & 0 \end{pmatrix}$$

Then, we see that candidate E is the only potential winner, since it is the only candidate for which $p[E, X] \geq p[X, E]$ for all other candidates X . Thus, E is the election winner. Writing out the pairwise path strength matchups, we generate the following social preference ordering; $E \succ A \succ C \succ B \succ D$.

For a very detailed account of the Schulze Method, consult reference [3], the first in five extensive papers written by Markus Schulze, inventor of this method.

5.11 Ranked Pairs Method

The Ranked Pairs method (also called Tideman method, since it was developed by Nicolaus Tideman) [4] uses preferential ballots and the margin of victory matrix to rank pairwise matchups based on margin of victory (thus taking into account both number of victories, and strength of victories). It essentially follows a three step process:

- First, determine the margin of victory matrix (i.e. calculate all pairwise matchups). For the purposes of this exposition, we'll assume (as is reasonable in large elections) that the nonzero entries of the margin of victory matrix are unique. Otherwise, tiebreaking votes may be required.
- Second, rank each pairwise matchup by margin of victory, largest to smallest.
- Finally, determine the preferential ordering by mandating (“locking-in”) each pairwise matchup (beginning with the largest margin of victory) unless adding a matchup would create an intransitive cycle in the preference order (for example $A \succ B$, $B \succ C$, and $C \succ A$).

The process is best understood with an example election. Consider the following margin of victory matrix:

$$\begin{pmatrix} & A & B & C & D \\ A & 0 & 17 & -7 & 25 \\ B & -17 & 0 & 13 & 5 \\ C & 7 & -13 & 0 & 3 \\ D & -25 & -5 & -3 & 0 \end{pmatrix}$$

We then rank the various pairwise defeats by magnitude:

Pairwise Result	Margin of Victory
$A \succ D$	25
$A \succ B$	17
$B \succ C$	13
$C \succ A$	7
$B \succ D$	5
$C \succ D$	3

We can then begin locking in the results in order (removing the from the ranking as we go along). We'll proceed either until a full preference ordering has been determined, or we encounter a cycle-producing pairwise result:

The first three lines go smoothly; in our preference ordering, we must have that $A \succ D$, $A \succ B$, and $B \succ C$. Note that the last two results create a chain; by transitivity (rationality), we have that $A \succ B \succ C$. This is a requirement of the Ranked Pairs method. Thus, the next remaining entry would cause a cycle ($A \succ B \succ C \succ A$):

Pairwise Result	Margin of Victory
$C \succ A$	7
$B \succ D$	5
$C \succ D$	3

Therefore, we remove this result (disregarding it), and move on to the remaining pairwise matchups:

Pairwise Result	Margin of Victory
$C \succ A$	7
$B \succ D$	5
$C \succ D$	3

The rest is straightforward, since no more cycles occur; we add $B \succ D$ and $C \succ D$. Thus, the final social preference ordering is $A \succ B \succ C \succ D$.

5.12 Copeland Method

The Copeland Method is based on a simplified version of the margin of victory matrix called the **win-loss matrix**, which is simply the sign of the margin of victory matrix. For example, given the following margin of victory matrix:

$$\begin{pmatrix} 0 & 5 & 1 & -3 \\ -5 & 0 & 3 & -5 \\ -1 & -3 & 0 & -9 \\ 3 & 5 & 9 & 0 \end{pmatrix}$$

The associated win-loss matrix would then be:

$$\begin{pmatrix} 0 & 1 & 1 & -1 \\ -1 & 0 & 1 & -1 \\ -1 & -1 & 0 & -1 \\ 1 & 1 & 1 & 0 \end{pmatrix}$$

Given the win-loss matrix, we can define another measure of electoral success, the **Copeland score** to be the number of 1's in a candidate's row of the win-loss matrix. Conceptually, this is exactly the number of pairwise contests that the candidate wins. Determining a winner is simple; the candidate with the highest Copeland score wins. If there is a tie, we must either use a tiebreaking vote or declare a shared victory (this concept will be developed further later on). We'll refer to this method as *cp*. Consider the win-loss matrix above as an example, the candidates have the following copeland scores:

Candidate	Copeland Score
A	2
B	1
C	0
D	3

Thus, we have that $cp(E) = D$. Notice the reasonable outcome; the candidate which beats all others in pairwise competition was declared the winner. We'll revisit this concept when discussing conditions on voting systems. Let's consider another example win-loss matrix, in which we have a tie:

$$\begin{pmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix}$$

Here we see that each candidate has a Copeland score of 1. At this point we would have to utilize the tiebreaking vote or accept a shared victory among all three candidates.

5.13 Breaking Ties

While ties in the margin of victory matrix are rare (occurring roughly with probability equal to $\frac{1}{\sqrt{rtn}}$, n the number of voters) [1], they do cause problems in performing the algorithms required by several of our voting systems. There is no prescribed course of action, for example, when two preference orders have exactly equal Kemeny-Young scores, or if several Ranked Pairs entries are zero. Thus, we're motivated to develop a tiebreaking procedure to ensure that all non-diagonal entries in the margin of victory matrix are not only nonzero, but *unique*.

To do this, we introduce the following tie-breaking vote matrix. Given a small value ϵ (for numerical convenience, we use $\epsilon = 0.1$, $T_{ij} = \epsilon^i - \epsilon^j$). We can see clearly that this will not change the results of the election (since every entry is significantly less than one-half), and that it ensures that each entry of the new margin of victory matrix M' is unique when we add M and T .

6 System Conditions

6.1 Basics

Before we consider the various system conditions used to differentiate the major voting methods, we need to set forth a few extremely basic criteria that all reasonable voting methods should satisfy. We consider this concepts to be so vital/basic to the purpose of voting that they should be a first requirement for all voting systems, before we employ mathematics and philosophy to make decisions.

6.1.1 Voter and Candidate Symmetry

It is clear that in democratic voting systems, each voter should have equal influence and each candidate should have an equal chance. Rigorously, we mean that a rearrangement of the voters or candidates (by name or position) should have no effect on the election. This prevents any voter from having a vote that is worth more than another voter's.

We understand that in business situations, it may not be desirable to have equality in voting, but for the purposes of democratic voting (and this paper), we'll stick to that restriction.

6.1.2 Non-Dictatorship

Next, we require that a voting method take into account the votes of all the voters (this is essentially implied by voter symmetry above). Formally, if there are multiple voters in an election, we can not have the social preference equal to one voter's preference in all elections where that preference is held constant.

As an example of a method which fails both non-dictatorship and voter symmetry, consider the **Barry Votes Method**. In this method, the vote of Barry Wright, III (that's me) is equal to the social preference. No matter the votes of all the other voters, the social preference does not change (this violates non-dictatorship). Further, if I switch names with someone else, the value of our votes certainly change, violating voter symmetry.

6.1.3 Surjectivity

When considering voting systems which produce a full ranking of preferences (social preference ordering), we require that the voting system f be surjective. This means that for every preference ordering p on a set of candidates, there exists a set of votes (an election E) such that $f(E) = p$. To see that all of the election systems discussed so far exhibit surjectivity, let E be the election with a single voter, and have that voter vote p .

As an example of a method which fails surjectivity, consider the **Barry Wins Method**. In this method, no matter what votes are cast, Barry wins, and the remaining places are

chosen randomly. Thus, any preference ordering which does not list Barry first can not be achieved, violating surjectivity.

6.1.4 Resolution

Finally, we require that all voting methods make a choice! This means that for every election E , $f(E)$ has a value (typically a preference ordering). Note that we do not require this choice be deterministic. We'll discuss the question of deterministic systems later in this section.

6.2 Practical Conditions

Our consideration of voting methods must not be completely based on theoretical conditions; in the end, a successful voting system must be carried out and used in real elections, sometimes with millions of voters. Thought must also be taken to ensure that the *input* to the elections (the votes) accurately reflect the true preferences of the voters. The computer science phrase “garbage in, garbage out” certainly applies here. Thus, we have several practical considerations which are not vital to successful voting systems, but must be taken into account when making a selection among different systems.

6.2.1 Polynomial Running time

Especially in dealing with elections involving millions of voters and hundreds of candidates (for example, a worldwide ranking of the greatest college basketball players of all time), the running time of the voting system algorithm can be crucial. Often, elections require quick turnaround of results, and a system which takes three hours to compute results may be strongly favored over one with slightly better voting properties that takes three weeks to compile.

It is also important to distinguish the running time based on the number of voters and the running time based on the number of candidates. In most situations, the number of voters will be significantly greater than the number of candidates. Thus, a system which is $O(m^2)$ and $O(\exp(n))$ would be significantly worse than a system which is $O(\exp(m))$ and $O(n^2)$.

We'll omit the proofs (this is not an algorithms paper), but of the methods we've discussed, only the Kemeny-Young Method fails to have polynomial running time. [5]

6.2.2 Margin of Victory Methods

As mentioned above, margin of victory methods can determine the election winner/social preference ordering solely from the margin of victory matrix. This is often very convenient; a programmer can work based on having just a single matrix as input, and it also guarantees quadratic running time in n . While this is certainly not a necessary condition by any account, it's something worth noting, if only for the convenience.

6.2.3 Clear Instruction

Moving to the user/voter aspect of the methods; it is vital to have clear, easy-to-follow instructions for the voter. In most cases, this is the onus of whoever implements the system, not who creates the system. The classic example is the 2000 U.S. Presidential Election, a plurality election, in which many voters were confused by the setup of the ballot, and thus may have misrepresented their true preferences, calling the legitimacy of the results into question. The chief concern here is that voters will not correctly input their true preferences, which would make it impossible to generate a true social preference ordering.

While this is not the case for any of the methods we described (most simply ask for a ranking of the candidates), we can envision systems which require extremely complicated inputs from the voter, causing confusion based on the system, regardless of implementation. For example, a system which asked for a ranking *of all possible candidate rankings* must be considered suspect, as we can't reasonably expect voters to accurately interpret their true preferences precisely enough to correctly produce such a ranking.

6.2.4 Voter Purpose

Now, the question of voter honesty is broad, and an important topic in voting theory. At the moment, we only want to consider one aspect of voting systems which impacts voter honesty. The voter purpose condition requires that any true preference vote be meaningful to the system, that it establish preferences to be considered and calculated in generating the social preference.

Notably, the approval method fails this criterion. Consider the following mock election; a philanthropist is going to give everyone who votes in an election an amount of money. The choices/candidates are ten dollars, one hundred dollars, and one thousand dollars. Under the approval voting method, I am supposed to vote for "all candidates which I approve of." Disregarding the vagueness of this instruction, I approve of any amount of money, and thus my true preference would be to approve all. However, this effectively negates my vote, since I do not differentiate between the candidates. It's clearly in my best interest to only approve the thousand dollar option, since I prefer it more.

There is no purpose to approving all the candidates in an election (or none of them) from a strictly rational voting perspective, and this very well may induce dishonest voting, something we would like voting systems to avoid.

6.2.5 Transparency

Finally, for a voting system to be successful (and generate high turnout, which improves the accuracy of a social preference decision), it must be accepted by the electorate as fair. Now, we'll see that several of the systems we've described thus far are quite good, in terms of the theoretical conditions which they satisfy. Unfortunately, their mechanism is obfuscated by complicated mathematics and complex algorithms. This makes it difficult for the average voter (who does not have expertise in the study of voting systems) to trust these systems.

This will motivate our consideration of continuous conditions and method agreement testing; if we have a simple system which agrees with a more complex one in 99.99% of elections, it's reasonable to consider preferring the simple method.

6.3 The Majority Condition

A candidate C is a **majority winner** if more than half of all voters list C as their first preference (when considering non-preferential ballot systems, we interpret based on the true preferences of the voters).

A voting system f meets the **majority condition** if for all elections E which have a majority winner C , we have that $f(E) = C$. That is, the system never fails to elect a majority winner, if one exists. This condition stems from our experience with two-candidate elections, in which the winner is always the majority winner (and a majority winner always exists). Note that this condition implies the **two-candidate condition** which states that an election method should give the correct result (majority winner) when run on elections with only two candidates.

6.3.1 Proofs

We now give proofs that the following methods satisfy the majority criterion; Plurality, Instant Runoff, Instant Runoff Borda Count, Least Worst Defeat, Instant Runoff Least Worst Defeat, Kemeny-Young, Schulze, Ranked Pairs, and Copeland. Some of the proofs are quite straightforward.

Plurality:

Suppose there exists a candidate C which is a majority winner. This means that C receives more than $\frac{n}{2}$ votes. Thus, every other candidate must receive less than $\frac{n}{2}$ votes. This implies that candidate C has the most first-place votes, and thus C (the majority winner) is declared the election winner.

Instant Runoff:

Recall that the Instant Runoff algorithm checks for a majority winner at the beginning of each stage. Clearly, if there exists a majority winner C , it will be found and selected at the first stage, making it the election winner.

