

Are Condorcet and Minimax Voting Systems the Best?¹

Richard B. Darlington
Cornell University

Version 5b relative to arXiv numbering for this title. Nov. 15, 2019.

Here we correct some errors in the references and in the numbering of tables.

There are no real differences in content between versions 5 and 5b.

Abstract

For decades, the minimax voting system was well known to experts on voting systems, but was not widely considered to be one of the best systems. But in recent years, two important experts, Nicolaus Tideman and Andrew Myers, have both recognized minimax as one of the best systems. I agree with that. This paper presents my own reasons for preferring minimax. The paper explicitly discusses about 20 systems, though over 50 are known to exist.

Comments invited.

rbd1@cornell.edu

Copyright Richard B. Darlington

May be distributed free for non-commercial purposes

Keywords

Voting system

Condorcet

Minimax

1. Many thanks to Andrew Myers, Sharon Weinberg, Eduardo Marchena, my wife Betsy Darlington, and my daughter Lois Darlington, all of whom contributed many valuable suggestions.

Table of Contents (* indicates a noticeable revision from version 4)

1. Introduction *	3
2. The variety of voting systems	4
3. Some electoral criteria violated by minimax's competitors	6
Monotonicity	6
Strategic voting	6
Completeness	7
Simplicity	7
Ease of voting	7
Resistance to vote-splitting and spoiling	8
Condorcet consistency (CC)	8
Sensitivity to differences in vote changes (SDVC) *	8
How these electoral criteria support minimax	10
4. Dismissing eight criteria violated by minimax	10
4.1 The absolute loser, Condorcet loser, and preference inversion criteria *	10
4.2 Three anti-manipulation criteria	12
4.3 SCC/IIA	13
4.4 Multiple districts	13
5. Simulation studies on voting systems	14
5.1. Why our computer simulations use spatial models of voter behavior	14
5.2 Three artifacts that may produce Condorcet paradoxes in spatial models	15
5.3. Which voting systems deal best with the artifacts of Section 5.2?	17
5.4 Two studies on "genuine" Condorcet paradoxes *	18
5.5 Might inventing a Schulze tie-breaker make Schulze outperform minimax?	22
6. Strategic voting under minimax *	23
7. Six variants and close relatives of minimax, and a tie-breaker	25
8. More on the disadvantages of various systems	27
8.1 Why "simpler" ballots can actually make voting harder	27
8.2 Problems with positional voting systems	27
8.2.1 Examples of strategic voting under Borda and Coombs *	27
8.2.2 Is strategic voting a real danger?	29
8.2.3 Two Borda-based attacks on CC systems	29
8.3. Problems with rating-scale voting systems	31
8.3.1 Introduction and summary on rating-scale systems	31
8.3.2 Using MJ or RV, how many insincere voters are needed to tip a landslide election?	31
8.3.3 Why MJ almost forces voters to vote insincerely	33
8.3.4 Like MJ, RV can totally ignore majority rule	35
8.4 Problems with the Dodgson and Young systems	36
8.5 Problems with the Hare system (instant runoff voting or single transferable vote) *	38
References	40

1. Introduction and summary

What is the best voting system for choosing a single winner from several candidates? The Condorcet Internet Voting Service (CIVS), at civs.cs.cornell.edu, has managed over 13,000 nongovernmental elections and polls since its creation in 2003 by computer scientist Andrew Myers. In 2017 Myers switched CIVS's default voting system from the Schulze system to minimax, citing an earlier version of this paper as the main reason for the switch. A major article by Green-Armytage, Tideman, and Cosman (2016) reviewed 54 voting systems. It named the Condorcet-Hare system as best, with minimax and Hare as close runners-up. Tideman (2019) repeated this same conclusion. This paper summarizes my reasons for continuing to prefer minimax. Section 6 tells why I prefer minimax to the Condorcet-Hare system preferred by Tideman, and Section 8.5 tells why I prefer minimax to the Hare system now used in many elections in Maine, and recently adopted for many elections in New York City.

