



An Introduction to Vote-Counting Schemes

Jonathan Levin, Barry Nalebuff

The Journal of Economic Perspectives, Volume 9, Issue 1 (Winter, 1995), 3-26.

Your use of the JSTOR database indicates your acceptance of JSTOR's Terms and Conditions of Use. A copy of JSTOR's Terms and Conditions of Use is available at <http://www.jstor.org/about/terms.html>, by contacting JSTOR at jstor-info@umich.edu, or by calling JSTOR at (888)388-3574, (734)998-9101 or (FAX) (734)998-9113. No part of a JSTOR transmission may be copied, downloaded, stored, further transmitted, transferred, distributed, altered, or otherwise used, in any form or by any means, except: (1) one stored electronic and one paper copy of any article solely for your personal, non-commercial use, or (2) with prior written permission of JSTOR and the publisher of the article or other text.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

The Journal of Economic Perspectives is published by American Economic Association. Please contact the publisher for further permissions regarding the use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/aea.html>.

The Journal of Economic Perspectives
©1995 American Economic Association

JSTOR and the JSTOR logo are trademarks of JSTOR, and are Registered in the U.S. Patent and Trademark Office. For more information on JSTOR contact jstor-info@umich.edu.

©2001 JSTOR

An Introduction to Vote-Counting Schemes

Jonathan Levin and Barry Nalebuff

The design of an electoral system is fundamental to any democracy. Motivation for understanding how electoral system design matters comes from many directions: the creation of new constitutions in eastern Europe; the recent trauma of a three-way presidential election in the United States; the ongoing debate over U.S. redistricting and gerrymandering; the furor caused by Lani Guinier's call for a more representative voting system. The status quo adds another incentive. Plurality rule is pervasive even though it is a flawed system. Fortunately, there is no lack of suitable alternatives. One of the purposes of this symposium—and this paper in particular—is to help illustrate and motivate the remarkable variety of alternative mechanisms that aggregate individual preferences to decide an election. This overview offers a summary of sixteen distinct methods and demonstrates by example how some of the more complex systems work. The shorter papers that follow focus on some of the more popular alternatives, including approval voting, single transferable vote, the maximum likelihood rule, and more.

One can speculate on why alternatives to plurality rule have had such a difficult time being adopted. Part of the cause may be Arrow's general possibility theorem. Arrow (1951) demonstrates that any voting system applied to an unrestricted collection of voter preferences must have some serious defect; we must always choose between flawed alternatives. With conflicting theoretical guidance to help select the least-flawed option, people evaluate a system by its likely effect on the status quo outcome. Since those in power tend to want to

■ *Jonathan Levin is an M. Phil. candidate at Nuffield College, Oxford University, Oxford, Great Britain. Barry Nalebuff is Professor of Economics, School of Management, Yale University, New Haven, Connecticut.*

preserve the status quo—the status quo electoral system is what brought them into power—we should not be surprised by the difficulty of implementing electoral reform.

An electoral system has to balance multiple objectives: establishing legitimacy, encouraging participation, discouraging factionalization. We focus on the goal of representativeness. How do we go about choosing a representative outcome? What is meant by representation?¹ Should we examine the position of the elected representatives or the position of the policies that they pass? How do we trade off these objectives? We do not believe that there is any one right answer (although there are certainly many wrong ones).

The point of this paper is to move away from theoretical discussions of various properties; our aim is to provide the motivation behind different counting schemes. Are they better suited for choosing a single winner or for ranking the candidates? Do they tend to favor candidates with loyal minorities or candidates who are acceptable to all and the favorite of none? Are they simple enough to be used for a general political election where voters may only be familiar with one or two candidates, or are they more applicable in a board of directors situation where each voter possesses more detailed information? We try to give some examples of how and why they work and where they fail. We have also taken several of the methods and applied them to voting data gathered from British Union elections (data collected separately by N. Tideman and I. D. Hill). An interesting feature of these British elections is that voters are required to rank the candidates. As a result, knowing the voter ranking, we can simulate elections under a variety of electoral systems. It is perhaps remarkable that among the 30 elections we examined, with the exception of plurality rule and single transferable vote, none of the other seven alternatives considered gave a different top choice (see later section). The systems differed in the rankings of the lower candidates. This empirical regularity suggests a connection to some recent theoretical work (Caplin and Nalebuff, 1988, 1991): when voter preferences are sufficiently similar, a variety of voting systems lead to similar choices, and these choices have desirable properties. The difficulties in aggregating preferences arise in the case of a population with a lack of a consensus; this is the situation where the choice of electoral system can make the greatest difference and where apparently minor differences can directly influence the outcome.

The next section of our paper describes the basic information on which most vote-counting schemes rely: voter rankings and paired comparisons. We then describe five voting rules scored directly from the voter rankings, six paired-comparisons rules, and two additional rules derived from sports rankings. An Appendix then offers a few additional rules. We then attempt to point

¹To give just the most simple example, if the population is uniformly distributed between positions 0 and 1, and we are to choose three representatives, should they be equally spaced [0.25, 0.5, and 0.75], or should they be selected so as to minimize the average distance traveled to the nearest legislator [0.16, 0.5, and 0.83]?

out how differences between rules might affect candidates' strategies and the outcome of a real election. In the final section, we address different factors that distinguish the methods and attempt to provide a basis to choose between methods. In our description of the various vote-counting rules, we have benefited greatly from the excellent surveys done by Lowell Anderson (1990, 1994) and Nicolaus Tideman (1993).

Voter Rankings and Paired Comparisons

Before describing the various counting rules, we need to consider their general structure and the information they rely on. Even if we actually had knowledge of all the voters' utility functions, there is the quandary of trying to make interpersonal comparisons of utility. This leads us to take as data a voter's ordinal ranking of the various candidates. However, there is a fundamental problem of getting people to reveal their true preferences accurately, since all voting systems encourage strategizing. Even if we had truthful rankings, there remains the issue of how to take averages over these rankings. Our focus is on this last step. Most of the variety in electoral schemes comes from the choice of different metrics for measuring the distance between one ranking and another.