Instant Runoff Borda Count:

(Note that we are assuming the standard weighting system described in the initial development of the Borda Count.)

Suppose that there exists a candidate C which is a majority winner. If C is the only candidate, then clearly it is selected as the election winner. Otherwise, the Instant Runoff

Borda Count (IRBC) begins eliminating candidates. We prove that candidate C can not be eliminated (and thus, must be selected as the election winner).

Because IRBC eliminates the candidate with the lowest Borda Count during each round, the eliminated candidate can not have a score greater than the average score of all candidates. We begin by calculating this average score. Recall that each voter gives 1 point for their bottom-ranked candidate, 2 the second-worst candidate, and so on up to n points for the top-ranked candidate. Thus, the total number of points in an election in $\epsilon(m, n)$ (m candidates, n voters) is:

$$BC_{total} = \frac{m(m+1)}{2} \times n \quad (6)$$

And thus the average Borda Count score is given by:

$$BC_{average} = \frac{(m+1)n}{2} \quad (7)$$

Now, we compute the minimum possible score of a majority winner. Suppose the majority winner C has *exactly* $\frac{n}{2}$ first-place votes (by definition, C must have more first-place votes than this). Further, in the worst case, every other vote would place C last, so that C has $\frac{n}{2}$ first-place votes and $\frac{n}{2}$ last-place votes. The Borda Count of C is then:

$$BC(C) = m \times \frac{n}{2} + 1 \times \frac{n}{2} = \frac{(m+1)n}{2} \quad (8)$$

Clearly, adding any first-place votes increases the Borda Count, and thus $BC(C) > BC_{average}$, so that the majority winner can not be eliminated during any stage of the IRBC computation. Thus, IRBC satisfies the majority condition.

Least Worst Defeat:

Recall that the Least Worst Defeat method selects the candidate with the greatest minimum value in their row of the margin of victory matrix. If a candidate is a majority winner, it must win all pairwise matchups (if C is a majority winner, then $M_{CA} > 0$ for all $A \neq C$). Thus, the worst defeat for C is no defeat, a value of 0. Since every other candidate has a defeat (in particular, at the hands of candidate C), C must have the least worst defeat. Thus, C , the majority winner, must be selected as the election winner.

Instant Runoff Least Worst Defeat:

Similarly, Instant Runoff Least Worst Defeat eliminates the candidate with the overall worst defeat each round. Since a majority winning candidate has no defeats, it can never be eliminated, and thus will be declared the election winner.

Kemeny-Young:

Suppose that the Kemeny-Young method selects a social preference ordering in which the majority winner M is not selected election winner, for example $A \succ M \succ B \succ C$. Since M is a majority winner, it wins all pairwise matchups, in particular the matchup $M \succ A$. This

means that the value associated with $A \succ M$ is negative when calculating the Kemeny-Young score. Thus, we have (denoting $KY()$ to be the Kemeny-Young score):

$$KY(A \succ M \succ B \succ C) < KY(M \succ A \succ B \succ C) \quad (9)$$

This is a contradiction, thus, M must be selected as the election winner, and Kemeny-Young satisfies the majority condition.

Schulze:

Although the Schulze method is based on indirect defeats (via the path definition), in order for a candidate A to have a path to another candidate B , there must be at least one candidate which defeats B (for which $[X, B] > [B, X]$). But, if B is a majority winner, it wins all pairwise matchups, and thus no candidate can have a path to B . This means that for all other candidates Y , $p[B, Y] > p[Y, B]$, which implies that B must be the election winner.

Ranked Pairs:

Recall that the Ranked Pairs method is based on pairwise matchups. Since the majority winner wins all pairwise matchups, no entry in the Ranked Pairs sort will be able to rank another candidate over the majority winner. This means that the majority winner will be ranked above all other candidates in creating the social preference ordering, and is thus the election winner.

Copeland:

Again, any majority winner must win all pairwise contests, implying that it must have a maximal copeland score of $m - 1$. Thus, every majority winner must also be the election winner.

6.3.2 Counterexamples

We now provide counterexamples to show that the following methods fail the majority condition: Approval and Borda Count.

Approval:

Polarizing candidates, which collect many first and last place votes, are vulnerable to the approval voting method. Recall the example given previously:

True Preferences	Candidates Approved	Not Approved	Number of Votes
$A \succ C \succ B$	A, C	B	81
$B \succ C \succ A$	B, C	A	75

Candidate	Total Approvals
<i>A</i>	81
<i>B</i>	75
<i>C</i>	156

Here, the majority winner *A* is not selected as the election winner. Thus, the approval method fails the majority condition.

Borda Count:

A similar situation can occur in the Borda Count, where a candidate who gets consistently high rankings (even if they are not top rankings) can beat a majority winner which also has many last-place rankings. Consider the example given previously:

Preference Ordering	Number of Votes
$A \succ B \succ C \succ D$	5
$D \succ B \succ C \succ A$	2
$C \succ B \succ D \succ A$	1
$C \succ D \succ B \succ A$	1

Candidate	Score
<i>A</i>	24
<i>B</i>	26
<i>C</i>	22
<i>D</i>	18

Thus, candidate *B* wins out over the majority winner, candidate *A*. Note how this differs from the Instant Runoff Borda Count (in which the majority winner *A* was selected). In the Borda Count, the winner must have the highest total (and it is possible for a majority winner to *not* have the highest total). On the other hand, in the Instant Runoff Borda Count, the winner must avoid having the lowest total in each round (we proved that a majority winner can never have the lowest total).

6.4 The Condorcet Condition

We now consider a stricter version of the majority condition. We define a **Condorcet winner** to be a candidate which beats every other candidate in pairwise matchups. Then, a voting system meets the **Condorcet Condition** if for all elections with a Condorcet winner, *C*, the system selects the Condorcet winner as the election winner.

Notice that every majority winner is also a Condorcet winner. Thus, if a voting system passes the Condorcet condition, it also passes the majority condition (also, if a voting system fails the majority condition, it must fail the Condorcet condition).

The Condorcet condition is very significant, and the first major hurdle in our consideration of voting systems. It is both transparent and reasonable to suggest that a candidate which “defeats all comers” should be elected. It is also counterintuitive, in a sense, to elect a candidate which is not preferred over another candidate (the Condorcet winner). Further, we’ll see that Condorcet winners occur in a large portion of elections; this means that Condorcet methods agree on this large portion as well. Thus, the Condorcet condition is often used as a first filter in selecting an appropriate voting method.

6.4.1 Proofs

Now, we’ll give proofs that the following methods satisfy the Condorcet condition; Instant Runoff Borda Count, Least Worst Defeat, Instant Runoff Least Worst Defeat, Kemeny-Young, Schulze, Ranked Pairs, Copeland.

Instant Runoff Borda Count:

As with the majority condition, we’ll prove that a Condorcet winner must always have greater than the average Borda Count score, and thus can not be eliminated during the Instant Runoff Borda Count process. Recall the average Borda score is given by:

$$BC_{average} = \frac{(m+1)n}{2} \quad (10)$$

Now, given a Condorcet winner C , we have that C defeats every other candidate in more than $\frac{n}{2}$ votes. This means that C must get more than $\frac{n}{2} \times (m-1)$ Borda Count points, simply because of the number of candidates ranked below C on various votes. Further, since even a last-place vote earns 1 point, C (and every candidate) is guaranteed a base of n points. Together we have that;

$$BC(C) > \frac{n}{2} \times (m-1) + n = \frac{(m+1)n}{2} = BC_{average} \quad (11)$$

Thus, since the Condorcet winner always has an above average Borda Count, it can never be eliminated, and is thus the election winner.

Least Worst Defeat:

Again, since a Condorcet winner (like a majority winner) wins all pairwise matchups, and thus has no defeats. This means it has the maximal minimum value (in its row of the margin of victory matrix), and is selected election winner.

Instant Runoff Least Worst Defeat:

Similarly, because the Condorcet winner automatically has a least worst defeat of 0, it can not be eliminated at any stage of Instant Runoff Least Worst Defeat (since every other candidate has a defeat, and thus a negative least worst defeat). Thus, the Condorcet winner is always selected as election winner.

Kemeny-Young:

The proof is congruent to that of the majority condition. Suppose that the Kemeny-Young method selects a social preference ordering in which the Condorcet winner C is not selected election winner. Since C is a Condorcet winner, it wins all pairwise matchups, so there must exist another ordering with strictly greater Kemeny-Young score. This is a contradiction, thus, C must be selected as the election winner, and Kemeny-Young satisfies the Condorcet condition.

Schulze:

As with the majority condition, in order for a candidate A to have a path to another candidate B , there must be at least one candidate which defeats B (for which $[X, B] > [B, X]$). But, if B is a Condorcet winner, it wins all pairwise matchups, and thus no candidate can have a path to B . This means that for all other candidates Y , $p[B, Y] > p[Y, B]$, which implies that B must be the election winner.

Ranked Pairs:

As with the majority criterion, since the Condorcet winner wins all pairwise matchups, no entry in the Ranked Pairs sort will be able to rank another candidate over the majority winner. This means that the Condorcet winner will be ranked above all other candidates in creating the social preference ordering, and is thus the election winner.

Copeland:

As with the majority condition, because a Condorcet winner wins all pairwise matchups, it has the maximal (and unique) Copeland score of $m - 1$. Thus, every Condorcet winner must also be the election winner under the Copeland method.

6.4.2 Counterexamples

We'll provide counterexamples to show that the following methods fail the Condorcet condition; Plurality, Approval, Instant Runoff, and Borda Count.

Plurality:

While all majority winners are Condorcet winners, not all Condorcet winners are majority winners. In fact, a Condorcet winner can have *zero* first-place votes, as demonstrated in this example election:

Preference Ordering	Number of Votes
$A \succ B \succ C \succ D$	34
$C \succ B \succ D \succ A$	33
$D \succ B \succ A \succ C$	33

Now, look at the first place votes and the margin of victory matrix:

Candidate	First-Place Votes
<i>A</i>	34
<i>B</i>	0
<i>C</i>	33
<i>D</i>	33

$$\begin{pmatrix} & A & B & C & D \\ A & 0 & -32 & 34 & -32 \\ B & 32 & 0 & 34 & 34 \\ C & -34 & -34 & 0 & 34 \\ D & 32 & -34 & -34 & 0 \end{pmatrix}$$

Thus, we see that the Condorcet winner *B* comes in last place in the plurality election (candidate *A* has the most first-place votes, and is thus selected as the plurality election winner).

Approval:

Because the Approval method fails the majority criterion, it must also fail the Condorcet criterion. The counterexample given for the majority criterion is also a counterexample for the Condorcet criterion (the majority winner is also the Condorcet winner, and is not elected).

Instant Runoff:

As with the plurality method, Instant Runoff voting is highly dependent upon getting first place votes, which Condorcet winners do not necessarily have a large quantity of. Looking at the example election given to show that the plurality method fails the Condorcet condition again: votes, as demonstrated in this example election:

Preference Ordering	Number of Votes
$A \succ B \succ C \succ D$	34
$C \succ B \succ D \succ A$	33
$D \succ B \succ A \succ C$	33

Candidate	First-Place Votes
<i>A</i>	34
<i>B</i>	0
<i>C</i>	33
<i>D</i>	33

Because candidate *B* (the Condorcet winner) has zero first-place votes, it is eliminated in the first pass (and is thus not selected as the election winner).

Borda Count:

Because the Borda Count fails the majority criterion, it must also fail the Condorcet criterion. The counterexample given for the majority criterion is also a counterexample for the Condorcet criterion (the majority winner is also the Condorcet winner, and is not elected).

6.5 The Copeland Condition

We now consider one further extension of both the majority and Condorcet conditions, the Copeland condition. Recall that the Copeland score of a candidate is equal to the number of pairwise matchups it wins (also, the number of 1 in the candidate's row of the win-loss matrix, or the number of positive values in the candidate's row of the margin of victory matrix). A voting system meets the **Copeland condition** if it selects a candidate with the maximal Copeland score in all elections.

If there is a unique candidate with the maximal Copeland score, the method must select that candidate (the Copeland winner) as the election winner. If multiple candidates tie for the maximal Copeland score (as in the forthcoming example) the method need only select one of them (that is, it does not have to agree with the tiebreaking procedure of the Copeland method, or accept a shared victory).