In 2015 I was led reluctantly to the conclusion that the minimax system, invented independently by Simpson (1969) and Kramer (1977), appears to be the best voting system for these elections. My own prior guess was that the Coombs system was best. But when I devised five very different computer-simulation tests of a system's ability to pick the best winner, minimax won all of them by wide margins; see Darlington (2016) or Section 5 of the current paper. That led me to study closely the well-known minimax anomalies – artificial-data examples in which minimax clearly seems to select the wrong winner. As described in Section 4, I found good reasons to ignore every one of those anomalies. (As experts know well, every voting system has anomalies; minimax is by no means unique in having them.) Minimax has other advantages involving simplicity, transparency, voter privacy, and resistance to attempts at insincere strategic voting and other types of manipulation. Altogether, several different lines of evidence and reasoning lead separately to the conclusion that minimax is the best voting system for most elections, with other Condorcet-consistent systems being the next best.

Unfortunately, I must disappoint readers seeking a short and simple path to this conclusion. In the last 750 years analysts have proposed at least 54 voting systems (most in the last century), and there is no easy way to dismiss all of them at once. More detailed but now-dated descriptions of the area appear in Tideman (2006), Saari (2001), and Felsenthal (2012). This paper is designed to be useful to both beginners and experts in this area. The paper's sections are arranged in the order which I believe beginners will find most useful. But more knowledgeable readers can read Sections 3-8 in almost any order. Section 2 briefly describes 18 voting systems identified by Felsenthal (2012) as the ones most worthy of detailed analysis. Section 3 summarizes the anomalies and weaknesses of many of these systems. Section 4 analyzes eight anomalies in minimax, arguing that there are good reasons to ignore all of them. Sections 3-4 show that even before we examine computer simulations, there are several different ways to conclude that minimax is the best of the 18 systems examined by Felsenthal (2012).

Section 5 turns to computer simulations. It considers two different models of voter behavior – models I'll call "excellence" and "spatial." In the former, candidates differ on only a single dimension we might call "excellence." In a spatial model, political issues which divide voters are represented by spatial axes, each voter and each candidate appears as a dot in that space, and each voter is presumed to prefer the candidates closest to themselves. A large study of hundreds of real elections, by Tideman and Plassmann (2012), found that spatial models fitted this data far better than excellence models did. Section 5 describes several quite different spatial-model simulation studies, all of which minimax won by wide margins. Skeptical experts in voting theory are urged to read at least Sections 4 and 5.

Section 6 discusses minimax's resistance to insincere strategic voting. Section 7 examines several variants and close relatives of minimax, and chooses as best a system I call minimax-PM. Section 8 describes in more detail some of the problems with other voting systems, including a good deal of original material. Section 8 was placed last mainly because it's quite long and is summarized in Section 3.

Several sections also show the substantial and consistent superiority of Condorcet-consistent systems (not just minimax) over other voting systems. I do not claim to have established the superiority of minimax and other Condorcet-consistent systems for all time, beyond all possible future evidence and analysis. But I do claim that several very different lines of evidence and reasoning suggest this conclusion given our current knowledge.

2. The variety of voting systems

This section says a little about each of the 18 voting systems analyzed by Felsenthal (2012). Nearly everyone has used the very simple **plurality** (vote for one) system. Another system with very simple ballots is **approval voting**, in which each voter approves as many candidates as they wish, and the candidate with the most approvals wins. Also simple and moderately familiar is **plurality runoff**, which uses plurality voting followed by a runoff election between the top two candidates.

Several other systems use **ranked-choice ballots**, in which voters rank the candidates. Some of these systems are **positional** systems, meaning that the winner is determined by the ranks received by each candidate. Four positional systems are especially well known. The **Hare** system is also called instant runoff voting (IRV) or single transferable vote (STV). In that system the candidate with the fewest first-place ranks is removed, and the ranks of all other candidates are recomputed as if that one candidate had never run. That process is repeated until only one candidate is left. The Hare system is often called "ranked-choice voting" (RCV), but it is actually just one of several systems using ranked-choice ballots. The **Coombs** system is like Hare, except the candidate removed in each round is the one with the most last-place ranks rather than the one with the fewest first-place ranks. In the **Borda** system, we find for each candidate X the number of candidates ranked below X by each voter. Summing these values across voters gives each candidate's Borda count. The one with the highest Borda count is the winner. Thus if there are no tied or missing ranks, the Borda winner is the candidate with the highest mean rank. In the **Bucklin** system a candidate wins if they receive a majority of the first-place ranks. If there is no winner by that rule, each candidate's number of "high ranks" is expanded to include second-place ranks, and any candidate who receives "high ranks" from over half the voters is the winner. If there is still no winner, "high ranks" are expanded to include third-place ranks. And so on. If two or more candidates all receive "high ranks" from over half the voters, the one with the most high ranks is the winner.