From the start we should note that in a multi-candidate election, simple plurality rule throws away too much data. The voter's first choice is a poor summary statistic of preferences. Information regarding second and later choices is valuable in helping aggregate preferences. At the other extreme, we cannot reasonably ask people to vote in all possible pairwise elections. Even if there were only seven dwarfs running in a primary, this would require a voter to make 21 choices. Fortunately, we can easily infer pairwise preferences from ordinal voter rankings. For example, if a voter ranks the candidates $a > b > c$, we infer that a is preferable to b and c and that b is preferable to c . (Of course, making these inferred rankings imposes a consistency condition that voters might not obey, as discussed in Amartya Sen's companion paper.) With one exception (approval voting), all the methods we discuss use as their base data the voter's ranking of the individual candidates. In general, these rankings provide more information than we need to tally the election. For our purposes, we assume that voters rank all the candidates on their ballots, and do not score candidates as tied.²

²Many theorists have addressed the issue of how to deal with voters who either fail to rank some candidates, or rank two or more candidates as tied. Because of the immediate complications these issues generate, we try to avoid raising them in the general discussion. In a paired comparisons approach, there are basically two options when a voter ranks two or more candidates as tied. When we compare the candidates head-to-head, we can either ignore that voter, or give each candidate 1/2 point. If we do not count the voter, some head-to-head matches will have fewer total points than others. Since this leads to problems later, we could normalize all head-to-head scores so that a candidate's score is the percentage of the electorate won against the opponent. This method differs

Voting theorists from the Marquis de Condorcet to Kenneth Arrow have shared the conviction that we should judge candidates on the basis of their pairwise performance. A candidate who wins every head-to-head matchup is called a Condorcet winner, in honor of Condorcet's *Essay on the Application of Mathematics to the Theory of Decision-Making* (1785). A Condorcet winner will win an election under most of the rules we describe below (although in some instances a good argument exists for not choosing the Condorcet winner, as Peyton Young explains in this issue). Most of the difficult issues in vote counting arise when no Condorcet winner exists; instead there is a voting cycle, i.e., three candidates a, b, c such that a beats b , b beats c , and c beats a .³ Each method treats cycles differently, leading to discrepancies in how they rank the candidates.

Once we translate the voter rankings into paired comparisons, we often organize the information into a "paired-comparisons matrix." In this matrix, the ij th entry is the number of votes for i over j . The entries on the diagonal are all 0. Writing down this matrix actually loses information contained on individual ballots. That is, the matrix might show that candidate i received 45 votes over j , but it does not reveal whether those voters placed i first on their ballots and j last, or i sixth and j seventh. We may make an informed guess, but we cannot recover individual rankings from the paired-comparisons matrix.

Some voting rules further distill this information by only distinguishing the winner of each head-to-head comparison. They utilize a "win-loss matrix," where the ij th entry is 1 if more voters prefer candidate i to j and -1 if more voters prefer j to i . If there is a tie between candidates i and j , we enter a 0. Once again, all entries on the diagonal are 0 by definition.

The first five voting rules discussed are categorized as "rank-scoring" rules because they score directly from the rankings. The following six rules are "paired-comparisons" rules because they rely on that matrix. The final set of rules use a variety of different approaches.

1–5. Rank-Scoring Rules

1. Plurality Voting

Plurality voting is the most common method of ranking candidates in an election and is used in almost every political election in the United States.

from giving each candidate $1/2$ point in a tie. For example, if the score is 60 to 30 among 90 of the voters, with 10 ties, normalizing the scores results in a 66.6 to 33.3 election, while adding $1/2$ point makes the score 65 to 35. What if a voter fails to rank some candidates? In our empirical work computing various schemes, we assume a ranked candidate is preferable to an unranked candidate. This means a single unranked candidate losses every head-to-head comparison and is effectively ranked last. If two or more candidates are unranked, we may place them in a tie for last place on the ballot and then treat the tie vote in one of the two ways described above.

³If there is no Condorcet Winner, there must be a voting cycle, but not vice versa. An example of an election with a Condorcet winner and a cycle is $a > b$, $a > c$, $a > d$, $b > c$, $c > d$, $d > b$. Here a beats the other three candidates, who are in a cycle.

Under plurality election rules, each voter picks a single most preferred candidate. The candidate with the most votes wins. In an election with more than two candidates, a candidate does not necessarily need a majority vote to win. Bill Clinton won the 1992 presidential election with less than 50 percent of the vote.

When there are only two candidates in an election, plurality voting is the natural choice as an election rule: it is simply the rule of the majority. With three or more candidates, however, plurality voting can lead to disturbing results. Suppose there is an election with three candidates, two of whom have closely related views, while the third candidate has a radically different platform. Even if a large majority would choose either of the two related candidates over the third candidate, the third candidate might win if the majority split its votes between the two similar candidates.

In an indirect way, this coordination failure may have even changed the presidency of the United States. As told by Lowell Anderson (1990), the story starts with the 1966 Democratic primary for the governor of Maryland:

George P. Mahoney received about 40 percent of the vote while his two opponents, Thomas Finan and Carlton Sickles, each received about 30 percent. Both Finan and Sickles are relatively liberal, and Maryland is a relatively liberal state. Mahoney is an unabashed ultraconservative, and it is extremely unlikely that he could have beaten either Finan or Sickles in a one-on-one contest. Maryland is a heavily Democrat state. However, in the main election, many Democrats could not support the ultraconservative Mahoney, and sufficiently many crossed over to vote Republican that the Republican candidate won. It is widely believed that had either Finan or Sickles won the Democrat primary, then he would have beaten the (at that time) relatively obscure Spiro T. Agnew, in the main race. Agnew won, was later elected vice-president, and then resigned under pressure. Richard Nixon nominated Gerald Ford in Agnew's place and, when Nixon resigned, Ford became president.

The rest, as they say, is history.

While most presidential elections involve only two serious candidates, redistributing the votes that went to serious third-party candidates (Wallace, Anderson, Perot) would have had the potential to swing an election. The coordination failures of plurality rule tend to be most glaring in a primary, where one often finds more than two serious candidates.

However, we should emphasize that the number of candidates in an election is not exogenous. It is partly determined by the type of electoral system in use. One of the consequences of using plurality rule is that it leads the outcome towards a two-party system. This empirical result is known as Duverger's law. When there are more than two parties, people tend to abandon

the third party so as not to waste their votes.⁴ This may not be a bad thing; an electoral voting rule that did a better job of representing preferences over more than two candidates might result in a proliferation of candidates, and the resulting factionalization could be worse for the democracy in the long run (as Douglas Rae explores in this issue). Since the job of winnowing a large number of candidates in a primary is quite different from the job of choosing between a smaller number in the general election, society might consider using a different voting rule for primaries and for the general election.

1a. Multiple-Winner Extensions

If an election is to produce multiple winners, there are several extensions of plurality rule. One option, used to elect the legislature in Japan, is the “single nontransferable vote.” Everyone gets one vote, and we simply pick the top several candidates based on the plurality voting. Under cumulative voting, voters are given a number of votes equal to the number of winners desired.⁵ Depending on the rules, voters may or may not assign more than one of their votes to a single candidate. The candidates with the highest vote totals win the election.

In theory, cumulative voting is designed to foster minority representation. In an election with ten positions to fill, any group that controls 10 percent of the vote has the opportunity to elect one candidate. How these coalitions might form is less clear, and failure to coordinate could result in a loss of electoral power. In this regard, the single transferable vote electoral system (discussed below) is a more satisfactory method, although it is more complicated.