The following example shows that it is indeed possible (though we'll see later on, rare) to have a tie for the maximal Copeland score:

Preference Ordering	Voters
$A \succ B \succ C$	9
$B \succ C \succ A$	9
$C \succ A \succ B$	9

This yields the following margin of victory matrix and win-loss matrix:

$$\begin{pmatrix} & A & B & C \\ A & 0 & 9 & -9 \\ B & -9 & 0 & 9 \\ C & 9 & -9 & 0 \end{pmatrix}$$

$$\begin{pmatrix} & A & B & C \\ A & 0 & 1 & -1 \\ B & -1 & 0 & 1 \\ C & 1 & -1 & 0 \end{pmatrix}$$

Finally, we note that satisfying the Copeland condition implies satisfying the Condorcet condition (the special case in which there is a candidate with maximal Copeland score), which in turn implies the majority condition (the special case of a Condorcet winner which is also a majority winner). Thus we have the following:

$$\text{Copeland} \Rightarrow \text{Condorcet} \Rightarrow \text{Majority} \tag{12}$$

$$!Majority \Rightarrow !Condorcet \Rightarrow !Copeland \tag{13}$$

6.5.1 Proofs

We now give proofs that the following methods satisfy the Copeland condition; Copeland.

Copeland:

By definition, the Copeland method selects a candidate with the maximal Copeland score, and thus necessarily satisfies the Copeland condition (as it should). In fact, it is the only method studied here which satisfies the Copeland condition.

6.5.2 Counterexamples

We now provide proofs/counterexamples that the following methods fail the Copeland condition; Plurality, Approval, Instant Runoff, Borda Count, Instant Runoff Borda Count, Least Worst Defeat, Instant Runoff Least Worst Defeat, Kemeny-Young, Schulze, and Ranked Pairs.

Plurality:

Because the Plurality method fails the Condorcet condition, it must fail the Copeland condition. The counterexample given for the Condorcet condition is also a counterexample for the Copeland condition (the Condorcet winner also has the highest Copeland score).

Approval:

Because the Approval method fails the Condorcet condition, it must fail the Copeland condition. The counterexample given for the Condorcet condition is also a counterexample for the Copeland condition (the Condorcet winner also has the highest Copeland score).

Instant Runoff:

Because the Instant Runoff method fails the Condorcet condition, it must fail the Copeland condition. The counterexample given for the Condorcet condition is also a counterexample for the Copeland condition (the Condorcet winner also has the highest Copeland score).

Borda Count:

Because the Borda Count fails the Condorcet condition, it must fail the Copeland condition. The counterexample given for the Condorcet condition is also a counterexample for the Copeland condition (the Condorcet winner also has the highest Copeland score).

Instant Runoff Borda Count:

Now, Instant Runoff Borda Count does satisfy the Condorcet condition, so we must look to elections without a Condorcet winner to find a counterexample. Consider the following election, represented by the votes on four candidates:

Preference Ordering	Number of Votes
$B \succ D \succ A \succ C$	12
$D \succ A \succ C \succ B$	9
$A \succ C \succ B \succ D$	7
$A \succ B \succ D \succ C$	2

The margin of victory matrix and Copeland scores are thus:

$$\begin{pmatrix} & A & B & C & D \\ A & 0 & 6 & 30 & -12 \\ B & -6 & 0 & -2 & 12 \\ C & -30 & 2 & 0 & -16 \\ D & 12 & -12 & 16 & 0 \end{pmatrix}$$

Candidate	Candidates Defeated	Copeland Score
A	B, C	2
B	D	1
C	B	1
D	A, C	2

Now, we compute the current Borda scores to determine the first eliminated candidate:

Candidate	Borda Count
A	87
B	77
C	53
D	83

It's significant that candidate C is eliminated in the first pass, since candidate A beats C on every vote. This means as we calculate the second round of preferences and then scores, candidate A will lose more points than the other candidates (enough, in fact, to cause elimination):

Preference Ordering	Number of Votes
$B \succ D \succ A$	12
$D \succ A \succ B$	9
$A \succ B \succ D$	9

Candidate	Borda Count
A	57
B	63
D	60

This reduces us to two candidates, B and D . Even though B only has one pairwise victory, it is the crucial one; B defeats D and thus wins the election (Borda score $51 - 39$). Thus, a candidate without the highest Copeland score is selected as the election winner, and Instant Runoff Borda Count fails the Copeland condition.

Least Worst Defeat:

Because Least Worst Defeat satisfies the Condorcet condition, it will not fail the Copeland condition whenever a Condorcet winner is present. We must then consider a situation in which there is not a Condorcet winner. Consider the following margin of victory matrix:

$$\begin{pmatrix} & A & B & C & D & E \\ A & 0 & 1 & 3 & 1 & -7 \\ B & -1 & 0 & -5 & -3 & 3 \\ C & -3 & 5 & 0 & 1 & -3 \\ D & -1 & 3 & -1 & 0 & 5 \\ E & 7 & -3 & 3 & -5 & 0 \end{pmatrix}$$

We then have the following Copeland scores and Least Worst Defeats:

Candidate	Copeland Score
A	3
B	1
C	2
D	2
E	2

Candidate	Least Worst Defeat
A	-7
B	-5
C	-3
D	-1
E	-3

Thus, we see that the candidate with the highest Copeland score (candidate A) loses in Least Worst Defeat; although it only has one defeat, that defeat is quite severe. Therefore, Least Worst Defeat fails the Copeland condition.

Instant Runoff Least Worst Defeat:

In fact, the same example used for Least Worst Defeat is a counterexample for Instant Runoff Least Worst Defeat (since candidate A 's only defeat is the worst defeat among all candidates, and is thus eliminated in the first pass. Therefore, Instant Runoff Least Worst Defeat also fails the Copeland condition.

Schulze:

While the Schulze method is a Condorcet method, it can fail to elect a candidate with the highest Copeland score if its path strengths are comparatively weak. Consider the following example election [3], given as a pairwise vote matrix:

$$\begin{pmatrix} & A & B & C & D & E \\ A & 0 & 18 & 11 & 21 & 21 \\ B & 12 & 0 & 14 & 17 & 19 \\ C & 19 & 16 & 0 & 10 & 10 \\ D & 9 & 13 & 20 & 0 & 30 \\ E & 9 & 11 & 20 & 0 & 0 \end{pmatrix}$$

The Copeland scores are then as follows:

Candidate	Candidates Defeated	Copeland Score
<i>A</i>	<i>B, D, E</i>	3
<i>B</i>	<i>D, E</i>	2
<i>C</i>	<i>A, B</i>	2
<i>D</i>	<i>C, E</i>	2
<i>E</i>	<i>C</i>	1

Now, we summarize the path strengths in the form of a matrix:

$$\begin{pmatrix} & A & B & C & D & E \\ A & 0 & 18 & 20 & 21 & 21 \\ B & 19 & 0 & 19 & 19 & 19 \\ C & 19 & 18 & 0 & 19 & 19 \\ D & 19 & 18 & 20 & 0 & 30 \\ E & 19 & 18 & 20 & 19 & 0 \end{pmatrix}$$

Thus, we see that candidate *B* is the only potential winner, since $p[B, X] > p[X, B]$ for all other candidates *X*. Thus *B* (not holding a maximal Copeland score) is selected as the election winner, and the Schulze method fails the Copeland condition.

Kemeny-Young:

Again, we must look beyond elections with Condorcet winners in order to find a Copeland counterexample for the Kemeny-Young method. Because the Kemeny-Young method develops scores based on full preference orderings, when a Condorcet winner is present, the preference ordering with the maximal score must select it as election winner (otherwise, a simple switch would find an ordering with a higher score). In elections without a Condorcet winner, however, some pairwise matchup implied by the preference ordering must be false. By making certain pairs more damaging to the score, we can find elections which elect a winner without maximal Copeland score. Consider the follow example, as a margin of victory matrix:

$$\begin{pmatrix} & A & B & C & D \\ A & 0 & 1 & 1 & -11 \\ B & -1 & 0 & 1 & -11 \\ C & -1 & -1 & 0 & 25 \\ D & 11 & 11 & -25 & 0 \end{pmatrix}$$

First, we check the Copeland scores of each candidate:

Candidate	Candidates Defeated	Copeland Score
<i>A</i>	<i>B, C</i>	2
<i>B</i>	<i>C</i>	1
<i>C</i>	<i>D</i>	1
<i>D</i>	<i>A, B</i>	2

Notice that the lone victory of candidate *C* is of the greatest magnitude (in fact, equal to the magnitude of the other five pairwise victories combined). This makes it weight very strongly in the Kemeny-Young scoring process. Further, *D* then beats candidate *A* and *B* very strongly as well (in comparison to the victories of *A* and *B*. This makes it very difficult to place *D* behind *A* or *B*. Since it was already very difficult to place *C* behind *D*, transitivity moves *C* to first place, despite its low Copeland score. Now that we've explained the result intuitively, let's back it up with the raw Kemeny-Young scores:

Preference Ordering	Kemeny-Young Score
<i>A</i> > <i>B</i> > <i>C</i> > <i>D</i>	6
<i>A</i> > <i>B</i> > <i>D</i> > <i>C</i>	-44
<i>A</i> > <i>C</i> > <i>B</i> > <i>D</i>	4
<i>A</i> > <i>C</i> > <i>D</i> > <i>B</i>	26
<i>A</i> > <i>D</i> > <i>B</i> > <i>C</i>	-22
<i>A</i> > <i>D</i> > <i>C</i> > <i>B</i>	-24
<i>B</i> > <i>A</i> > <i>C</i> > <i>D</i>	4
<i>B</i> > <i>A</i> > <i>D</i> > <i>C</i>	-46
<i>B</i> > <i>C</i> > <i>A</i> > <i>D</i>	2
<i>B</i> > <i>C</i> > <i>D</i> > <i>A</i>	24
<i>B</i> > <i>D</i> > <i>A</i> > <i>C</i>	-24
<i>B</i> > <i>D</i> > <i>C</i> > <i>A</i>	-26
<i>C</i> > <i>A</i> > <i>B</i> > <i>D</i>	2
<i>C</i> > <i>A</i> > <i>D</i> > <i>B</i>	24
<i>C</i> > <i>B</i> > <i>A</i> > <i>D</i>	0
<i>C</i> > <i>B</i> > <i>D</i> > <i>A</i>	22
<i>C</i> > <i>D</i> > <i>A</i> > <i>B</i>	46
<i>C</i> > <i>D</i> > <i>B</i> > <i>A</i>	44
<i>D</i> > <i>A</i> > <i>B</i> > <i>C</i>	0
<i>D</i> > <i>A</i> > <i>C</i> > <i>B</i>	-2
<i>D</i> > <i>B</i> > <i>A</i> > <i>C</i>	-2
<i>D</i> > <i>B</i> > <i>C</i> > <i>A</i>	-4
<i>D</i> > <i>C</i> > <i>A</i> > <i>B</i>	-4
<i>D</i> > <i>C</i> > <i>B</i> > <i>A</i>	-6

Thus, we see that the winning social preference ordering is $C \succ D \succ A \succ B$, which elects C as the winner, proving that the Kemeny-Young method fails the Copeland condition.

Ranked Pairs:

Although Ranked Pairs is a Condorcet method, it can fail to elect a candidate with the highest Copeland score if its victories are of comparatively small magnitude. Consider the following example election among five candidates, given as a margin of victory matrix:

$$\begin{pmatrix} & A & B & C & D & E \\ A & 0 & -13 & -3 & 25 & 7 \\ B & 13 & 0 & 19 & -21 & 11 \\ C & 3 & -19 & 0 & -1 & 9 \\ D & -25 & 21 & 1 & 0 & 5 \\ E & -7 & -11 & -9 & -5 & 0 \end{pmatrix}$$

Now, let's summarize the candidates defeated and the Copeland scores:

Candidate	Candidates Defeated	Copeland Score
A	D, E	2
B	A, C, E	3
C	A, E	2
D	B, C, E	3
E		0

Now, we rank the pairwise matchups by magnitude to apply the Ranked Pairs method. We put matchups that cause an intransitive cycle in boldface, and build the preference ordering as we work down the list;

Pairwise Matchup	Magnitude	Current Preference Ordering
$A \succ D$	25	$A \succ D$
$D \succ B$	21	$A \succ D \succ B$
$B \succ C$	19	$A \succ D \succ B \succ C$
$B \succ A$	13	$A \succ D \succ B \succ C$
$B \succ E$	11	$A \succ D \succ B \succ C, B \succ E$
$C \succ E$	9	$A \succ D \succ B \succ C \succ E$
$A \succ E$	7	$A \succ D \succ B \succ C \succ E$
$D \succ E$	5	$A \succ D \succ B \succ C \succ E$
$C \succ A$	3	$A \succ D \succ B \succ C \succ E$
$D \succ C$	1	$A \succ D \succ B \succ C \succ E$

Thus, we have that the final social preference ordering is $A \succ D \succ B \succ C \succ E$, which makes candidate A the election winner, but a candidate without the highest Copeland score in the election. Thus, Ranked Pairs fails the Copeland condition.