Some non-positional systems use rating scales. In **range voting**, each voter rates each candidate on a multi-point scale, often ranging from 0 to 100, and the candidate with the highest mean rating wins. The **majority judgment** system also uses a rating scale, though the points on the scale have non-numeric verbal labels like "excellent" and "poor," and the number of points on the scale is much smaller – usually 5 to 8. The winner is the candidate with the highest median rating. A tie-breaker described in Section 8.3.1 breaks nearly all ties. These two systems are the best-known rating-scale systems.

If ballots use ranks or rating scales, election managers can use the ballots to run a two-way race between each pair of candidates. If a candidate wins all their two-way races by majority rule, they are called a **Condorcet winner**. A voting system has **Condorcet consistency** (CC) if any Condorcet winner is

named the winner. Surprisingly, none of the voting systems just described has CC, although all of them seem reasonable at first glance.

A **Condorcet paradox** occurs if there is no Condorcet winner. There is then a **Condorcet cycle** of three or more candidates. In a typical cycle, A beats B by majority rule, and B beats C, but C beats A. In the simplest possible example, one voter ranks three candidates as ABC, one as BCA, and one as CAB. Then A beats B 2:1, B beats C, and C beats A. Large studies by Gehrlein (2006) and Tideman and Plassmann (2012) both found that Condorcet paradoxes occur only occasionally in real data, but they are common enough that it's important to handle them as well as possible.

All the CC systems below are described by Felsenthal. CC systems typically allow tied ranks. If a voter fails to rank a candidate, they are typically presumed to rank them below anyone whom they did rank explicitly. I agree with those rules. All these systems name any Condorcet winner as winner, so we describe only what they do under a Condorcet paradox. In that event the **Black** system names the Borda winner as winner. The **Copeland** winner is the candidate who wins the largest number of their two-way races. It was recently discovered that this system had been proposed by 1283, for electing church officials, by Cardinal Ramon Llull; see Colomer (2013). In the **Young** system, we find for each candidate X the smallest number of voters who would have to be removed to turn X into a Condorcet winner, and the winner is the one for whom that number is smallest. In the **Dodgson** system, an "interchange" occurs if a voter switches two candidates whom they had ranked adjacently, as when a voter changes the ABCD ranking to ACBD or ABDC. The Dodgson winner is the candidate who would need the fewest interchanges, summed across voters, to become a Condorcet winner. In the **Nanson** system, all candidates with Borda counts below the mean are removed and Borda counts are then recomputed. That process is repeated until only one candidate is left. Nanson is the only system mentioned here which has CC even though it never explicitly examines the results of all possible two-way races. That's because a Condorcet winner doesn't always have the highest of all Borda counts, but it never has a Borda count below the mean of those counts, and thus is never eliminated by Nanson.

Minimax has CC. If there is no Condorcet winner, we find each candidate's largest margin of loss in their two-way races. Call that margin LL for "largest loss." The candidate with the smallest LL is the winner. Some writers define the size of candidate X's loss in a two-way race not as a margin but as the number of voters who voted against X. The two definitions may yield different minimax winners if there are ties or missing ranks. My reasons for defining losses as margins appear in a discussion of "minimax-WV" in Section 7. That section also describes minimax-PM, a new tie-breaking system for minimax which is simple and intuitively reasonable, breaks nearly all ties, and is shown by simulations to pick better winners than random tie-breakers.