2. Approval Voting

Approval voting was invented in the 1970s. It is the only method we consider that does not require voters to rank the candidates. Instead, voters approve or disapprove each candidate on the ballot—that is, they select some subset of candidates. Candidates are ranked by the number of voters who approve them, and the highest ranked candidate (or candidates if more than one winner is required) wins. Thus with three candidates *A*, *B* and *C*, a voter has the opportunity either to vote for *A* or against *A* (by approving both *B* and *C*). Approving all three candidates is equivalent to not voting since it has no differential impact. Approval voting is the most frequently adopted alternative to plurality rule and is now used by several professional associations to elect their officers (Brams and Nagel, 1991). Robert Weber’s companion article focuses on this voting method.

⁴This argument is a bit specious, since nearly all votes are wasted in the sense that they do not determine the outcome of an election. And votes are equally well wasted on a party with a large lead. Evidence from the ‘94 elections suggests that Duverger’s law may be breaking down.

⁵Endorsement of cumulative voting was one of the “radical” positions taken by Lani Guinier. For a popular discussion, see *The New York Times Magazine*, Feb. 27, 1994.

A variant of approval voting can be used if we do not require a fixed number of winners, but rather a level of acceptance. We can set a quota and declare as winners any candidates reaching that quota. The election of baseball players to the Hall of Fame employs this form of approval voting. Sportswriters vote up or down on each candidate, and candidates who receive a certain (previously specified) number of votes head for Cooperstown.⁶

3. Runoff Voting

In runoff voting, or a “double election,” voters rank the candidates, and votes are tabulated just as in a plurality election. If one candidate commands a majority vote, that candidate wins the election. If no candidate receives more than half the votes, we have a runoff election between the two candidates. The winner of the head-to-head runoff wins the election. Again, variations exist to runoff voting. In New York City primaries, the top two candidates engage in a runoff if the top candidate receives less than 40 percent of the vote in the primary. In France, candidates must also garner a certain minimum vote in order to participate in the runoff election. This has sometimes led to the peculiar result of a runoff election with only one candidate on the ballot, which then leads to imaginative and colorful voter responses (Rosenthal and Sen, 1980)!

In a multi-candidate election, runoff voting can prevent some of the potentially skewed results generated by a plurality count. Suppose we have four candidates, three similar candidates who evenly divide 70 percent of the vote, and a radically different candidate who commands a 30 percent minority. The 30 percent candidate would win outright under plurality, but would suffer a sound defeat in a runoff against any one of the other three candidates.

4. Single Transferable Vote

The single transferable vote, also known as “Cincinnati Rules” or “Hare voting,” extends the logic of runoff voting by eliminating no more than one candidate at a time. A single transferable vote election with only one winner is sometimes called the “alternative vote.” We consider this case first. Voters begin by ranking the candidates. If any candidate is ranked first by a majority, that candidate wins immediately. If no winner exists, we eliminate the candidate with the fewest first place votes and tabulate the ballots again as if that candidate never existed. This means votes for the losing candidate get redistributed to one of the remaining candidates. If still no candidate controls a majority, we again delete the candidate with the fewest votes and repeat the process. Eventually, barring a tie, a winner must emerge.

When there are many candidates and multiple winners, the algorithm for a single transferable vote becomes considerably more complex. However, it still

⁶Actually, there is an upper bound on the number of players a sportswriter may list on his or her ballot. The limit is ten players, which is high enough that except in extraordinary years, the vote ceiling is rarely a factor.

works quite well in theory and in practice, finding use in New York City school board and Cambridge city council elections. With multiple winners, each one cannot control a majority, but rather should control some dominant share. Following the logic of the majority winner, we want this share (or quota) to be as small as possible, yet big enough to allow no more than the desired number of winners. For example, with two winners, we would require a winning candidate to control at least 34 percent of the vote. This is the smallest percentage that could be gained by no more than two candidates.⁷

Now suppose there are w winners and a corresponding quota q . The election runs as above until we find the first winner. When a candidate exceeds the vote quota, single transferable vote rules specify that the winner's surplus—that is, the number of votes the candidate receives over the quota—be redistributed to the voters' second-place choices. We call the surplus s , where $s = \text{total votes received} - q$, and divide it up using the following method: the winner receives a score of q and is then deleted from the ballots. Next, we tabulate who receives the transferred votes, but count them not as one full point, but as $s/(s + q)$. This weighting ensures that the sum of all candidate's scores remains equal to the number of votes originally cast, which in turn guarantees that we do not exceed w winners (recall that we chose the quota big enough so that not more than w candidates could control a share q). If the transferred votes result in another winner, we repeat the surplus distribution process for the new winner. The election ends when w candidates reach the quota.⁸

One attractive feature of single transferable vote in a multi-winner election stems from the fact that it theoretically leads to proportional representation. A united minority can elect candidates in proportion to the size of the minority, ensuring diverse representation and avoiding tyranny of the majority. Suppose a region can elect six representatives. Any candidate who can control 17 percent of the vote will win, so that of course a dominant majority candidate will win, but beyond that, a candidate ranked second or third by a majority may or may not defeat a candidate with a small but loyal minority base. Also, unlike cumulative voting, there is no need to coordinate. If two candidates appeal to the same group, some voters can rank them 1 and 2 and the others 2 and 1, and the group will not lose its voice. In particular, single transferable vote avoids gerrymandering district lines to ensure minority representation (as

⁷In general, we define a quota $q = \lfloor n/(w + 1) \rfloor + 1$, where n is the number of voters, w is the desired number of winners, and the bracket notation $\lfloor x \rfloor$ means "greatest integer less than x ." This choice achieves the goal of having the smallest possible quota such that no more than w candidates may exceed it. This discussion oversimplifies the quota issue; N. Tideman's companion paper in this issue offers a more thorough examination.

⁸In the Cambridge and New York elections, votes are hand counted, a practice fixed by law. Because of this, rather than assign transferred votes a weight of $s/(s + q)$, $s/(s + q)$ ballots are randomly selected as transferred votes, and count as one full vote. Clearly, this is less preferable, because it adds an element of randomness to the election. Even with this time-saving procedure, Cambridge city council elections have taken as long as a week to be counted.

discussed in Pildes and Niemi, 1993). If congressional elections were done on a statewide basis using single transferable vote, there would be no need to engage in the time-consuming and controversial act of redistricting following each census.

5. Coombs Voting

Clyde Coombs (1964) suggests a variant of the alternative vote. In successive rounds, instead of deleting the candidate with the fewest first-place votes, we eliminate the candidate with the most last-place votes. The election ends when only the desired number of winners remains. Unlike Hare voting, where candidates may qualify as winners in early rounds, Coombs voting always requires the full number of rounds. Duncan Black (1958, p. 69) describes an identical scheme called “exhaustive voting,” which requires voters to approve all but one of the remaining candidates in each round. After each round of voting, the candidate with the fewest votes is eliminated. Again, the election ends when we reach the desired number of winners.⁹

Coombs voting loses some of its appeal in the light of analysis by Myerson (1993) and Cox (1990). Under Coombs voting, a candidate who takes a stance favored by a majority of voters may lose if all the other candidates take the opposite position. For examples, suppose 75 percent of the electorate favors a certain proposition, but only one of five candidates supports it. The views of the candidates are identical on all other issues. Then 25 percent of the voters will rank the differing candidate last, while the majority, the 75 percent, if they split their last place votes equally, give each of the other four candidates less than 20 percent of the last place votes. Even though 75 percent of the electorate agreed with the proposition, the candidate in favor will be the first eliminated.