6.6 Monotonicity

The concept of monotonicity is straightforward, and its justification is intuitive. In simple terms, we demand that increasing a candidate's position on a ballot not decrease their position in the resultant social preference ordering. Changing our vote in a way which signifies a stronger preference for candidate A shouldn't decrease the chances of A winning the election, or hurt the final position of A in the social preference ordering. More formally;

A voting method f satisfies the **monotonicity condition** if for any election E , for all candidates C , increasing the position of C on any ballot in E (thus creating a new election \overline{E}) does not cause the position of C in $f(E)$ to worsen in $f(\overline{E})$.

I can not find a good justification to allow non-monotonic voting systems, as this would cause discord among the votership. The voters should not have to consider the mathematics of the system in expressing their preferences; this is likely to cause a misrepresentation of true preferences, whether intentional or unintentional.

6.6.1 Proofs

We now provide proofs that the following methods satisfy the monotonicity condition; Plurality, Approval, Borda Count, Least Worst Defeat, Kemeny-Young, Schulze, Ranked Pairs, and Copeland. Many of the proofs are relatively straightforward.

Plurality:

The plurality position is only dependent on the number of first-place votes a candidate receives. Given a candidate A , increasing the position of A on some ballots can not decrease the total number of first-place votes A gets, nor can it increase the number of first-place votes any other candidate X gets. Therefore, the social preference ordering position of candidate A can not decrease (since no other candidate can increase their total number of first-place votes), and the Plurality method satisfies the monotonicity condition.

Approval:

Similarly, the approval vote depends only on the number of ballots on which a candidate appears above the approval line. Increasing the position of a candidate A on some ballots can not decrease the number of approvals A receives, nor can it increase the number of approvals any other candidate X earns. Thus, A 's position in the social preference ordering can not decrease, and the approval method satisfies the monotonicity condition.

Borda Count:

Consider a candidate A with some Borda Count score. Increasing the position of A on any ballot strictly increases the Borda Count score of A , and strictly decreases the Borda Count score of some other candidate X . Therefore, A can not be passed by any other can-

didate, and its position in the social preference ordering can not decrease, so that the Borda Count satisfies the monotonicity condition.

Least Worst Defeat:

Again, consider a candidate A . Increasing its position on any number of ballots can not increase its pairwise defeat with any other candidate X , and further, it strictly increases the pairwise defeat corresponding to M_{AX} for some candidate X . Therefore, the worst defeat of A can not get worse, and the worst defeats of all other candidates X can not improve. This implies that A 's position in the social preference ordering can not decrease, and thus Least Worst Defeat satisfies the monotonicity condition.

Kemeny-Young:

Recall that the Kemeny-Young method ranks all possible social preference orderings by adding up the values of the margin of victory matrix corresponding to the pairwise defeats implied by the social ordering. Suppose we have a candidate A , and the social preference ordering with the highest Kemeny-Young score is;

$$X_1 \succ X_2 \succ \dots X_j \succ A \succ Y_1 \succ Y_2 \succ \dots \succ Y_k \quad (14)$$

Now, if we simply increase the position of A on some set of ballots (without rearranging any of the other candidates with respect to each other), only margin of victory matrix elements of the form M_{AZ} can change. The values which do not involve candidate A are left unchanged.

Because the above preference order has the maximal Kemeny-Young score, it has a higher Kemeny-Young score than any ordering which permutes the set (X_i, Y_l) , leaving A fixed. Representing the contribution of only the margin of victory matrix terms excluding candidate A as K , we can write the Kemeny-Young score of the social preference ordering as:

$$K + \sum_{i=1}^j M_{X_i A} + \sum_{l=0}^k M_{A Y_l} \quad (15)$$

After increasing the position of A on some ballots, the values in $M_{X_i A}$ are left unchanged or decrease, while the values in $M_{A Y_i}$ are left unchanged or increase. Because K is unchanged, and increasing values of $M_{A Y_i}$ only reinforces A being ranked above all Y_i in the preference ordering, we must only consider the change in values of $M_{X_i A}$.

If these values decrease, the only improvement we can make is to replace them with values $M_{A X_i}$, which actually implies an improvement in A 's position in the preference ordering. Thus, the Kemeny-Young method satisfies the monotonicity condition.

Schulze:

A full proof that the Schulze method is monotonic is presented in [1], Markus Schulze's recent account of the Schulze Method. The proof requires a different notation and formalism to frame the method than that developed here. Thus, we give only a sketch of the proof, and refer the reader to Schulze's paper for a more complete treatment.

Using the path heuristic of the Schulze method, we must only consider two possibilities; that our target candidate A is already a potential winner ($p[A, X] \geq p[X, A]$ for all other candidates X , or not. If not, the proof is trivial; a candidate can only improve from not being a potential winner. Now, if A is a potential winner, and we increase the position of A on ballots, we must show both that A remains a potential winner and that the number of potential winners does not increase (covering the interpretation in which the potential winners have an equal, random chance of being selected election winner).

Again, recall that the Schulze method is based on the strength of paths between two candidates (defined as the maximal weakest defeat among all paths between the two candidates). Suppose that we increase the position of candidate A , a potential winner, on some set of ballots. The pairwise defeats $[A, X]$ can not decrease. This implies that the weakest defeat along any path beginning at candidate A can not decrease, and the strengths of paths $p[A, X]$ all can not decrease.

Further, since the candidate A was moved up on ballots, without changing the relative position of any other candidates, the strengths of paths $p[X, A]$ can not increase. This means that any candidate who was not previously a potential winner can not become a potential winner after the vote change. Thus, the Schulze method is indeed monotonic.

Ranked Pairs:

Recall that the Ranked Pairs method operates by ranking the pairwise victories in descending order of magnitude. It then locks in each successive pairwise comparison, unless there is a conflict with an already locked pairwise matchup. Suppose we have a candidate A , and we increase the position of A of some set of ballots. This means that the magnitude of any pairwise victories of A can not decrease, and the magnitude of any pairwise losses can not increase. This means no pairwise defeat of A can rise on the list, and no pairwise victory of A can fall on the list. Therefore, candidate A can not fall on the preference ordering, which means that Ranked Pairs satisfies the monotonicity condition.

Copeland:

Recall that the Copeland method only depends on the number of pairwise matchups each candidate wins. Suppose we take a candidate A and increase its position on some set of ballots. This can not decrease the number of pairwise matchups it wins, and further, it can not increase the number of pairwise matchups any other candidate X wins. Therefore, the position of A in the social preference ordering can not decrease, and thus the Copeland method is monotonic.

6.6.2 Counterexamples

We now provide counterexamples to show that the following methods fail the monotonicity condition; Instant Runoff, Instant Runoff Borda Count, and Instant Runoff Least Worst De-

feat. Note that all methods which fail this condition eliminate candidates in succession, and the dependence on the order of elimination will cause them to fail the monotonicity condition.

Instant Runoff Voting:

As mentioned above, the order of elimination of candidates is important in instant runoff type voting systems. Increasing the position of a candidate on several ballots can change the order in which other candidates are eliminated, and this can effect the elimination of the assisted candidate. Consider the following example election, looking at the preference orderings and first-place votes:

Preference Ordering	Number of Votes
$A \succ B \succ C$	10
$B \succ A \succ C$	9
$C \succ A \succ B$	5
$C \succ B \succ A$	7

Candidate	First-Place Votes
A	10
B	9
C	12

We see that candidate B is eliminated first, leaving 19 votes for A and 12 votes for C , making A the election winner (the social preference ordering is $A \succ C \succ B$). Consider what happens if all of the voters who initially voted ($C \succ A \succ B$) change their votes to $A \succ C \succ B$, thus improving their preference of candidate A ;

Preference Ordering	Number of Votes
$A \succ B \succ C$	10
$A \succ C \succ B$	5
$B \succ A \succ C$	9
$C \succ A \succ B$	0
$C \succ B \succ A$	7

Candidate	First-Place Votes
A	15
B	9
C	7

The changed votes cause candidate C to be eliminated first, setting up a head-to-head competition between A and B , which A loses! Thus, candidate A goes from first to second in the social preference ordering after receiving additional support, violating the monotonicity condition.

Instant Runoff Borda Count:

We follow a similar line of thought in constructing a counterexample for the Instant Runoff Borda Count. We begin with a set of preference orderings which setup a favorable second round head-to-head matchup for our target candidate, and then disrupt this matchup by increasing the votes of the target candidate. Consider the following set of preferences and the associated Borda Count scores:

Preference Ordering	Number of Votes
$A \succ C \succ B$	9
$B \succ A \succ C$	1
$B \succ C \succ A$	6
$C \succ B \succ A$	3

Candidate	Borda Count
A	38
B	36
C	40

The elimination of candidate B then leads to a head-to-head matchup between A and C , which A wins (10 votes to 9), producing a social preference ordering of $A \succ C \succ B$. Now, suppose that all those who voted $B \succ C \succ A$ change their votes to $B \succ A \succ C$, increasing their opinion of candidate A ;

Preference Ordering	Number of Votes
$A \succ C \succ B$	9
$B \succ A \succ C$	7
$B \succ C \succ A$	0
$C \succ B \succ A$	3

Candidate	Borda Count
A	44
B	36
C	34

This time, candidate C is eliminated, instigating a head-to-head matchup between A and B , which A loses (again 10 votes to 9), producing a social preference ordering of $B \succ A \succ C$. Thus, Instant Runoff Borda Count also fails the monotonicity condition.

Instant Runoff Least Worst Defeat:

Finally, we consider Instant Runoff Least Worst Defeat, which functions in a similar way to the previous two counterexamples. Consider the following margin of victory matrix and worst defeats:

$$\begin{pmatrix} & A & B & C \\ A & 0 & 6 & -5 \\ B & -6 & 0 & 5 \\ C & 5 & -5 & 0 \end{pmatrix}$$

Candidate	Worst Defeat
<i>A</i>	-5
<i>B</i>	-6
<i>C</i>	-5

Thus, candidate *B* is eliminated, making candidate *C* the election winner (as $M_{CA} = 5$). Now, consider the following modified margin of victory matrix, in which *C* wins more pairwise matchups against *A* (signifying an swapping of *A* and *C* on ballots which ranked *A* above *C*);

$$\begin{pmatrix} & A & B & C \\ A & 0 & 6 & -7 \\ B & -6 & 0 & 5 \\ C & 7 & -5 & 0 \end{pmatrix}$$

Candidate	Worst Defeat
<i>A</i>	-7
<i>B</i>	-6
<i>C</i>	-5

The changed ballots cause candidate *A* to be eliminated first, setting up a head-to-head matchup between *B* and *C*, which *B* wins ($M_{BC} = 5$), causing a drop in *C*'s position in the preference order, violating monotonicity.

6.7 Clone Invariance

A set of candidates C_i are **clones** if all voters rank them equivalently compared to all other candidates. This means that for every voter, the preference ordering must not rank any other candidate between two members of C_i . For example, the preference ordering

$$A \succ D \succ C_2 \succ C_4 \succ C_1 \succ C_3 \succ B \tag{16}$$

treats the C_i as clones, while the preference ordering

$$C_3 \succ C_4 \succ A \succ C_1 \succ C_2 \succ D \succ B \tag{17}$$

does not treat the C_i as clones. If all voters treat a set of candidates C_i as clones, then they are, in fact, clones.

Clones are significant from a practical and political perspective. Suppose that the results of

a voting system could be altered by introducing clones of a particular candidate (politically, this might mean introducing candidates with very similar platforms, for example). Formally, a voting method satisfies the **clone invariance condition** if a candidate C clones itself to a set C_i , then the chance of C winning the election is equal to the chance of any of the C_i winning the election.

This is important, as if cloning candidates provided some advantage, politicians would actively seek clones and have them run, creating insincere slates of candidates (which surely disrupts the voters' abilities to vote their true preferences amongst the legitimate candidates). On the other hand, if cloning candidates provides a disadvantage, then candidate who take common (or reasonable, or productive) stances could be unfairly penalized. If both can occur in different situations, we have a mess of confusion.

Before we consider which methods satisfy this condition, we must distinguish between two types of clones. We define a set of candidates C_i to be **internal clones** if the margin of victory matrix amongst the C_i is the zero matrix. This means that, with respect to each other, the voters have no preference on the clones. If then, the clones would tie as the election winner (for example, if they all tie for the highest Copeland score), we can imagine one of them is selected at random as the winner. Any set of clones which are not internal clones will be called **external clones**. We then consider two types of clone invariance; clone invariance is defined as above, with no restrictions on the clones. **Internal clone invariance** will, expectedly, consider only internal clones. Obviously, clone invariance implies internal clone invariance (which can thus be considered a weaker form of clone invariance).