So far we have named 9 voting systems without CC, and 6 with CC. The **Kemeny** and **Schwartz** systems are also included by Felsenthal (2012). Both have CC. The Schwartz system is the only one of Felsenthal's 8 CC systems which fails the Pareto criterion. That means it can select one winner X even though Y is preferred to X by every single voter. That failure seems especially unacceptable to me. The Schwartz system is otherwise quite similar to the Schulze system, which satisfies the Pareto criterion and is used much more widely than Schwartz. The Schwartz system is not ultimately recommended by Felsenthal or in an earlier major review of voting systems by Tideman (2006). So we shall in effect substitute Schulze for Schwartz in the rest of this paper. We later describe the major disadvantages of the Kemeny and Schulze systems but do not describe their computations, which are fairly complex.

The last of Felsenthal's 18 systems is **successive elimination**. If there are c candidates this system uses a series of $c-1$ two-way votes, with the loser of each vote eliminated from all later votes. This system is sometimes used within legislatures to choose among several versions of a bill. But it is clearly unsuitable for ordinary elections, both because it requires so many separate votes and because the final outcome can depend heavily on the order of the votes, which may be arbitrary or manipulated by insiders. This system is not mentioned again, leaving 17 systems to analyze: 3 with very simple ballots, 4 positional systems, 2 with rating scales, and 8 with CC.

3. Some electoral criteria violated by minimax's competitors

A voting system is said to suffer an *anomaly* if an artificial example can be created which demonstrates some clearly undesirable property of the system, such as if giving candidate X a higher rank can sometimes make X *lose*. An *electoral criterion* is a rule requiring the absence of a particular anomaly. Electoral theorists have widely agreed for decades that every known voting system, and indeed every possible system, suffers from at least one anomaly, and often from several. The extensive literature on this topic almost never asks how likely any anomaly is to occur in real data, since there has been no very practical way to answer such questions. I'll distinguish anomalies from *weaknesses* or *problems*, for which there is no real doubt about the frequency of the problem. For instance, vote-splitting is a ubiquitous problem in plurality elections with more than two candidates; candidates A and B might both be able to beat C in two-way races, but they both lose to C in a plurality election because they split their vote. Of the items discussed in the rest of this section, I would classify most as "weaknesses" because there is little serious doubt about their potential frequency. But the first item (nonmonotonicity) is more of an anomaly.

Monotonicity. Define a simple vote change as a change by one or more voters who all vote identically and then all change their votes identically. A voting system is nonmonotonic if a simple vote change ever changes the election's outcome in the opposite direction from the obvious intention of those voters. Thus, a system is nonmonotonic if any of these outcomes could ever occur:

1. Candidate X changes from a winner to a loser when a simple vote change raises X's ranking without changing the ranks of other candidates relative to each other.
2. X changes from a loser to a winner when the change lowers X's ranking without changing the ranks of other candidates relative to each other.
3. X changes from a winner to a loser when new identical voters appear, all ranking X first.
4. X changes from a loser to a winner when new identical voters appear, all ranking X last.

Felsenthal and Nurmi (2017) show that of the voting systems in Section 2, the only ones with monotonicity are plurality, approval, Borda, range voting, and minimax. It is not known how common nonmonotonicity is in real data, but the mere fact that it's possible might provide to some voters a convenient excuse to not bother voting. So monotonic systems are definitely preferable.

Strategic voting. Voters may sometimes gain some advantage by voting insincerely or strategically. For instance, suppose A and B are a voter's favorite candidates in that order, but the voter rates or ranks B at the very bottom to try to prevent B from beating A. That maneuver is called "burying." Nonmonotonicity may occasionally offer other opportunities for strategic voting. The Gibbard-Satterthwaite theorem, discussed in Tideman (2006, 143-148), says essentially that in single-winner elections with more than two candidates, no reasonable rank-based voting system can be entirely resistant to strategic voting in all circumstances. But the degree of resistance can vary greatly. In

a mild case, strategic voting might be effective only given three conditions: (1) there are special circumstances, such as one or even two Condorcet cycles, (2) the number of cooperating strategic voters must at least approximate the margin of victory they wish to overcome, and (3) those voters must know quite accurately how others are voting, including the special circumstances like Condorcet cycles. If they guess wrong about this, they may hurt themselves. In such cases, strategic voting is probably little real danger. Sections 4.2 and 6 illustrate cases like this involving minimax. But in the worst case, none of these conditions is required – a landslide election with no special circumstances might be tipped by just a few strategic voters who need not know how others are voting, and who run no risk that their strategic maneuvers will backfire. Section 8.3.2 gives artificial-data but realistic examples in which landslide elections, in which one candidate wins by a 10% margin under sincere voting, are tipped if 1% of voters vote strategically under range voting, or 1.4% under majority judgment. Section 8.2.1 gives less extreme but fairly similar examples for the Borda and Coombs systems, and Section 8.4 for Dodgson.