6–11. Paired-Comparisons Rules

6. Borda Voting

Jean-Charles de Borda, in a paper that marks the beginning of serious study in voting theory, proposed this elementary rule to the French Academy of Sciences in 1770. Suppose there are k candidates. Voters submit their rankings, and candidates receive $k - 1$ points for every first place vote, $k - 2$ points for every second place vote, and so on. We rank the candidates by their total points. Borda voting is a familiar scheme and is widely used to rank candidates or teams—the AP football poll is one example. Since every point a candidate receives may be considered a head-to-head vote against some other candidate, Borda scores are equal to the total number of head-to-head votes a

⁹It is easy to imagine other reverse elimination schemes. For example, the Nanson rule, which eliminates candidates based on their Borda scores, appears in the Appendix.

candidate receives. This means we can count Borda scores by writing a paired-comparisons matrix and summing the rows to generate the candidates' scores.¹⁰

Interestingly, a Condorcet winner will not necessarily win a Borda election. We show a very simple example of this phenomenon with the pairwise-comparisons matrix and three candidates:

	<i>a</i>	<i>b</i>	<i>c</i>	Borda	<i>a</i>	<i>b</i>	<i>c</i>
<i>a</i>	0	51	51	Scores	102	114	84
<i>b</i>	49	0	65	<i>a</i> is a Condorcet winner but			
<i>c</i>	49	35	0	<i>b</i> wins under Borda rules.			

Here, the Condorcet winner *a* defeats *b* and *c* by a minimal margin, 51-49, while *b* beats *c*, 65-35. This election raises the question of whether one winner is intrinsically better than the other. As *b* beats *c* so decisively, we can be fairly sure that *b* is preferable to *c*, but less sure about the contests involving *a*. On a different day, with slightly different voter turnout, *b* or *c* might well beat *a*, but it is unlikely that *c* would beat *b*. In such cases, the Borda winner may have a better claim on the election (as Peyton Young discusses in this issue).

We should note that there are an infinite number of possible variations to the basic Borda scheme.¹¹ For example, the choice of a linear point system is somewhat arbitrary; there is no reason to suppose that ordinal rankings should translate neatly into linear preferences. By manipulating the point values, we could arrive at dozens of alternative schemes.

Additional variations to Borda can be created by coupling it with other systems such as Coombs voting. E. J. Nanson, an Australian mathematician, worked on voting theory between 1875 and 1922. He proposed a variation of Borda voting that applies the Borda score in successive rounds. Under the Nanson rule, any candidate with a below-average Borda score is eliminated.¹² We then recompute Borda scores and repeat the elimination procedure. The last remaining candidate wins the election. This resembles Coombs voting as described earlier, except that low Borda scores substitute for last-place votes.

7. Copeland Voting

A. H. Copeland (1951) proposed the obvious paired-comparisons scheme where candidates are scored by their win-loss record across all head-to-head competitions. Employing the win-loss matrix, we sum the rows to determine each candidate's Copeland score. The sum of row *i* is equal to the number of pairwise wins candidate *i* has minus the number of losses. Equivalently, we might rank candidates by their winning percentages.

¹⁰Obviously, from its original description, Borda voting could also be thought of as a rank-scoring rule.

¹¹Duncan Black (1958) proposed a compromise between Borda and Condorcet. Under the Black rule, we compare candidates based on their head-to-head performances. If a Condorcet winner exists, that candidate wins the election. Otherwise, the candidate with the highest Borda score wins.

¹²If there are *n* voters and *k* candidates, the "average" score is $n(k/2)$.

An example of Copeland scoring in sports occurs when competitors play a round-robin, and the teams or individuals are ranked by their number of victories. The first rounds of the Olympic hockey competition and the World Cup finals use this method.¹³ When we use a Copeland system, we must also adopt a contingency plan for ties, since they may occur in the case of a cycle. Both the Olympics and the World Cup use total goals as a tiebreaker.

We might say intuitively that since professional sports leagues rank teams by winning percentages, they are using Copeland rules; but, if we consider each game between two teams as equivalent to the decision of a single voter, they are actually using Borda scoring. That is, teams are ranked by their total number of victories (or votes) over all head-to-head contests. In baseball, for example, teams in the same league play each other either 12 or 13 times each season, and 12 or 13 decisions are scored—under Copeland scoring, the winner of the season series would get one point. Of course in sports, rather than having ballots with rankings, we are truly interested in observing the $n(n-1)/2$ pairwise contests separately.

A Copeland winner must defeat all other candidates if a Condorcet winner exists. Even when there is no Condorcet winner, the Copeland winner dominates every other candidate in the sense that he either defeats him directly or defeats a third candidate who defeats him (Maurer, 1978; Miller, 1980). We say he can defeat any other candidate through a chain of length one or two.¹⁴ However, just as several candidates might tie for first in a Copeland election, more than one candidate could dominate each of the others through a chain of length one or two.

8. Minimum Violations

Minimum violations is the first computationally difficult method we cover. After voters rank the candidates, we tabulate the head-to-head votes to determine the winner of each matchup. We then consider each permutation of the candidates a , b , c , and so on, as a potential ranking. The best ranking, according to the minimum violations criteria, is the permutation that has the fewest “contradictions,” where by contradiction we mean that a candidate with lower rank defeats a candidate with higher rank.

One method to compute a minimum violations ranking is to use the win-loss matrix and maximize the sum of the entries above the diagonal by permuting the rows and columns. From this algorithm, we can see that the winning list of candidates will never display the feature that a candidate is ranked *immediately* above another who defeated her head-to-head. If this were the case, the list that switches only these two candidates would result in one less violation and no other changes.

¹³In order to encourage aggressive offensive play, the World Cup now treats ties in a novel manner, awarding both teams one point. When one team wins a game, they receive three points to the loser's zero. Devaluing a tie is meant to provide both teams with incentives to play more exciting soccer.

¹⁴However, Niemi and Riker (1976) note that under certain conditions, the Copeland method can choose as the winner a candidate who loses by near unanimity to the Borda winner.

If there are no cycles, the minimum violations results will be identical to the Copeland ranking, and this order will be a zero violations ranking. Problems can arise if there are voting cycles, because in this instance a zero violations ranking does not exist, the best ranking is not necessarily unique, and the minimum violations criterion provides no particular way of resolving these ambiguities.