6.7.1 Proofs

We now provide proofs that the following methods satisfy clone invariance; Approval, Instant Runoff, Instant Runoff Least Worst Defeat, Schulze, and Ranked Pairs. We also provide proofs that the following methods satisfy internal clone invariance (but not general clone invariance); Least Worst Defeat

Approval:

Suppose we have an election in which voters provide approval lines. On any particular ballot, if a candidate C is approved, then if C is replaced by a set of clones C_i , all of the clones must be approved. Similarly, if C is not approved, then none of the clones can be approved. This means that the approval scores of each of the clones is equal to the original approval score of C . Of course, the approval scores of all the other candidates are left unchanged. Therefore, the clones (collectively) will occupy the same spot in the approval social preference ordering as did the original candidate C . This means that if C was the approval winner, one of the clones will be, and if not, none of the clones can be. Thus, the approval method satisfies clone invariance (and therefore satisfies internal clone invariance).

Instant Runoff:

In the process of instant runoff voting, the candidate with the fewest number of first place votes is eliminated during each round. When a candidate is eliminated, its votes are redistributed to the next candidate on the ballot.

First, we note that while, as in the plurality case, if a candidate C is replaced by a set of clones C_i , the first-place votes of C are split amongst the clones. This may indeed cause *some* of the clones to be eliminated earlier than C was in the original election. However, because the C_i are clones, as long as the eliminated candidate is not the last member of the set, all of the first-place votes are redistributed to other clones. This implies two things; when only a single clone remains, it will have all of the first place votes the original candidate C did at that point, and no other candidate can gain first-place votes from the clones until all of them are eliminated.

These means that we can consider the election involving clones in the following way; each round ends when a non-clone is eliminated, or when the last clone is eliminated. By the above observations, this election must proceed in the same order as the original instant runoff election. Thus, the instant runoff method satisfies clone invariance (and therefore satisfies internal clone invariance).

Least Worst Defeat:

Suppose we have a candidate C replaced by a set of internal clones, $\{C_1, \dots, C_k\}$. Because the clones are internal, the C_i block of the margin of victory matrix is the k by k zero matrix. Further, for every other candidate X , the margin of victory matrix entry $M_{XC_i} = M_{XC}$ for all i . The implication of this is that the worst defeat of all candidates remains unchanged (though the frequency of it may be repeated). Further, the worst defeat of all of the clones is equal to the worst defeat of the original candidate C . Thus, Least Worst Defeat satisfies internal clone invariance.

Instant Runoff Least Worst Defeat:

Just as with Instant Runoff Voting, the introduction of clones does not change the value of the elimination condition for any of the candidates.

Suppose we have an election with some candidate C replaced by an arbitrary set of clones C_i . Now, if a candidate X had a defeat of M_{XC} against C , then for each C_i it has the same defeat $M_{XC_i} = M_{XC}$. Thus, the introduction of clones can not alter the worst defeat of any candidate. Further, if a clone is eliminated, the worst defeats of a non-clone candidate do not change unless that clone was the last clone, which is equivalent to the original candidate being eliminated.

Further, if the clones are not internal clones, then the worst defeat of C is greater than or equal to the worst defeat for all C_i , as the internal defeats could be worse than the original worst defeat of C (we'll see that this causes standard Least Worst Defeat to fail general clone invariance, but not internal clone invariance). However, once a single clone is left, it's worst defeat must be equal to the worst defeat of the original candidate C . This, combined

with the fact that the worst defeats of all other candidates remains constant, implies that the last clone eliminated will be eliminated in the position that the original candidate C was (with respect to the non-clone candidates). Thus, Instant Runoff Least Worst Defeat satisfies clone invariance (and therefore satisfies internal clone invariance).

Schulze:

Again, for a full proof of clone invariance, we refer the reader to [1], Markus Schulze's exposition on the Schulze method, which establishes a more intricate formalism for studying the Schulze method. We give a sketch of the proof here.

Suppose we have a candidate A replaced by a set of clones A_i . If the candidate A was not a potential winner before, then none of the clones can be potential winners (as there must exist a candidate X such that $p[A_i, X] < p[X, A_i]$). If A was a potential winner, we show that at least one of the A_i must be a potential winner.

Because the candidates are clones, since A is a potential winner, each A_i satisfies $p[A_i, X] \geq p[X, A_i]$ for all X not in the set of clones. Further, we can consider the Schulze method as applied to the set of clones. By definition, there must be at least one potential winner. Therefore, at least one clone must be an overall potential winner, since it satisfies $p[A_i, X] \geq p[X, A_i]$ for all X not equal to A_i . Thus, the Schulze method is indeed clone invariant (and thus internal clone invariant).

Ranked Pairs:

Suppose we have a candidate C which is replaced by an arbitrary set of clones C_i . First, the pairwise defeats of other candidates against C must be equal to their pairwise defeats against each of the C_i . This means that where a pairwise matchup (such as $C \succ X$ or $X \succ C$) would be listed in the ranked pairs count, we simultaneously include either $C_i \succ X$ or $X \succ C_i$ for all i .

Our other concern is the relationships between the C_i themselves. Consider the results of ranked pairs run on the C_i themselves. Because each of the C_i behave identically with respect to all other candidates, the results of the original ranked pairs election and of the C_i election are independent. Thus, the victor in the C_i election will be the highest ranked clone, since no other candidate can selectively negate the ranked pairs validity of any clone candidate.

This means that the entire set of clone candidates collectively finish in the same position that the original candidate C did, and thus, Ranked Pairs satisfies clone invariance (and also internal clone invariance).

6.7.2 Counterexamples

We now provide counterexamples that the following voting methods fail all clone invariance; Plurality, Borda Count, Instant Runoff Borda Count, Kemeny-Young, Copeland. We also provide counterexamples to show that the following methods fail clone invariance (though they satisfy internal clone invariance); Least Worst Defeat. Recall that if a counterexample is found using internal clones, this proves that the method fails all clone invariance.

Plurality:

When candidates generate clones, the number of first-place votes the candidate originally received may be split amongst the clones, causing a decrease in plurality score. Consider the following election, given the preference orderings and the first-place votes:

Preference Order	Number of Votes
$A \succ B \succ C$	15
$B \succ A \succ C$	9
$C \succ B \succ A$	6

Candidate	First-Place Votes
A	15
<i>B</i>	9
<i>C</i>	6

Clearly, candidate *A* is elected the plurality winner. Suppose, now, that candidate *A* is replaced by the set of internal clones $\{A_1, A_2, A_3\}$. The votes then stand as follows:

Preference Order	Number of Votes
$A_1 \succ A_2 \succ A_3 \succ B \succ C$	5
$A_2 \succ A_3 \succ A_1 \succ B \succ C$	5
$A_3 \succ A_1 \succ A_2 \succ B \succ C$	5
$B \succ A_1 \succ A_2 \succ A_3 \succ C$	3
$B \succ A_2 \succ A_3 \succ A_1 \succ C$	3
$B \succ A_3 \succ A_1 \succ A_2 \succ C$	3
$C \succ B \succ A_1 \succ A_2 \succ A_3$	2
$C \succ B \succ A_2 \succ A_3 \succ A_1$	2
$C \succ B \succ A_3 \succ A_1 \succ A_2$	2

Candidate	First-Place Votes
A_1	5
A_2	5
A_3	5
B	9
<i>C</i>	6

Because the clones A_i have split their first-place votes three ways, candidate B now has the most first-place votes, and wins the plurality election. Thus, the plurality method fails internal clone invariance, and therefore fails clone invariance.

Borda Count:

Constructing a counterexample involving external clones is relatively straightforward for the Borda Count; consider the following example;

Preference Ordering	Number of Votes
$A \succ B \succ C$	10
$B \succ A \succ C$	6

Candidate	Borda Count
A	42
B	38
C	16

And thus, we see that candidate A wins the election. However, if B introduces a clone, which it beats on all ballots (perhaps by notifying supporters and instructing them how to vote), B can take control of the election;

Preference Ordering	Number of Votes
$A \succ B_1 \succ B_2 \succ C$	10
$B_1 \succ B_2 \succ A \succ C$	6

Candidate	Borda Count
A	52
B_1	54
B_2	36
C	16

By leveraging the patsy clone, B_1 now wins the election, violating clone invariance.

However, it is not as straightforward to find a counterexample involving internal clones, since they can not claim as many defeats over their own clones. It helps, for our analysis, to note that a candidate's Borda Count is equal to the number of voters (since even last place scores one point) plus the number of candidates defeated (summed over all ballots). We'll consider the change in the number of candidates defeated following the introduction of clones.

Suppose that in an initial election with n voters, candidate A has a higher Borda Count than candidate B . Candidate B then tries to control the election by introducing internal clones, so that B is replaced by $\{B_1, \dots, B_k\}$. By the required symmetry present with internal clones, each clone gains $\frac{n(k-1)}{2}$ defeats, and thus adds this many Borda Count points. Now, suppose that candidate A defeats B on x ballots (note that $x > \frac{n}{m}$, where m is the number

of candidates, since A receives a higher total Borda score than B). Then, candidate A gains $x(k - 1)$ defeats, and adds this many Borda Count points.

Now, for a reversal in preference ordering to occur (that is, for all of the B_i , which have the same score, to overtake candidate A), we then require

$$\frac{n(k - 1)}{2} > x(k - 1) + (BC(A) - BC(B)) \quad (18)$$

where $BC(A)$ and $BC(B)$ are the old Borda scores of A and B . Solving for x , we then have the following inequality;

$$\frac{n}{m} < x < \frac{n}{2} - \frac{BC(A) - BC(B)}{k - 1} \quad (19)$$

But further, as the difference between $BC(A)$ and $BC(B)$ increases, so to must the lower bound on x (as the $\frac{n}{m}$ gives the smallest possible victory). Thus, increasing the initial victory of A constricts the inequality on both ends, severely decreasing the number of available counterexamples. We now provide such a counterexample;

Preference Order	Number of Votes
$A \succ C \succ B$	8
$B \succ A \succ C$	14

We see that $m = 3$ and $n = 22$. Rounding to integers, we see that;

$$8 \leq x < 11 - \frac{BC(A) - BC(B)}{k - 1} \quad (20)$$

Thus, if we use a set of two clones, the difference in initial Borda Scores must be less than or equal to two;

Candidate	Borda Count
A	52
B	50
C	30

We see that candidate A wins, and by exactly two points. Thus, introducing a set of internal clones should cause a reversal of A and B in the social preference ordering;

Preference Order	Number of Votes
$A \succ C \succ B_1 \succ B_2$	4
$A \succ C \succ B_2 \succ B_1$	4
$B_1 \succ B_2 \succ A \succ C$	7
$B_2 \succ B_1 \succ A \succ C$	7

Candidate	Borda Count
A	60
B_1	61
B_2	61
C	38

As expected, both clone candidates tie for the highest Borda Score, and thus one will be selected as the election winner. This shows that the Borda Count indeed fails internal clone invariance.

Instant Runoff Borda Count:

Unlike Instant Runoff voting, in which first place votes could not transfer until all clones were eliminated, in Instant Runoff Borda Count, the presence of clones can change the order in which candidates are eliminated. Of course, this can cause a change in the overall results. Recall that the Borda Count is equivalent to summing the rows of the margin of victory matrix. Thus, consider the following margin of victory matrix and the sum of the rows;

$$\begin{pmatrix} & A & B & C \\ A & 0 & 2 & -3 \\ B & -2 & 0 & 6 \\ C & 3 & -6 & 0 \end{pmatrix}$$

Candidate	Row Total
<i>A</i>	-1
<i>B</i>	4
<i>C</i>	-3

Thus, candidate *C* is eliminated, setting up a head-to-head matchup between *A* and *B*, which candidate *A* wins, making *A* the election winner. Now, suppose *A* is replaced by a set of four internal clones A_i ;

$$\begin{pmatrix} & A_1 & A_2 & A_3 & A_4 & B & C \\ A_1 & 0 & 0 & 0 & 0 & 2 & -3 \\ A_2 & 0 & 0 & 0 & 0 & 2 & -3 \\ A_3 & 0 & 0 & 0 & 0 & 2 & -3 \\ A_4 & 0 & 0 & 0 & 0 & 2 & -3 \\ B & -2 & -2 & -2 & -2 & 0 & 6 \\ C & 3 & 3 & 3 & 3 & -6 & 0 \end{pmatrix}$$

Candidate	Row Total
A_1	-1
A_2	-1
A_3	-1
A_4	-1
<i>B</i>	-2
<i>C</i>	6

By cloning the candidate that *C* originally defeated, and *B* originally lost to, we are able to eliminate *B* in the first round. This sets up a situation in which *C* is the Condorcet

winner of the remaining candidates, and thus C is the election winner, rather than one of the clones. Thus, Instant Runoff Borda Count fails the internal clone invariance condition, and therefore clone invariance in general.