Completeness. I'll call a ballot "incomplete" if it doesn't allow a voter to express a preference between every pair of candidates. Among the systems considered here, the only ones suffering this problem are plurality, approval, and plurality with runoff. For instance, suppose in an approval election a voter would rank the candidates ABCD, and that voter chooses to approve A and B. But if it later turns out that the top two candidates are A and B, or are C and D, the voter has in effect been prevented from voting in the two-way race which determines the final winner. Incomplete ballots are often assumed to make voting easier, but Section 8.1 explains why that's not so. Thus, those ballots are unacceptable. Incomplete systems were consistently among the worst performers in all the simulation studies mentioned in Section 5; see for instance Tables 2 and 3 in that section.

Of our 17 voting systems, minimax is the only one passing all of the three criteria discussed so far. But many systems also fail other criteria named below.

Simplicity. A system fails this criterion if the average citizen would find it difficult even to understand the system or explain it to others. In my opinion the well-known Kemeny and Schulze systems clearly fail this criterion. As mentioned in Section 2, the *goals* of the Dodgson and Young systems are easily described. But papers cited in Section 8.4 show their computations can be a challenge even for modern computers. Thus, these two systems also fail our simplicity criterion.

Ease of voting. Plurality and approval ballots are often assumed to make voting easier than ranked-choice or rating-scale ballots. But plurality ballots certainly aren't easy for a voter who prefers some candidates over others but likes two or more top candidates equally well. And Section 8.1 explains why approval ballots can actually make voting very difficult for the voter who wants to maximize their influence, as most voters do. But ranked-choice ballots can make voting surprisingly easy if tied and missing ranks are allowed. Then voting will be even easier than plurality voting for anyone who likes two or more top candidates equally well, and also easier than approval voting (see Section 8.1). If there are many candidates, voting may be quite difficult with rating scales or with ranked-choice ballots if tied and missing ranks are not allowed. That's often the case with the Hare and Coombs systems, since those systems must know at each stage of the elimination process which one candidate is each voter's top or bottom choice among those remaining. CC systems allow tied and missing ranks, so voting may be easier with them than with any other system.

Resistance to vote-splitting and spoiling. These problems arise primarily with plurality voting and with Hare. A *Smith set* is a set of candidates, any one of whom could beat any candidate outside the

set in a two-way race using majority rule (MR). In plurality voting, the vote-splitting problem arises if the candidates in a Smith set are so similar to each other that they all lose because they split the votes of the voters attracted to those candidates. The similar “spoiler” problem arises if there is one candidate X who could beat anyone else in a two-way race, but one or more less-popular candidates take some of X’s votes so that X loses to some other candidate. The next paragraph shows how Hare shares these same problems.

Condorcet consistency (CC). Several different lines of reasoning can lead to the conclusion that CC is essential. Many people simply consider the point to be self-evident. Others may be persuaded by the fact that all non-CC systems are eliminated by the needs for monotonicity, completeness, and reasonable resistance to burying. Still others are persuaded by the argument that CC is politically necessary just because so many ordinary citizens consider it important. This was illustrated by the 2009 mayoral race in Burlington, VT, a liberal university town which had recently adopted the Hare system. In that race, Hare eliminated the Democratic candidate before either the Republican or Progressive candidate (who won), even though the Democratic candidate was the Condorcet winner, beating each of five other candidates in two-way races. Voters found that bizarre and promptly voted to abandon the Hare system. Section 8.5 explains how this can happen, and why we might expect to see it fairly often.