9. Ranked Pairs

Condorcet's (1785) seminal paper expresses the idea that a candidate who would defeat each of the others head-to-head should win the election. If no such candidate exists, then large majorities should take precedence over small majorities in breaking cycles. In his own words, the general rule was "to take successively all the propositions that have a majority, beginning with those possessing the largest. As soon as these first decisions produce a result, it should be taken as a decision, without regard for the less probable decisions that follow." How is this idea implemented? Consider all the possible lists that order the candidates from top to bottom. Find the largest margin of victory in any pairwise match—and then eliminate all potential rankings that contradict this preference. For example, if the largest victory is for candidate a over candidate b , eliminate all potential rankings which place b above a . Next, consider the second largest margin of victory, and eliminate all potential rankings that disagree. Continue this process until only one ranking remains.

With only three candidates, this method is well defined and is equivalent to ignoring the election with the smallest margin of victory. The problem is that in elections with four or more candidates, considering the largest unconsidered margin of victory may, at some point, force us to eliminate all remaining potential rankings (by locking in a cycle). Condorcet does not discuss this possibility, an omission which has led to criticism and some confusion.

T. N. Tideman suggests one solution to the dilemma of cycles: simply skip over a head-to-head result that will lock in a cycle. Tideman further notes that if ties exist, there may be more than one potential ranking left even after we have considered all victories. In this case, we declare a tie among the candidates who have first-place ranks in the remaining potential rankings. Tideman calls this scheme "ranked pairs."

In the pairwise-comparisons matrix below, we first lock in $a > b$, then $b > c$. This implies that $a > c$. The next largest victory is $c > a$, but this locks in a cycle, so we ignore that head-to-head result. We then lock in $a > d$, $b > d$, $c > d$, which leaves us with a final ranking $a > b > c > d$.

	a	b	c	d
a	0	61	41	51
b	39	0	60	51
c	59	40	0	51
d	49	49	49	0

We further discuss this example in the Kemeny-Young section.

10. Simpson-Kramer Min-Max Rule

The Simpson-Kramer min-max rule adheres to the principles offered by Condorcet in that it emphasizes large majorities over small majorities. A candidate’s “max” score is the largest number of votes against that candidate across all head-to-head matchups. The rule selects the candidates with the minimum max score. A Condorcet winner will always be a min-max winner. When there is a cycle, we can think of the min-max winner as being the “least-objectionable” candidate. It is the person whose biggest defeat is the closest to 50 : 50. Large majorities take precedence over small majorities in that we are willing to ignore defeats if they are close enough to 50 : 50.

While the min-max rule works well for choosing the winner of an election, it may be less effective as a ranking technique, especially when a dominant Condorcet winner exists. In the election below, candidate *a* defeats all others by a large margin. As a result, the other candidates are ranked solely by their performances against *a*, while their head-to-head matchups are disregarded. Despite beating *b* and *c* head-to-head, candidate *d* places last in the election by virtue of having the poorest score against *a*.

	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>					
<i>a</i>	0	61	63	64	Min-Max	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>
<i>b</i>	39	0	49	45	Scores	39	61	63	64
<i>c</i>	37	51	0	44					
<i>d</i>	36	55	56	0					

11. Kemeny-Young Method

H. Peyton Young (1988) concluded that Condorcet intended to rank candidates in the order with which the most voters agree. In other words, we should maximize the number of head-to-head votes that agree with the ranking (or minimize the number of votes that disagree). As it turns out, Kemeny (1959) had suggested a method identical to the Young interpretation.

Kemeny-Young voting is very similar to the minimum violations method, except that it emphasizes decisive wins over smaller majority margins. Again, begin by considering all possible rankings of the candidates. However, instead of maximizing the sum of the entries in the upper diagonal of the win-loss matrix, Kemeny-Young voting uses the paired-comparisons matrix. The example nearby demonstrates this process. The left-hand pairwise-comparisons matrix shows the original rankings. The sum of entries in the upper right half is $51 + 45 + 58 + 53 + 44 + 49 = 300$. Note that a cycle exists in these rankings: $c > a$, $d > c$, and $a > d$. Kemeny-Young resolves this cycle by placing $c > a$, $a > d$ and $c > d$. In the right-hand matrix, the sum of the upper

diagonal is $55 + 49 + 47 + 58 + 51 + 56 = 316$.

	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>		<i>c</i>	<i>a</i>	<i>d</i>	<i>b</i>	
<i>a</i>	0	51	45	58		<i>c</i>	0	55	49	47
<i>b</i>	49	0	53	44	⇒	<i>a</i>	45	0	58	51
<i>c</i>	55	47	0	49		<i>d</i>	51	42	0	56
<i>d</i>	42	56	51	0		<i>b</i>	53	49	44	0

How do these rankings make *c* end up first? After all, although *c* beats *a* decisively, it loses to both *b* and *d*. Well, *c* must beat *a*, because the largest majorities must be respected. However, *c* cannot be *immediately* ahead of either *b* or *d*. Just as in the minimum violations ranking, a candidate cannot be ranked directly above another candidate who defeated her head-to-head. If *c* must be ahead of *a*, but cannot be one place ahead of either *b* or *d*, then the one possible position for *c* is first. Because *a* defeats *d* decisively, we place *a* ahead of *d*, and the rest follows. If *b* or *d* complain that they “should” have won the election, based on pairwise results with *c*, the Kemeny-Young answer is that both *b* and *d* clearly deserve to be defeated by *a*, and thus have a weak claim to victory.

When there is a Condorcet winner, Kemeny-Young and minimum violations both rank that winner first. Differences appears when there are cycles. The minimum violations winner in the above election is not *c*, but *a*. Minimum violations actually ranks *c* last.

It is also interesting to compare the min-max and Kemeny-Young methods. Again, both will rank a Condorcet winner first. But if there is a cycle, the Kemeny-Young and the min-max rule may not agree on the winner; and even if a Condorcet winner exists, they may not rank the lower candidates in the same order. In the example above, the min-max ranking is *cabd*; the Kemeny-Young rank is *cadb*. Examples can readily be created where the two methods pick dramatically different winners. In the example in the ranked-pairs section, candidate *d*, who receives 49 votes in each pairwise competition, is the min-max winner, and yet places last under Kemeny-Young rules. This disparity makes sense in that *a*, *b*, and *c* are locked in a cycle. It would be hard to make an argument that candidate *d* should be inserted into the middle of the cycle rather than placed above or below it. Interestingly, Caplin and Nalebuff (1988) prove that in a spatial voting model—where voters each have a most preferred point and each candidate occupies a position in Euclidean space—this outcome could not occur. In particular, the ranked-pairs winner will agree with the min-max winner because no cycle can occur if we restrict attention to majorities larger than the biggest margin involving the min-max winner.

In this election, candidate *a* dominates the other candidates, and receives by far the highest power ranking. Although *b* beats *c* convincingly, 7-3, *c* earns the highest power score by virtue of stronger performances against *a* and *d*. As a comparison, the Borda method would agree with the power rank method, while the Kemeny-Young and ranked-pairs rules would rank *b* above *c*. The min-max rule, the Copeland rule, and the minimum violations method would place them in a tie.