Least Worst Defeat:

While Least Worst Defeat does satisfy internal clone invariance, it does not satisfy clone invariance in general. While internal clones have a zero margin of victory matrix, this is not true for clones in general. Consider the following election, as a margin of victory matrix;

$$\begin{pmatrix} & A & B & C \\ A & 0 & 5 & 7 \\ B & -5 & 0 & 3 \\ C & -7 & -3 & 0 \end{pmatrix}$$

Clearly, since A is a Condorcet winner, it has no defeats, and thus is the Least Worst Defeat winner. Suppose now that A is replaced by a set of clones, each of which claims one victory and one defeat over the other clones:

$$\begin{pmatrix} & A_1 & A_2 & A_3 & B & C \\ A_1 & 0 & 11 & -17 & 5 & 7 \\ A_2 & -11 & 0 & 13 & 5 & 7 \\ A_3 & 17 & -13 & 0 & 5 & 7 \\ B & -5 & -5 & -5 & 0 & 3 \\ C & -7 & -7 & -7 & -3 & 0 \end{pmatrix}$$

Because each of the similar candidates loses heavily to another, they have the three worst defeats, and thus none of them is selected election winner (previous loser candidate B is now the winner). Thus, Least Worst Defeat does not satisfy general clone invariance.

Kemeny-Young:

Consider the following election, in which the introduction of a set of clones changes the winner (and improves the cloned candidate's position in the social preference ordering). We give the margin of victory matrix, and associated Kemeny-Young scores;

$$\begin{pmatrix} & A & B & C \\ A & 0 & 2 & -5 \\ B & -2 & 0 & 3 \\ C & 5 & -3 & 0 \end{pmatrix}$$

Preference Order	Kemeny-Young Score
$A \succ B \succ C$	0
$A \succ C \succ B$	-6
$B \succ A \succ C$	-4
$B \succ C \succ A$	6
$C \succ A \succ B$	4
$C \succ B \succ A$	0

Thus, we see that candidate B is the winner, and that candidate A is selected last in the social preference ordering. Consider what happens now, if A is replaced by a set of two internal clones, A_1 and A_2 . The margin of victory matrix is then:

$$\begin{pmatrix} & A_1 & A_2 & B & C \\ A_1 & 0 & 0 & 2 & -5 \\ A_2 & 0 & 0 & 2 & -5 \\ B & -2 & -2 & 0 & 3 \\ C & 5 & 5 & -3 & 0 \end{pmatrix}$$

We write the Kemeny-Young scores without subscripts on the A candidates, since permuting A_1 and A_2 has no effect on the Kemeny-Young score:

Preference Order	Kemeny-Young Score
$A \succ A \succ B \succ C$	-3
$A \succ B \succ A \succ C$	-7
$A \succ B \succ C \succ A$	3
$B \succ A \succ A \succ C$	-11
$B \succ A \succ C \succ A$	-1
$B \succ C \succ A \succ A$	9
$A \succ A \succ C \succ B$	-9
$A \succ C \succ A \succ B$	1
$A \succ C \succ B \succ A$	-3
$C \succ A \succ B \succ A$	7
$C \succ A \succ A \succ B$	11
$C \succ B \succ A \succ A$	3

And thus we see that candidate C is now the winner, and that the previous winner B has fallen to last place in the social preference ordering. Further, candidate A has improved their position in the preference order by being replaced by a set of clones, and therefore the Kemeny-Young method fails internal clone invariance (and general clone invariance).

Copeland:

Because internal clones introduce a margin of victory matrix full of ties, internal clones can be at a significant disadvantage in a Copeland election, which is predicated on number of pairwise victories. Consider the following win-loss matrix, and associated Copeland scores;

$$\begin{pmatrix} & A & B & C & D & E \\ A & 0 & -1 & 1 & 1 & 1 \\ B & 1 & 0 & 1 & -1 & -1 \\ C & -1 & -1 & 0 & 1 & 1 \\ D & -1 & 1 & -1 & 0 & 1 \\ E & -1 & 1 & -1 & -1 & 0 \end{pmatrix}$$

Candidate	Candidates Defeated	Copeland Score
A	C,D,E	3
<i>B</i>	<i>A,C</i>	2
<i>C</i>	<i>D,E</i>	2
<i>D</i>	<i>B,E</i>	2
<i>E</i>	<i>B</i>	1

Thus, we see that candidate *A* is the Copeland election winner. Now, suppose *A* is replaced by a set of three internal clones. The results then change:

$$\begin{pmatrix} & A_1 & A_2 & A_3 & B & C & D & E \\ A_1 & 0 & 0 & 0 & -1 & 1 & 1 & 1 \\ A_2 & 0 & 0 & 0 & -1 & 1 & 1 & 1 \\ A_3 & 0 & 0 & 0 & -1 & 1 & 1 & 1 \\ B & 1 & 1 & 1 & 0 & 1 & -1 & -1 \\ C & -1 & -1 & -1 & -1 & 0 & 1 & 1 \\ D & -1 & -1 & -1 & 1 & -1 & 0 & 1 \\ E & -1 & -1 & -1 & 1 & -1 & -1 & 0 \end{pmatrix}$$

Candidate	Candidates Defeated	Copeland Score
<i>A</i> ₁	<i>C,D,E</i>	3
<i>A</i> ₂	<i>C,D,E</i>	3
<i>A</i> ₃	<i>C,D,E</i>	3
B	A₁, A₂, A₃, C	4
<i>C</i>	<i>D,E</i>	2
<i>D</i>	<i>B,E</i>	2
<i>E</i>	<i>B</i>	1

We can see that since *B* defeated *A* in the original election, it gains two points in the Copeland score, since it beats all clones of *A*. On the other hand, the clones tie each other, and thus gain no defeats. Since *B* is now the election winner, we see that Copeland fails internal clone invariance, and thus clone invariance in general.

6.8 Loser Independence

The concept of loser independence is relatively simple (and similar to another common condition, the independence of irrelevant alternatives). Intuitively, given an election, with election

winner C , if a losing candidate were to drop out (if it is convenient, imagine that all candidates can accurately predict the results of the election), this should not change the result of the election. Formally, if the removal of any losing candidate causes the election winner to change in any election, the voting method fails the **loser independence condition**. Note that, equivalently, adding another outcome (which does not become the election winner) should not change the election outcome.

This is important, since if a candidate knows that they will lose, they may be tempted to drop out (via a bribe) to benefit another candidate who would not win if they stayed in the election. Allowing this type of corruption certainly seems like a bad idea. Unfortunately, every method we've discussed so far (except the Approval method, which has other weaknesses) fails this condition.

It turns out that (requiring a deterministic outcome) the Condorcet and loser independence conditions are incompatible. This is very similar to the famous Arrow's Paradox. [6], which incited much of the study of preferential ballot voting systems (since it was the first widely published work to show that two "vital" conditions on a voting system were incompatible, leading many to claim that no perfect voting system exists).

The proof of this is relatively simple. Suppose a voting method f satisfies both the Condorcet condition and the loser independence condition. Now consider the following election, represented by a margin of victory matrix, where x, y , and z are positive integers.

$$\begin{pmatrix} & A & B & C \\ A & 0 & x & -z \\ B & -x & 0 & y \\ C & z & -y & 0 \end{pmatrix}$$

Now, since there is no Condorcet winner, without loss of generality, suppose that $f(E) = A$ (that is, the election winner is A). By loser independence, if B drops out of the election, A must remain the winner, but since the remaining margin of victory matrix is;

$$\begin{pmatrix} & A & C \\ A & 0 & -z \\ C & z & 0 \end{pmatrix}$$

By the Condorcet condition C must be the election winner, a contradiction. Thus, no Condorcet method can also satisfy the loser independence condition.

6.8.1 Proofs

Only one method discussed here satisfies the condition of loser independence: Approval Voting, given a caveat. We must assume that the removal of a candidate C from the race does not change the voter's set of approved candidates (other than, of course, removing C). In this case, the number of approval votes for each candidate is unchanged, and since C

was not the approval winner, the approval winner is therefore unchanged. This means that approval voting indeed satisfies the condition of loser independence.

6.8.2 Counterexamples

We now provide counterexamples demonstrating that the non-Condorcet methods (except Approval voting) we have discussed fail the condition of loser independence. Recall from above that all deterministic Condorcet methods fail the loser independence condition. The counterexamples have a similar spirit to the counterexamples used for the monotonicity condition; we eliminate a candidate that creates an unfavorable head-to-head matchup for the winner (note that while Condorcet methods fail loser independence, a Condorcet winner can not typically be ousted by the removal of a losing candidate).

Plurality:

Consider the following election, as a set of ballots and the corresponding number of first-place votes;

Preference Order	Number of Votes
$A \succ B \succ C$	7
$B \succ A \succ C$	4
$C \succ B \succ A$	6

Candidate	First-Place Votes
A	7
B	4
C	6

Thus, candidate A is the plurality winner. Now, suppose that candidate C is removed from the race. The preferences and first-place votes are then recalculated for a simple majority election:

Candidate	First-Place Votes
A	7
B	10

And now, candidate B is the winner, since in all votes for candidate C , B was the second choice. Thus, plurality fails the loser independence condition.

Instant Runoff Voting:

Consider the following election, as a set of ballots and the corresponding first-place votes;

Preference Ordering	Number of Votes
$A \succ B \succ C \succ D$	5
$B \succ D \succ A \succ C$	4
$C \succ D \succ B \succ A$	6

Candidate	First-Place Votes
A	5
B	4
C	6
D	0

In this case, candidate D is eliminated, leading to the next round of Instant Runoff Voting. The final two rounds progress as follows:

Preference Ordering	Number of Votes
$A \succ B \succ C$	5
$B \succ A \succ C$	4
$C \succ B \succ A$	6

Candidate	First-Place Votes
A	5
B	4
C	6

Candidate	(First-Place) Votes
A	9
C	6

Thus, candidate A is the election winner. Now, consider how the election would change if losing candidate C were removed from the race. This would produce a three candidate Instant Runoff election involving A , B , and D , which proceeds as follows:

Preference Ordering	Number of Votes
$A \succ B \succ D$	5
$B \succ D \succ A$	4
$D \succ B \succ A$	6

Candidate	First-Place Votes
A	5
B	4
D	6

Candidate	(First-Place) Votes
A	5
D	10

And we see that candidate D (the last place candidate when C remained in the election) is now the election winner, ousting A . Thus, Instant Runoff Voting fails the loser independence condition.

Borda Count:

We can use an election essentially equivalent to the counterexample for the Plurality method as a counterexample for the Borda Count. Consider the following set of ballots and the associated Borda Count scores;

Preference Ordering	Number of Votes
$A \succ B \succ C$	3
$A \succ C \succ B$	4
$B \succ A \succ C$	4
$C \succ B \succ A$	6

Candidate	Borda Count
A	35
B	34
C	33

Thus, we see that candidate A is selected as the election winner. However, if losing candidate C drops out, we have the following head-to-head matchup;

Candidate	Number of Votes
A	7
B	10

And we see that candidate B now wins the election, showing that the Borda Count also fails the loser independence condition.

6.9 Summary Table

We now summarize which methods satisfy each of the conditions we've discussed thus far, as a reference for decision making.

Voting Method	Majority	Condorcet	Copeland	Monotonicity
Plurality	Yes	No	No	Yes
Approval	No	No	No	Yes
Instant Runoff	Yes	No	No	No
Borda Count	No	No	No	Yes
Instant Runoff Borda Count	Yes	Yes	No	No
Least Worst Defeat	Yes	Yes	No	Yes
Instant Runoff Least Worst Defeat	Yes	Yes	No	No
Kemeny-Young	Yes	Yes	No	Yes
Schulze	Yes	Yes	No	Yes
Ranked Pairs	Yes	Yes	No	Yes
Copeland	Yes	Yes	Yes	Yes

Voting Method	Internal Clone Inv.	Clone Inv.	Loser Independence
Plurality	No	No	No
Approval	Yes	Yes	Yes
Instant Runoff	Yes	Yes	No
Borda Count	No	No	No
Instant Runoff Borda Count	No	No	No
Least Worst Defeat	Yes	No	No
Instant Runoff Least Worst Defeat	Yes	Yes	No
Kemeny-Young	No	No	No
Schulze	Yes	Yes	No
Ranked Pairs	Yes	Yes	No
Copeland	No	No	No

6.10 Conditions Philosophy

We make a brief philosophical aside at this juncture. As we define a list of reasonable conditions for voting systems to follow, it may be tempting to pick a favorite system, and then argue the merits of the conditions which it satisfies. Instead, we recommend beginning with conditions, choosing a set of them based on their merit, then deriving what systems in fact hold all the conditions (if any). Especially while tabling practical considerations, it is best to begin by considering what properties are most desirable. As we develop probabilistic measures of some of these conditions, we will have to reevaluate this approach, because we will not only have to compare the merits of different conditions, but quantify these differences in order to make rational decisions. We'll revisit this concept in the conclusions section. First, we need to compile more information.