Sensitivity to differences in vote changes (SDVC). I invented this criterion. Suppose candidate X is a Condorcet winner, or could become one if just one voter changed their vote to put X at the top. Suppose any other candidate would need at least k such changes to become a Condorcet winner. A voting system fails SDVC if it would ever make X a loser no matter how large k is – even in the thousands or millions. I think of SDVC as a useful variant of CC. Both those criteria focus on the number of vote changes each candidate needs to become a Condorcet winner. But SDVC says that if the differences in those numbers are large enough, one of them needn’t actually be 0 to determine the winner. SDVC could accept a system which violates CC: if A was a CC winner, and B lost one two-way race by just a few votes, and system Q named B the winner, then Q would violate CC but not SDVC. Despite this relative “laxness” of SDVC, the only of our 17 systems which satisfy SDVC are Dodgson, Young, and minimax. Section 8.4 shows that Dodgson and Young fail criteria of simplicity, monotonicity, and transparency (not discussed here in Section 3), and Dodgson is highly susceptible to burying.

Many electoral theorists express little interest in examples with artificial data, having learned over the years that clever mathematicians seem to be able to create artificial examples that appear to prove almost anything they wish. But almost none of those examples give any explanation of how the example might arise in real life. I feel that examples should be considered seriously if they do include such an explanation, so I include one for this example. The explanation also helps us think about the example. A village must choose among three alternatives, such as three possible locations for a new park, highway, or municipal building. Three mayoral candidates A, B, C all advocate different choices, while a fourth candidate D expresses no preference and says the issue needs more study. If the three positions were equally popular and this were the only issue, then D would be the Condorcet winner, because the roughly two-thirds of the voters favoring A or B would rank D above C, and D would also beat A and B by about 2:1 ratios. But suppose D favors expanding the police force, the other candidates don’t, and most voters don’t. Suppose this issue is important enough to voters that D loses to each of the others by one vote. And suppose there is a large Condorcet cycle among A, B, and C, with A beating B, B beating C, and C beating A. As mentioned in Section 2, both Gehrlein (2006, pp. 31-58) and Tideman and Plassmann (2012) found that Condorcet cycles do arise in real elections.

Table 1. Hypothetical election results

Pattern frequencies		Margin of defeat for the first-named candidate in each two-way race					
Pattern	Freq	A<C	B<A	C<B	D<A	D<B	D<C
A B C D	101	-101	101	101	101	101	101
D A B C	101	-101	101	101	-101	-101	-101
B C A D	101	101	-101	101	101	101	101
D B C A	101	101	-101	101	-101	-101	-101
C A B D	101	101	101	-101	101	101	101
D C A B	100	100	100	-100	-100	-100	-100
Totals	605	201	201	203	1	1	1
Vote changes		101	101	102	1	1	1

These circumstances could generate the results shown in Table 1, which appears in slightly different form as Table 1 in Darlington (2016). Six voting patterns appear in column 1; each voter's favorite choice appears first and their least favorite last. The pattern frequencies appear in column 2. All other columns are derived from these two columns. In those later columns, each candidate loses the races in which they are listed first, and wins their other races. Thus, within Table 1 the "<" symbol can be read as "loses to," as in "A loses to C," and each column total is the first-named candidate's margin of defeat in that race. An entry in the table is negative if the first-named candidate beats the other in that entry's pattern, because that entry lowers the first-named candidate's margin of defeat.

Table 1 shows that A, B, and C each win two of their three two-way races but lose one. D is the minimax winner "by a mile," since the Totals in Table 1 show that D's largest loss LL is 1 while all other LL values are over 200.

Let each candidate's "vote change index" be the minimum number of vote changes that candidate would need to become a Condorcet winner. In this example, each candidate's vote change index is the smallest integer which exceeds half that candidate's LL value, because each vote change can lower LL by 2. These values appear in the bottom row of Table 1. In real elections, with tied ranks, incomplete ballots, and write-ins, it may be difficult to compute vote change indexes exactly. But that doesn't prevent us from using vote change indexes to discuss simple artificial examples like this one. We see that D needs only one vote change to become a Condorcet winner, whereas any other candidate needs over 100.

If all 6 frequencies in Table 1 were increased by the same amount, even by thousands or millions, the vote change indexes of candidates A, B, and C would all rise by that same amount, while D's index would remain at 1. So, any voting system violates SDVC unless it names D the winner. These facts bridge the gap between Table 1 and the earlier mentions of thousands or millions of vote changes. Expressing the necessary vote changes as proportions of the total number of voters, we can say that for any number of total voters, every candidate except D would need over 15% of the voters to change their votes, and their LL values would all exceed 30% of the number of voters.