13. Jech Method (Maximum Likelihood Estimation)

The Jech method, also proposed by Zermelo (1929), Bradley and Terry (1952), and Ford (1957), is a probabilistic ranking method. Jech's (1983) goal is to provide an ordering of the candidates that not only demonstrates whether *i* is superior to *j*, but also by how much. Each candidate *i* is given a strength T_i . Given these strengths, Jech defines a probability matrix p where the probability that *i* defeats *j* in a pairwise election (or game) is $p_{ij} = T_i / (T_i + T_j)$. Jech makes two assumptions: first, that odds multiply and second, that expected wins equal actual wins. Jech proves that only one matrix exists that satisfies this criteria. Candidates (or teams) receive a score equal to their expected winning percentage. In an election, this ranking is equivalent to the Copeland method. This can be viewed as an independent argument for using the Jech method. For sports teams, Jech's method works even if not all teams play one another. In this case, Jech's method predicts the number of wins a team might have in a full round-robin tournament (essentially an expected Copeland score).

Choosing an Appropriate Voting System

In the examples we looked at, there was surprisingly little difference between the winners under the various election methods. We tested nine of the methods—plurality, single transferable vote, Borda, Copeland, min-max, Kendall-Wei, power ranking, minimum violations, and Kemeny-Young—using voting data from British union elections. All elections were multi-candidate elections where voters ranked their preferences. With the exception of plurality rule, all the other methods obtained similar results. Plurality rule frequently resulted in a different winner than the other methods, and single transferable vote occasionally led to a different outcome. The other methods essentially differed only when there was a voting cycle, and even then, it did not affect the winner. While rankings were not identical, a Condorcet winner typically emerged, and this winner tended to be the Borda winner as well.

Given the variety of election formats, and that their results do not differ so dramatically, it is natural to ask how to choose among them. The answer may depend on which features of a voting rule are most important in a particular situation. We consider five aspects that distinguish the various methods.

Level of Complexity

A voting method should be relatively simple and transparent, both for voters and for those calculating the winner. The tolerable level of complexity depends on how many voters and candidates there are, who the voters are, and what purpose the election serves.

Simplicity helps explain why plurality voting is so widespread and why approval voting and Borda voting are the two most frequently used alternatives. The U.S. presidential election is something of an anomaly to this principle of simplicity—hence the periodic calls for a direct (plurality) election to replace the electoral college. In some cases, we accept greater complexity to gain accuracy. For example, the professional tennis rankings and *The New York Times* college football rankings (forms of power ranking) both require a computer for calculation. They can account for a great deal of information, much more than if we ranked tennis players simply by total matches won or tournaments won.

Voting in a single transferable vote election is straightforward; the main hurdle would be explaining the vote-counting procedure to voters and establishing the legitimacy of the system. Here, charts can be extremely useful. Edwin Newman and Miles Rogers (1952), in their analysis of the 1951 Cambridge city election, compare cumulative voting outcomes where the voters are given between 1 and 9 votes (there were nine city council seats). In their chart, presented as Table 1, we see how the candidates move up and down the rankings as voters are allowed more choices. The graph clearly demonstrates the importance of second, third and lower-place choices in capturing voter preferences. Note how candidate 20 rises from seventh to second place, candidate 06 rises from eighteenth to eighth position, and candidate 14 drops from fifth to twelfth.

The issue of how voters express preferences on a ballot is separate from how the preferences are combined. When a corporate board votes on a set of well-defined proposals, it would be reasonable to ask the board members to rank their preferences. A board member would presumably be well informed and interested, so we could use a more involved paired-comparisons technique to determine the winning choice. On the other hand, in an election for student government, with many candidates and mostly disinterested voters, ranking candidates could be capricious. Instead, we might ask voters to approve or disapprove each candidate, or approve a set number, or just pick their favorite. Simplicity is a relative term.

Voter Strategies

In an election, the possibility always exists for voters to cast their ballots strategically, rather than in accordance with their true preferences. All voting systems are susceptible to strategizing. For example, in a three-candidate plurality election, a voter might vote for her second choice if her preferred

Table 1
Plurality Count with Each Number of Choices

<i>Rank</i>	<i>P.R. Count</i>	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>	<i>5</i>	<i>6</i>	<i>7</i>	<i>8</i>	<i>9</i>
1	03	03	03	03	03	03	03	03	03	03
2	25	25	25	20	20	20	20	20	20	20
3	05	23	20	25	05	05	05	05	05	05
4	20	05	05	05	25	25	25	25	11	11
5	23	14	23	23	23	23	11	11	25	25
6	11	26	26	26	26	11	23	23	23	23
7	14	20	14	07	11	26	26	26	26	26
8	07	11	11	11	07	07	07	07	07	06
9	26	17	07	14	14	14	14	14	06	07
10	17	27	17	17	17	17	17	08	08	08
11	27	08	27	27	08	08	08	17	14	18
12	08	07	15	08	27	02	06	06	18	14
13	15	15	08	15	02	27	02	18	17	17
14	01	01	01	02	15	06	18	02	02	16
15	02	02	02	01	06	18	27	27	16	02
16	18	19	19	18	18	15	15	15	27	27
17	19	18	18	06	01	19	19	16	15	15
18	06	06	06	19	19	01	16	19	19	19
19	12	12	12	12	16	16	01	01	01	01
20	16	16	16	16	12	12	12	12	12	12
21	13	13	13	13	13	13	13	04	04	04
22	21	21	09	09	09	09	04	13	13	13
23	24	09	21	21	24	04	09	09	09	24
24	09	24	24	24	21	24	24	24	24	09
25	04	04	04	04	04	21	21	21	21	10
26	10	10	10	10	10	10	10	10	10	21
27	22	22	22	22	22	22	22	22	22	22

Notes: Chart showing rank of candidates under assumption of vote by varying number of choices. Numbers in the table refer to candidates.

candidate lags far behind in preelection polls. An honest vote might be considered “wasted.” Another person might vote for the lagging candidate as a protest vote, even if she did not prefer that candidate over the others. Or suppose two of the three candidates—the ones with relatively similar views—are running neck and neck. Supporters of the first might rank their candidate first and the other last, even if in truth they would have chosen the other as their second choice. In recognition of this danger of strategizing, Borda wrote that his method was appropriate only for honest voters. For this reason, Borda counting works well when there is little incentive to strategize—for example, in the AP sports poll where writers, in general, have no vested interest in the outcome of the poll.

Voter strategizing provides a counterargument to the benefit of simplicity; the more complex schemes are harder to strategize and by their very complexity may help promote honest behavior.