7 Simulation of Measurable Conditions

7.1 Converting Binary Conditions into Continuously Measurable Conditions

As we stated at the beginning of the paper, one goal is to refine how conditions of voting systems are measured. Just as preferential ballot voting systems utilize more information than single vote systems; continuous measures are more illuminating than binary ones. An increase in information is certainly valuable to those making decisions concerning the merits of our various preferential ballot voting systems. Our job, then, is to determine the appropriate extensions of existing conditions which accurately quantify the extent to which voting systems satisfy those conditions. We note that this quantization is philosophically distinct from quantifying the relative importance of the different conditions; to make rational decisions we need some measure of both, and unfortunately the latter is necessarily a philosophical question.

7.2 Condorcet and Copeland Ratios

As we've analyzed above, Condorcet winners do not occur in every election, and in an election with m candidates, the maximal Copeland score is not always $m - 1$. We must take this into account when measuring the effectiveness of voting systems with respect to the Condorcet and Copeland conditions.

Before, we determined that a method is a Condorcet method if it selects the Condorcet winner whenever it exists. While the probability of the existence of Condorcet winners has been well studied [7], the probability of non-Condorcet methods selecting the Condorcet winner (when it exists) has not been significantly covered. In this spirit, we define the **Condorcet Ratio** to be the fraction of possible elections with Condorcet winners in which a method actually selects the Condorcet winner. Because this fraction is dependent both on the number of candidates and voters, we denote it $C_{ratio}(m, n)$.

Similarly, we can not expect a method to select a winner with a Copeland score of $m - 1$ when the maximal score in a particular election is $m - 3$. Thus, we define the **Copeland Ratio** as the average value of the Copeland score of the alternative selected by a method divided by the maximal Copeland score (in each particular election). For example, a method which picked a winner with Copeland score 4 when there was a candidate with Copeland score 5 would have a 0.8 ratio for that particular election. The Copeland Ratio is then the average of these ratios over all possible elections. Again, because of the dependence on m , the number of candidates, and n , the number of voters, we denote it $CP_{ratio}(m, n)$.

To measure these ratios, we run a large number of random elections through all of the voting methods. Intuitively, we uniformly select a random preference order from the set of $m!$ possible preference orderings for each voter. Thus, we approximate a uniform random distribution over the set of possible votes. The rationale here is that random votes represent a worst case scenario, both in the proportion of Condorcet winners present, and in unearthing examples which test the framework of the various methods. As we saw in the

sections above, many of the counterexamples to voting conditions are quite contrived, and thus may not appear in many actual elections. Using a large sample of random elections helps us to avoid that bias, as well as set a baseline for how the methods perform in actual election situations.

The following results correspond to random elections run with $n = 1001$ voters, for the indicated number of candidates and indicated election sample size:

3 candidates, 1001 voters, 100000 elections:

Voting Method	Condorcet Ratio	Copeland Ratio
Plurality	0.764	0.871
Approval	0.599	0.755
Instant Runoff	0.944	0.961
Borda Count	0.898	0.951
Copeland	1	1
IRLWD	1	1
IRBC	1	1
Kemeny-Young	1	1
Least Worst Defeat	1	1
Ranked Pairs	1	1
Schulze	1	1

4 candidates, 1001 voters, 50000 elections:

Voting Method	Condorcet Ratio	Copeland Ratio
Plurality	0.653	0.845
Approval	0.547	0.778
Instant Runoff	0.877	0.911
Borda Count	0.870	0.957
Copeland	1	1
IRLWD	1	0.993
IRBC	1	0.994
Kemeny-Young	1	0.997
Least Worst Defeat	1	0.994
Ranked Pairs	1	0.995
Schulze	1	0.995

5 candidates, 1001 voters, 20000 elections:

Voting Method	Condorcet Ratio	Copeland Ratio
Plurality	0.567	0.823
Approval	0.517	0.794
Instant Runoff	0.802	0.849
Borda Count	0.852	0.961
Copeland	1	1
IRLWD	1	0.984
IRBC	1	0.987
Kemeny-Young	1	0.992
Least Worst Defeat	1	0.987
Ranked Pairs	1	0.990
Schulze	1	0.989

6 candidates, 1001 voters, 15000 elections:

Voting Method	Condorcet Ratio	Copeland Ratio
Plurality	0.506	0.809
Approval	0.497	0.808
Instant Runoff	0.736	0.789
Borda Count	0.849	0.964
Copeland	1	1
IRLWD	1	0.977
IRBC	1	0.981
Kemeny-Young	1	0.988
Least Worst Defeat	1	0.982
Ranked Pairs	1	0.984
Schulze	1	0.983

7 candidates, 1001 voters, 10000 elections:

Voting Method	Condorcet Ratio	Copeland Ratio
Plurality	0.460	0.796
Approval	0.475	0.812
Instant Runoff	0.658	0.722
Borda Count	0.839	0.965
Copeland	1	1
IRLWD	1	0.974
IRBC	1	0.977
Kemeny-Young	1	0.985
Least Worst Defeat	1	0.979
Ranked Pairs	1	0.981
Schulze	1	0.981

The first thing we notice is that Condorcet methods which do not satisfy the Copeland condition (that is, all Condorcet methods except the Copeland method) in a binary fashion have very high Copeland ratios. This means that while we can readily find counterexamples to

the Copeland condition, they occur extremely rarely. Naturally, they occur more frequently as the number of candidates increase (and the complexity of the margin of victory matrix increases). In the case of three candidates, no examples exist, since the only possibilities are a Condorcet winner (which must be selected), or a three-way tie with Copeland score. The reason for the high Copeland ratio among these methods is that counterexamples can only occur when there is no Condorcet winner. We summarize these results in terms of the number of candidates:

Voting Method	$m = 3$	$m = 4$	$m = 5$	$m = 6$	$m = 7$
IRLWD	1	.993	.984	.977	.974
IRBC	1	.994	.987	.981	.977
KY	1	.997	.992	.988	.985
LWD	1	.994	.987	.982	.979
RP	1	.995	.990	.984	.981
Sch.	1	.995	.989	.983	.981

Second, we can analyze the continuous performance of the non-Condorcet methods with respect to the Condorcet Ratio. This information is summarized below, again in terms of the number of candidates;

Voting Method	$m = 3$	$m = 4$	$m = 5$	$m = 6$	$m = 7$
Plur.	.764	.653	.567	.506	.460
App.	.599	.547	.517	.497	.475
IRV	.944	.877	.802	.736	.658
BC	.898	.870	.852	.849	.839

We see that Instant Runoff and Borda Count perform significantly better than the more commonly used Plurality and Approval methods (which quickly drop below fifty percent success). Instant Runoff Voting performs very well for a small number of candidates, but quickly deteriorates, while the Borda Count remains very consistent, having a Condorcet Ratio over eighty percent for all simulations studied. Unsurprisingly, the Copeland Ratio scores follow very similar trends for the non-Condorcet methods.

8 Agreement Simulations

8.1 Philosophy

One of the merits of binary conditions is that when systems follow them, we gain a guaranteed subset of outcomes. All Condorcet methods, for example, will yield the same election winner in elections with Condorcet winners. Given that a set of clone invariant methods yield the same result on a particular election, we know that the methods will also agree on any election which is a “clone extension” of the original. When measuring conditions continuously, we often lose this certainty. A Condorcet Ratio of 0.88 doesn’t tell us which 12 percent of Condorcet winners are missed, and two methods with Condorcet Ratios of 0.88 could differ on as much as 24 percent of elections with Condorcet winners, or they could agree on all of them. The information simply isn’t there.

Philosophically, if two methods agree on *all possible elections*, they should be considered the same method, since their outputs are identical. We’ll see that some seemingly distinct methods actually do agree in all cases, for certain numbers of candidates. However, it is a reasonable question to ask what level of agreement is significant. Would policymakers be satisfied in replacing a very complicated, expensive system with an simpler, cheaper one if they agreed 99 percent of the time? Again, the philosophical questions can’t be considered without the statistical information, thus we turn to simulations to determine the relative agreement of our various preferential ballot voting systems.

8.2 Simulation of Random Elections

While we used random elections in the previous sections as a worst case scenario in simulating actual election behavior, here we use random elections in order to best approximate the exact percentage of elections on which two different voting methods agree. Again, we simulate using a random distribution of votes over the set of possible preference orderings. As the sample of elections increases, we approach an even distribution over all possible elections (for a given number of voters and candidates), and thus approach the true percentage of agreement.

We have the following random election simulations for $n = 1001$ voters and the given number of candidates and sample size of elections:

3 candidates, 1001 voters, 100000 elections:

	<i>Plur.</i>	<i>App.</i>	<i>IRV</i>	<i>BC</i>	<i>Cope.</i>	<i>IRLWD</i>	<i>IRBC</i>	<i>KY</i>	<i>LWD</i>	<i>RP</i>	<i>Sch.</i>
<i>Plur.</i>	1.000	0.527	0.730	0.745	0.685	0.711	0.716	0.727	0.727	0.727	0.727
<i>App.</i>	0.527	1.000	0.545	0.573	0.531	0.558	0.561	0.566	0.566	0.566	0.566
<i>IRV</i>	0.730	0.545	1.000	0.813	0.860	0.907	0.908	0.902	0.902	0.902	0.902
<i>BC</i>	0.745	0.573	0.813	1.000	0.817	0.830	0.846	0.876	0.876	0.876	0.876
<i>Cope.</i>	0.685	0.531	0.860	0.817	1.000	0.912	0.912	0.912	0.912	0.912	0.912
<i>IRLWD</i>	0.711	0.558	0.907	0.830	0.912	1.000	0.984	0.954	0.954	0.954	0.954
<i>IRBC</i>	0.716	0.561	0.908	0.846	0.912	0.984	1.000	0.970	0.970	0.970	0.970
<i>KY</i>	0.727	0.566	0.902	0.876	0.912	0.954	0.970	1.000	1.000	1.000	1.000
<i>LWD</i>	0.727	0.566	0.902	0.876	0.912	0.954	0.970	1.000	1.000	1.000	1.000
<i>RP</i>	0.727	0.566	0.902	0.876	0.912	0.954	0.970	1.000	1.000	1.000	1.000
<i>Sch.</i>	0.727	0.566	0.902	0.876	0.912	0.954	0.970	1.000	1.000	1.000	1.000

Immediately, we see that four methods are equivalent for three candidate elections (this is relatively unsurprising, since there are only six possible transitive preference orderings). Kemeny-Young, Least Worst Defeat, Ranked Pairs, and the Schulze method all agree (and satisfy both the Condorcet and Copeland conditions) for these elections. We find that the other Condorcet methods all agree over ninety percent of the time, the most divergent being the Copeland method (because of the shared victory implied in cases without a Condorcet winner). As we might expect, the Condorcet ratio is a good predictor for the agreement of Condorcet and non-Condorcet methods; Instant Runoff and Borda Count agree with the Condorcet methods in more than eighty percent of elections, while Plurality and Approval average approximately seventy percent and fifty-five percent agreement, respectively.

4 candidates, 1001 voters, 50000 elections:

	<i>Plur.</i>	<i>App.</i>	<i>IRV</i>	<i>BC</i>	<i>Cope.</i>	<i>IRLWD</i>	<i>IRBC</i>	<i>KY</i>	<i>LWD</i>	<i>RP</i>	<i>Sch.</i>
<i>Plur.</i>	1.000	0.433	0.589	0.621	0.529	0.579	0.585	0.598	0.598	0.598	0.599
<i>App.</i>	0.433	1.000	0.455	0.517	0.439	0.487	0.491	0.498	0.500	0.498	0.500
<i>IRV</i>	0.589	0.455	1.000	0.709	0.725	0.806	0.810	0.804	0.799	0.803	0.800
<i>BC</i>	0.621	0.517	0.709	1.000	0.717	0.762	0.781	0.818	0.822	0.818	0.824
<i>Cope.</i>	0.529	0.439	0.725	0.717	1.000	0.826	0.826	0.826	0.826	0.826	0.826
<i>IRLWD</i>	0.579	0.487	0.806	0.762	0.826	1.000	0.967	0.918	0.909	0.915	0.908
<i>IRBC</i>	0.585	0.491	0.810	0.781	0.826	0.967	1.000	0.945	0.936	0.942	0.937
<i>KY</i>	0.598	0.498	0.804	0.818	0.826	0.918	0.945	1.000	0.975	0.986	0.978
<i>LWD</i>	0.598	0.500	0.799	0.822	0.826	0.909	0.936	0.975	1.000	0.982	0.997
<i>RP</i>	0.598	0.498	0.803	0.818	0.826	0.915	0.942	0.986	0.982	1.000	0.985
<i>Sch.</i>	0.599	0.500	0.800	0.824	0.826	0.908	0.937	0.978	0.997	0.985	1.000

As the set of possible margin of victory (and win-loss) matrices grows more complex, we begin to see a divergence of results. For elections with more than three candidates, there are no methods which are equivalent, although the same set of four methods (Kemeny-Young, Least Worst Defeat, Ranked Pairs, and Schulze) do agree over ninety-seven percent of the time (with Least Worst Defeat and Schulze disagreeing on only three elections per thousand). The Condorcet Instant Runoff methods (IRLWD and IRBC) continue to behave coherently, agreeing more than ninety-six percent of the time. We also see a quick deterioration in the

agreement percentages of the Plurality method, concomitant with its decreasing Condorcet ratio.