Of the 17 voting systems mentioned in Section 2, the only ones naming D the winner in this example are minimax, Dodgson, Young, and plurality. But plurality fails SDVC in other examples. For instance, suppose voters differ only on a left-right dimension, and candidate A gets the 33% of votes farthest left, B gets the next 33% of the leftmost votes, and C gets the 34% farthest right. Then C wins under plurality even though B would beat either A or C by over 3 million votes if there were 10 million voters. Section 8.4 shows that Dodgson and Young have other problems. The Condorcet Internet Voting Service (CIVS), at civs.cs.cornell.edu, is the only free, highly secure, open-source voting service managed by a prominent researcher – Andrew Myers, who is also editor-in-chief of TOPLAS, the premier journal on computing languages. CIVS offers users the choice of several CC voting systems, most of which are newer and lesser known than the systems listed in Felsenthal (2012). But of all the systems offered by CIVS, minimax is the only one satisfying SDVC.

How these electoral criteria support minimax. The preceding discussion shows that we will reject all of the aforementioned 17 systems except minimax if we adopt any one of the following sets of criteria:

- Monotonicity, completeness, and substantial resistance to strategic voting
- Monotonicity and either CC or SDVC
- SDVC and either monotonicity or simplicity or transparency

Section 4.5 mentions several newer voting systems which all satisfy CC but violate SDVC. Thus, the criteria of this section provide several ways to reject all the minimax competitors discussed by Felsenthal, but only one way to reject the newer systems mentioned in Section 4.5.

4. Dismissing eight criteria violated by minimax

This section describes the eight electoral criteria which Felsenthal (2012) identifies as being violated by minimax, and explains why all those violations can be ignored. Much of the material in previous sections will be familiar to experts in voting theory, but most of the material in Section 4 is original.

4.1 The absolute loser, Condorcet loser, and preference inversion criteria

Felsenthal (2012) describes 16 electoral criteria, and identifies five of them as criteria whose violation is widely regarded as “especially intolerable.” One of these is the **absolute loser criterion**, which states that no candidate should ever win if they are ranked last by over half the voters. Another is the **Condorcet loser criterion**, which says that no candidate should ever win if they lose all their two-way majority-rule races against other candidates. Minimax violates both of those, and those violations lead Felsenthal to dismiss minimax as unacceptable. This section discusses these two criteria, plus a third – the **preference inversion criterion**, which says that no voting system is acceptable if it ever names as winner a candidate whom it would also name as winner if every voter’s ranking of the candidates were inverted.

Example 3.5.11.1, on pages 61-62 of Felsenthal (2012), shows that minimax can violate all three of these criteria. In this example we have five groups of voters, who have the voting patterns shown in cell 4 of Table 2. Those groups include 2, 3, 3, 1, 2 voters respectively. In the various two-way races, A loses to B by 3 votes, B loses to C by 3 votes, and C loses to A by 5 votes. D loses to each of the others by

just one vote, so D is the minimax winner. But D is ranked last by 6 of the 11 voters and first by only 5, so D loses all their two-way races by one vote. Thus, D is an absolute loser and Condorcet loser, so minimax violates those criteria. If all ranks were reversed, D would be a Condorcet winner and would win under almost all voting systems including minimax, so minimax also violates the preference reversal criterion.

Table 2. Artificial data showing how the positions of 4 candidates and 5 sets of voters on 3 policy dimensions can produce a large ideal winner anomaly if group sizes are 2, 3, 3, 1, 2 respectively

Positions of 4 candidates A-D on 3 policy dimensions	Positions of 5 sets of voters on those same 3 dimensions	Euclidean distance of each set of voters from each candidate, computed from figures at left				Candidate rankings based on these distances	Mean voter-candidate distances, using group sizes in title
		A	B	C	D		
A 36 91 0	54 40 89	104.14	105.02	104.34	3.00	DACB	A 80.62
B 7 10 0	51 38 88	103.82	102.29	103.89	1.00	DBAC	B 81.71
C 92 1 0	57 22 16	73.88	53.85	43.84	73.93	CBAD	C 77.87
D 52 38 88	44 35 24	61.45	50.70	63.53	64.57	BACD	D 42.29
	69 67 10	42.01	84.81	70.60	84.94	ACBD	