Candidate Strategies

Electoral systems may influence candidate strategies as well. In plurality voting, a minority candidate could win if two similar candidates split the majority vote. Because of this, some candidates would do better to generate strong minority support instead of trying to attract a majority vote. Cox (1987, 1990) confirms this suspicion, finding that candidates do best by solidifying minority support rather than taking centrist positions. Myerson (1993) shows that in a two-candidate plurality election, candidates fare best when exactly half the electorate are promised benefits above the average. As the number of candidates increases, Myerson demonstrates that the optimum strategy shifts: candidates do better by promising the bulk of the resources to a smaller and smaller minority, seeking to cultivate a small bloc of strong support rather than widespread appeal.

When there are only two candidates, the approval voting and plurality rule lead to the same incentives. As the number of candidates increases, the incentives are different under approval voting. Here, to win broad approval, the optimum campaign strategy is to spread the bulk of the resources over a larger and larger segment of the electorate.

For single transferable vote elections, more than one possible equilibrium exists. Candidates might opt to seek majority support, but they might also choose to cultivate minority support at the expense of the majority. Suppose all but one candidate offer a distribution of benefits to voters that equals the distribution that each of these candidates would have granted in a plurality election with four candidates. The final candidate offers a uniform distribution of benefits, seeking majority support. She will be the very first candidate eliminated, even though head-to-head she would defeat any of the others. If the final candidate chooses the optimum strategy for a five-candidate plurality election, she will rank first in each stage until only four candidates remain, but

will face elimination at this stage, as all the other candidates have taken the optimum four-candidate strategy.

All paired-comparisons methods, including the Borda rule, offer candidates precisely the same incentives. Rather than encouraging candidates to seek minority support, candidates in a paired-comparisons election do best to appeal to a majority.¹⁵ We see this intuitively if we look at the head-to-head comparisons as simultaneous two-candidate plurality elections, all of which are symmetric. A candidate wants to do well in each of these individual head-to-head contests, hence the two-candidate plurality strategy.

Ranking vs. Picking a Winner

In an election to fill an office, we care solely about determining one winner. In other elections, we need to elect several candidates. In a weekly football poll, the number one team matters, but we also want to know the rankings of the other teams as well. These different goals suggest different voting rules. Methods that use a good deal of the information available on the ballots, such as Kendall-Wei, power rank, Jech, Borda, Copeland, and Kemeny-Young work particularly well for ranking all the candidates. Each of these takes into account a candidate's or a team's head-to-head performance against each of the other teams. Kendall-Wei and Copeland count only wins and losses; the more complex of these ranking techniques, Jech, power rank, and Kemeny-Young, consider the margins of victory as well. Single transferable vote works well only when the quota is adjusted based on the number of candidates we seek to rank.

Plurality, approval voting, and min-max are better suited for choosing the winner of the election than ranking the candidates. Plurality works poorly for ranking mainly because it also takes into account only first-place rankings, and thus discards much of the available information. For example, a candidate who ranks second on every ballot, but first on none, probably deserves a relatively high rank, yet plurality gives that candidate zero votes. Min-max works poorly on the lower rankings when one candidate wins by an enormous margin. All the remaining candidates are then ranked solely by their performance against the strong candidate—because that will be their minimum vote—regardless of how they fare against each of the others. A candidate with only one head-to-head loss might wind up last with a bad enough defeat against the winner.

Minority Support and the Safe Choice

Suppose a business wants to introduce a new brand of cereal. They do not want to market a cereal that everyone picks as the second or third best but no one picks first. Consumers will buy the cereal they like best. It doesn't matter whether they like a cereal second best or fifth best; they won't buy it. The

¹⁵Similarly, following Cox (1987), candidates have an incentive to adopt positions close to the political center. Cox uses a spatial model of voter preferences and shows that there is an equilibrium of candidate strategies when candidates adopt the position of the median voter.

company should market a cereal that generates strong preferences, even if some tasters rank it very low. Counting systems such as plurality or weighted Condorcet, which emphasize strong preferences, work better for product testing than methods such as min-max or Borda, which reward consensus and wide approval.

Elections are different from product testing. The usual goal is a candidate who appeals to a large portion of the electorate, rather than one who draws forceful support from a strong minority. When the objective is to pick a “safe” choice, systems like approval voting or min-max will favor the least objectionable over more controversial candidates. Similarly, Borda rewards strong preferences to an extent, but a winning Borda candidate will probably not have alienated a large minority. A candidate who places second or third on each ballot stands a strong chance in a Borda, approval or min-max election, even though such a candidate has little hope of winning a weighted Condorcet or plurality election.

Conclusion

If the multitude of available vote-counting systems leaves you overwhelmed, you are not alone. Several of the relatively simple schemes—such as single transferable vote, Borda counting, min-max, Kemeny-Young, power ranking, and approval voting—possess qualities that are often desirable. Others, such as Copeland, minimum violations, or Jech may be useful for specific purposes. The diversity of questions we might ask about vote-counting procedures, their biases and their outcomes, is enormous. Many researchers, following Arrow’s (1951) lead, have concerned themselves with stating various desirable or undesirable criteria and attempting to classify systems by these means. In contrast, there are far fewer papers written concerning the application of alternative systems in practice.

Despite the wealth of alternative voting mechanisms documented here, plurality rule remains the overwhelming favorite choice. Why do alternatives to plurality rule have such a difficult time being adopted?

Part of the answer is probably that when voter preferences are sufficiently similar, a variety of voting systems will lead to similar choices, and these choices will have desirable properties. In this case, the choice between voting systems will seem to make little difference. In many other cases, society does not have a consensus about what the goals of an electoral system should be. One is to establish legitimacy of the victor. A second is to encourage participation. A third may be to discourage the formation of a large number of political parties. A fourth is to assure a representative political system although the very definition of “representativeness” is part of the problem. For participants in the system, a fifth goal is often that “their side” improve its chances of winning.

No voting system will satisfy everyone on all of these dimensions; the choice is always between flawed alternatives. Between the entrenched power of the status quo, and with conflicting theoretical guidance to help select a second-best alternative, it is not surprising that electoral reform is difficult to implement. Most of the variety in electoral schemes comes from the choice of different metrics for measuring the distance between one ranking and another.

The Marquis de Condorcet (1785) believed that in an election, there exists some underlying truth to be discovered—that one candidate should be the “true” winner. We suggest a more modest conclusion, but one that might stop us from looking for some perfect Holy Grail of a voting method that meets all needs: a voting system can’t find a consensus when none exists.

Appendix

Additional Voting Rules

Young Method

Peyton Young, whose interpretation of Condorcet’s work we have already described, proposed a separate method. Under the Young method, if there is no Condorcet winner, each candidate receives a score equal to the largest subset of voters for which that candidate is a Condorcet winner. The Young winner is the candidate with the highest score. In the example below, b is the Young as well as min-max winner.