5 candidates, 1001 voters, 20000 elections:

	<i>Plur.</i>	<i>App.</i>	<i>IRV</i>	<i>BC</i>	<i>Cope.</i>	<i>IRLWD</i>	<i>IRBC</i>	<i>KY</i>	<i>LWD</i>	<i>RP</i>	<i>Sch.</i>
<i>Plur.</i>	1.000	0.363	0.474	0.524	0.428	0.482	0.488	0.501	0.504	0.502	0.505
<i>App.</i>	0.363	1.000	0.388	0.482	0.389	0.440	0.444	0.452	0.454	0.453	0.455
<i>IRV</i>	0.474	0.388	1.000	0.618	0.614	0.707	0.713	0.711	0.700	0.705	0.702
<i>BC</i>	0.524	0.482	0.618	1.000	0.662	0.711	0.732	0.775	0.779	0.775	0.783
<i>Cope.</i>	0.428	0.389	0.614	0.662	1.000	0.765	0.766	0.771	0.767	0.769	0.767
<i>IRLWD</i>	0.482	0.440	0.707	0.711	0.765	1.000	0.953	0.889	0.866	0.880	0.866
<i>IRBC</i>	0.488	0.444	0.713	0.732	0.766	0.953	1.000	0.924	0.901	0.911	0.904
<i>KY</i>	0.501	0.452	0.711	0.775	0.771	0.889	0.924	1.000	0.948	0.966	0.954
<i>LWD</i>	0.504	0.454	0.700	0.779	0.767	0.866	0.901	0.948	1.000	0.953	0.992
<i>RP</i>	0.502	0.453	0.705	0.775	0.769	0.880	0.911	0.966	0.953	1.000	0.960
<i>Sch.</i>	0.505	0.455	0.702	0.783	0.767	0.866	0.904	0.954	0.992	0.960	1.000

6 candidates, 1001 voters, 15000 elections:

	<i>Plur.</i>	<i>App.</i>	<i>IRV</i>	<i>BC</i>	<i>Cope.</i>	<i>IRLWD</i>	<i>IRBC</i>	<i>KY</i>	<i>LWD</i>	<i>RP</i>	<i>Sch.</i>
<i>Plur.</i>	1.000	0.319	0.390	0.461	0.363	0.411	0.418	0.430	0.433	0.428	0.433
<i>App.</i>	0.319	1.000	0.330	0.460	0.359	0.409	0.415	0.427	0.430	0.426	0.431
<i>IRV</i>	0.390	0.330	1.000	0.545	0.532	0.618	0.625	0.625	0.613	0.617	0.615
<i>BC</i>	0.461	0.460	0.545	1.000	0.633	0.673	0.697	0.742	0.751	0.740	0.754
<i>Cope.</i>	0.363	0.359	0.532	0.633	1.000	0.719	0.722	0.732	0.725	0.728	0.726
<i>IRLWD</i>	0.411	0.409	0.618	0.673	0.719	1.000	0.936	0.865	0.828	0.849	0.828
<i>IRBC</i>	0.418	0.415	0.625	0.697	0.722	0.936	1.000	0.907	0.870	0.884	0.872
<i>KY</i>	0.430	0.427	0.625	0.742	0.732	0.865	0.907	1.000	0.916	0.940	0.921
<i>LWD</i>	0.433	0.430	0.613	0.751	0.725	0.828	0.870	0.916	1.000	0.923	0.991
<i>RP</i>	0.428	0.426	0.617	0.740	0.728	0.849	0.884	0.940	0.923	1.000	0.930
<i>Sch.</i>	0.433	0.431	0.615	0.754	0.726	0.828	0.872	0.921	0.991	0.930	1.000

7 candidates, 1001 voters, 10000 elections:

	<i>Plur.</i>	<i>App.</i>	<i>IRV</i>	<i>BC</i>	<i>Cope.</i>	<i>IRLWD</i>	<i>IRBC</i>	<i>KY</i>	<i>LWD</i>	<i>RP</i>	<i>Sch.</i>
<i>Plur.</i>	1.000	0.275	0.332	0.404	0.319	0.369	0.373	0.386	0.387	0.386	0.389
<i>App.</i>	0.275	1.000	0.288	0.424	0.329	0.378	0.385	0.396	0.395	0.392	0.396
<i>IRV</i>	0.332	0.288	1.000	0.477	0.456	0.538	0.542	0.543	0.529	0.532	0.531
<i>BC</i>	0.404	0.424	0.477	1.000	0.613	0.648	0.671	0.721	0.737	0.715	0.741
<i>Cope.</i>	0.319	0.329	0.456	0.613	1.000	0.693	0.696	0.709	0.702	0.704	0.703
<i>IRLWD</i>	0.369	0.378	0.538	0.648	0.693	1.000	0.929	0.843	0.804	0.828	0.804
<i>IRBC</i>	0.373	0.385	0.542	0.671	0.696	0.929	1.000	0.884	0.846	0.858	0.850
<i>KY</i>	0.386	0.396	0.543	0.721	0.709	0.843	0.884	1.000	0.899	0.917	0.906
<i>LWD</i>	0.387	0.395	0.529	0.737	0.702	0.804	0.846	0.899	1.000	0.891	0.989
<i>RP</i>	0.386	0.392	0.532	0.715	0.704	0.828	0.858	0.917	0.891	1.000	0.900
<i>Sch.</i>	0.389	0.396	0.531	0.741	0.703	0.804	0.850	0.906	0.989	0.900	1.000

The trends of divergence continue for five, six and seven candidate elections. Unsurprisingly, no two methods behave equivalently for such elections. However, we do notice that Least Worst Defeat and Schulze continue to show a very coherent response, agreeing in nearly ninety-nine percent of all elections through seven candidates. The set of methods which were equivalent for three candidates remain similar, agreeing on approximately ninety percent of elections. On the other hand, the non-Condorcet methods begin to disagree more often than not, not only with Condorcet methods, but with each other. It is telling that the Plurality and approval methods agree less than forty percent of the time with every method except for the Borda Count. We would expect the divergence of methods to magnify as the number of candidates increases further.

9 Conclusions

Given all of the data we've collected, we would like to make a confident decision about the "best" available election method. Unfortunately, as we've alluded to throughout the paper, this simply isn't possible. We don't have normative measures of the relative importance of different voting conditions, and in the absence of a method which satisfies all conditions (which we've proven can not exist), we must make some compromises. The main purpose of this thesis has been to provide a more complete set of information on which to base decisions; we feel strides have been taken to do just that. However, if only for the exercise or as an example, it is also important to put this new information to practical use.

9.1 The Importance of the Condorcet Condition

The Condorcet condition is often the first voting system condition to be considered in a classroom setting [1] or in voting theory texts [2], [6], [7]. There are several good reasons to use the Condorcet condition as a necessary condition for a good voting system, and we use it as a first test in selecting a voting system:

First, intuitively it is reasonable to select the candidate which beats every other candidate in pairwise matchup. To select another candidate seems to accept a minority view as the correct social choice (as a minority must perceive this candidate to be better than the Condorcet winner).

Second, Condorcet winners are common. Consider the following data (from [7]), the percentage of Condorcet winners is uniform random elections (as conducted in the simulations of this paper):

Number of Candidates	2	3	4	5	6	7	8	9
Probability of Condorcet	1	.91226	.82452	.74869	.68477	.63082	.61010	.54547

This means that in elections with fewer than ten candidates, a Condorcet winner should occur at least half of the time. This is a significant subset of elections that we can agree on. Choosing a Condorcet method guarantees an expected result for all such elections. Further, even in elections with more than ten candidates, in most practical cases, the election will involve less than ten (probably closer to three) candidates with a legitimate chance of winning (frontrunners), and thus will likely behave as an election with fewer candidates, increasing the chance of having a Condorcet winner.

Finally, suppose we pit a Condorcet winner against the winner of any other voting system in a head-to-head runoff election. The Condorcet winner will win *every time*, which is a powerful statement to the social desire for a Condorcet winner.

Thus, we take the Condorcet condition as necessary, eliminating the common methods of Plurality, Approval, Instant Runoff, and Borda Count from our short-list of possible "best"

voting methods.

9.2 Monotonicity, Clone Invariance, and Loser Independence

Because the majority condition is implied by the Condorcet condition, we do not need to consider it further. We also discount the Copeland condition since it is tailored to the Copeland method, and serves better as a continuous measure, rather than a binary measure. This leaves us with three conditions to differentiate our Condorcet methods; monotonicity, clone invariance, and loser independence. We will not make any attempt here to justify one of these conditions as more important than the others; this depends too much on personal philosophy and the specific political and practical considerations involved in different election landscapes. Instead, we'll simply see which methods satisfy the most of the conditions.

Unfortunately, as we showed previously, the Condorcet condition is incompatible with the Loser Independence condition. This means we can not use it to distinguish between Condorcet methods (and is a motivation to develop a related continuous measure). Thus, we look to our summary table for any methods which satisfy both monotonicity and clone invariance.

Three such methods exist; Ranked Pairs, Schulze, and Least Worst Defeat (which only satisfies internal clone invariance).

9.3 Agreement and Practicality

Currently, we don't have a good way to distinguish between these methods. We do know that they mutually agree a large percentage of the time, and that Least Worst Defeat and the Schulze method agree especially often. Without other conditions (which would be, admittedly, arbitrarily chosen to eliminate methods), or a classification of the elections which cause disagreement, we are left at somewhat of a standstill. We can, fortunately, be confident that any one of these three systems will serve election purposes well, given the set of good conditions which they satisfy, and their relative agreement.

We are then motivated to think practically. If we are to introduce a new system of voting to a skeptical public, it behooves us to choose the simplest, clearest method. There can be no doubt that this is Least Worst Defeat. Voting reform is a tricky political process as it stands; choosing a mathematically complex system would only make the prospect of change more frightening. Thus, because of the set of conditions it satisfies, its agreement with other theoretically desirable methods, and its practical application, we recommend Least Worst Defeat as the "best" available new voting method.

9.4 Future Research

In this thesis, we've made strides towards developing a better library of information and data with which to compare preferential ballot voting systems. There are three main areas in which we would like to extend this work, to build a more complete profile of information.

- Developing continuous measures of conditions other than the Condorcet and Copeland Ratios. Finding and simulating reasonable measures of Monotonicity, Clone Invariance, and Loser Independence will allow us to better compare Condorcet methods which fail one or more of these conditions.
- Classifying elections which cause divergent results from different voting methods. Our simulation of the agreement matrices revealed the extent to which different methods agree or disagree on possible elections. We can extend this idea further by examining those elections for which elections disagree. Understanding these examples may allow someone to make a qualitative decision between similar methods (such as Least Worst Defeat and Schulze).
- Exploring non-deterministic systems. Dr. Hubert Bray has developed a non-deterministic voting method which satisfies most all of the conditions considered here (except the Copeland condition) for elections with less than 6 candidates. Understanding these systems may lead to the discovery of a system which satisfies more theoretical properties, and this gain could be weighed against a loss of determinism.

10 References

- 1: Bray, H. Lecture Notes; Game Theory and Democracy, Duke University, Fall 2008.
- 2: Saari, D. Basic Geometry of Voting. Springer-Verlag, Berlin, 1995.
- 3: Schulze, M. A New Monotonic, Clone-Independent, Reversal Symmetric, and Condorcet-Consistent Single-Winner Election Method (Draft). <http://m-schulze.webhop.net/schulze1.pdf>. 2009.
- 4: Tideman, T.N. Independence of Clones as a Criterion for Voting Rules. Social Choice and Welfare 4: pp. 185-206, 1987.
- 5: Bartholdi, III, J., Tovey, C. A., and Trick, M. A. Voting Schemes For Which it Can Be Difficult to Tell Who Won the Election. Social Choice and Welfare, Vol. 6, No. 2: pp. 157-165, 1989.
- 6: Hodge, J., Klima, R. The Mathematics of Voting and Elections: A Hands-On Approach. American Mathematical Society, Providence, RI, 2005.
- 7: Gehrlein, W. Condorcet's Paradox. Springer-Verlag, Berlin, 2006.