#Voters	Ballot		a	b	c	a	b	c	
2	$a > b > c$	Young	a	b	c	a	0	6	4
3	$b > c > a$	Scores	3	9	7	b	5	0	7
2	$b > a > c$					c	7	4	0
4	$c > a > b$								

Here, a defeats the others if we count the two ballots that rank a first and any one other ballot. To compute b ’s score, we count the 5 first-place ballots and any 4 of the others, for a score of 9. Finally, for c we count the 4 first-place ballots and any three of the others, for final scores of 3, 9, and 7. The final ranking is bca . This method is similar to the min-max rule and will produce the same winner so long as the number of last-place votes each candidate receives is at least as large as that candidate’s largest loss margin in a head-to-head match (Tideman, 1993, ch. 13).

Dodgson Rule

Charles Dodgson, mathematician and author of *Alice in Wonderland* (under the pen name Lewis Carroll) studied elections and lawn tennis tournaments at

Oxford in the late 19th century. He agreed with Condorcet that a candidate who is undefeated head-to-head should win an election, but felt that if a Condorcet winner did not exist, the election should be called off. In discussing this opinion, he proposed this method of scoring candidates: if there is no Condorcet winner, the Dodgson rule holds that a candidate's score should be defined to be the total number of inversions on individual ballots necessary to make that candidate a Condorcet winner. More than one inversion may be needed on some ballots to make a candidate a Condorcet winner. This sounds difficult to compute, and it is.

A simplified Dodgson rule can be calculated directly from the paired-comparisons matrix. To compute the simplified Dodgson scores, we write the paired-comparisons matrix and for each candidate sum the vote difference in each of that candidate's head-to-head losses. The simplified Dodgson winner is the candidate who has the smallest total differential over all losses. The simplified Dodgson method bears a close similarity to the Borda rule, in that if we were to tabulate total differential over wins and losses, that is to sum the differential over all wins and subtract the simplified Dodgson scores, we would arrive at the Borda rankings.

Estimated Centrality

Discussed by T. N. Tideman and I. G. Good in 1976 (Tideman, 1993), estimated centrality assumes that each candidate occupies some spatial position, and that voters likewise have some ideal preferred point in space. Under estimated centrality, the candidate who occupies that point in space closest to the "median," or zero moment, of the voters' ideal points wins the election. (We might, in fact, take any moment of the voter's ideal points—here we take the zero moment.) Good and Tideman have proven that it is possible to use the fraction of voters who place the candidates in each of the possible orders to compute the estimated centrality winner, but the actual computation is still quite difficult.

References

- Ali, I., W. D. Cook, and M. Kress, "On the Minimum Violations Ranking of a Tournament," *Management Science*, June 1986, 32, 660–72.
- Anderson, Lowell, "Voting and Paired Comparisons." Extracted from Grotte, J., L. Anderson, and M. Robinson, "Selected Judgmental Methods in Defense Analyses." Institute for Defense Analyses paper P-2387, 1990.
- Anderson, Lowell, "How to Take Votes: New Ideas on Better Ways to Determine Winners," research draft, Institute for Defense Analyses, 1994.
- Arrow, Kenneth J., *Social Choice and Individual Values*. New Haven: Yale University Press, 1951.
- Black, Duncan, *The Theory of Committees and Elections*. London: Cambridge University Press, 1953.
- Board of Elections, Cincinnati, Ohio, "Proportional Representation Count: Rules and Instructions to Employees," 1955.
- Bradley, R. A., and M. E. Terry, "The Rank Analysis of Incomplete Block Designs. I. The Method of Paired Comparisons," *Biometrika*, 1952, 39, 324–35.
- Brams, Steven, and Jack Nagel, "Approval Voting in Practice," *Public Choice*, August 1991, 71, 1–17.
- Caplin, Andrew, and Barry Nalebuff, "On 64%-Majority Rule," *Econometrica*, July 1988, 56, 787–814.
- Caplin, Andrew, and Barry Nalebuff, "Aggregation and Social Choice: A Mean Voter Theorem," *Econometrica*, January 1991, 59, 1–24.
- Condorcet, Marquis de, "Essay on the Application of Mathematics to the Theory of Decision Making," 1785. In Baker, K., ed., *Condorcet: Selected Writings*. Indianapolis: Bobbs-Merrill, 1976.
- Coombs, Clyde, *A Theory of Data*. New York: Wiley, 1964.
- Copeland, A. H., "A 'Reasonable' Social Welfare Function," mimeo, University of Michigan Seminar on Applications of Mathematics to the Social Sciences, 1951.
- Cox, Gary W., "Electoral Equilibrium under Alternative Voting Institutions," *American Journal of Political Science*, February 1987, 31, 82–108.
- Cox, Gary W., "Centripetal and Centrifugal Incentives in Electoral Systems," *American Journal of Political Science*, November 1990, 34, 903–35.
- David, H. A., "Ranking the Players in a Round Robin Tournament," *Review of the International Statistical Institute*, 1971, 39:2, 137–47.
- Ford, L. R., "Solution of a Ranking Problem from Binary Comparisons," *American Mathematical Monthly*, 1957, 64, 28–33.
- Goddard, Steven T., "Rankings in Tournaments and Group Decision-Making," *Management Science*, December 1983, 29, 1384–92.
- Jech, Thomas, "The Ranking of Incomplete Tournaments: A Mathematician's Guide to Popular Sports," *American Mathematical Monthly*, 1983, 90, 246–66.
- Kemeny, John, "Mathematics Without Numbers," *Daedalus*, Fall 1959, 88, 571–91.
- Maurer, Stephen B., "The King Chicken Theorems," working paper, Princeton University Mathematics Department, 1978.
- Miller, N., "A New Solution Set for Tournaments and Majority Voting," *American Journal of Political Science*, February 1980, 24, 68–96.
- Myerson, Roger, "Incentives to Cultivate Favored Minorities under Alternative Electoral Systems," *American Political Science Review*, December 1993, 87, 856–69.
- Newman, Edwin B., and Miles Rogers, "PR Voting: An Analysis of the 1951 Cambridge City Election," unpublished (available from Barry Nalebuff), 1952.
- Niemi, Richard, and William Riker, "The Choice of Voting Systems," *Scientific American*, 1976, 234:6, 21–7.
- Pildes, Richard, and Richard Niemi, "Expressive Harms, 'Bizarre Districts,' and Voting Rights: Evaluating Election-District Appearances after Shaw," working paper, University of Michigan Law School, 1993.
- Rosenthal, Howard, and Subrata Sen, "Electoral Participation in the French Fifth Republic," *American Political Science Review*, March 1973, 67, 29–54.
- Stob, Michael, "A Supplement to 'A Mathematician's Guide to Popular Sports,'" *American Mathematical Monthly*, May 1984, 91, 277–82.
- Tideman, Nicolaus, "Collective Decisions and Voting," draft, March 1993.
- Young, H. Peyton, "Condorcet's Theory of Voting," *American Political Science Review*, December 1988, 82, 1231–44.
- Zermelo, E., "Die Berechnung der Turnier-Ergebnisse als ein Maximumproblem der Wahrscheinlichkeitsrechnung," *Math. Z.*, 1929, 29, 436–60.