Now consider a spatial model in which voters and candidates differ from one another on three opinion dimensions, such as liberalism-conservatism on economic issues, foreign-policy issues, and social issues like abortion and gay rights. Each voter and each candidate is represented as a dot in that space, according to their opinions on those issues. We assume that each voter ranks the candidates by their closeness to themselves, with closest candidates ranked highest. Suppose the scores of the four candidates on those three opinion dimensions are as shown in cell 1 of Table 3, and the five groups of voters fall in the space as shown in cell 2. The three-dimensional version of the Pythagorean theorem shows that the distances between the candidates and the five groups are as shown in cell 3. Then the voters in the various groups will rank the candidates as shown in cell 4, and the mean distance of each candidate from all 11 voters is shown in cell 5. We see in cell 5 that D is by far the best candidate, with a mean distance from the voters far below that of any other candidate.

Thus, in this example, the three criteria in the title of this section all require the rejection of the candidate who is in fact the best candidate. So those criteria must be dismissed. And we reached this conclusion using what may be the best-known single example claiming to show that minimax is unacceptable.

Consider again the example in Table 1, in Section 3. Candidate D in that example is the only candidate whose selection satisfies the SDVC criterion, but selecting D violates the three criteria of this section. That gives another reason for rejecting these three criteria – they all conflict with SDVC.

4.2 Three anti-manipulation criteria

This section considers three other electoral criteria from Felsenthal's list of eight criteria violated by minimax. Felsenthal calls them the no-show, twin, and truncation criteria. I'll call them "anti-manipulation" criteria, since they are all intended to prevent various kinds of illicit electoral manipulation. A voting system violates the no-show criterion if a voter can sometimes benefit by abstaining rather than voting sincerely. The twin criterion is very similar: it prohibits an anomaly in which two people who rank candidates the same find it beneficial for one of them to vote and for the other to abstain. The truncation criterion is also somewhat similar: it is violated if a voter can benefit by ranking only their first few choices rather than all of them.

Felsenthal (2012, p. 63) presents a 19-voter artificial example in which 5 voters rank four candidates in the order DBCA, 4 others rank them BCAD, 3 others ADCB, 3 others ADBC, and 4 others CABD. He credits the example to Hannu Nurmi. The minimax winner here is B. In the no-show and twin anomalies, three voters from the last pattern choose to abstain. This changes the winner to A, whom those voters prefer to B, thus violating those criteria. In the truncation anomaly, all four voters in the last pattern truncate their ballots, giving only their first two choices. This changes the minimax winner to C, who is their first choice. These are interesting as mathematical curiosities, but there are seven reasons I find them unconvincing as examples of possible real-world manipulations. *First*, the data contain two Condorcet cycles: one with candidates A, B, and C, and one with A, C, and D. But even one cycle is quite rare in real data. *Second*, all the manipulations took 3 or 4 of the 19 voters, yet the key margins of victory (both before and after the manipulations) were all by one vote each. Thus, there is nothing like our examples of burying or max-and-min in Sections 8.2.1 and 8.3.2, in which a few strategic voters reverse very large margins of victory. *Third*, to know their maneuvers would benefit them, the strategic voters would have to know almost exactly how everyone else would vote, including the two Condorcet cycles. But such certainty is rare; recall that on US Election Day in November 2016, some respected analysts were saying with 99% confidence that Hillary Clinton would win the presidency. As we'll see in Section 8, burying and max-and-min may require no such knowledge. *Fourth*, if the strategic voters guess wrong, they would very likely harm themselves, since they are withholding all or part of their votes. *Fifth*, the voters must be persuaded to actually execute this complex and counter-intuitive plan. It makes me imagine a college math class with a full blackboard. *Sixth*, in most real elections there would have to be hundreds or thousands of strategic voters to whom this complex plan must be explained. The plan would surely end up in local newspapers, eliciting a mixture of outrage, amazement, and ridicule. This would alienate some voters who had previously planned to vote for the candidate benefited by the scheme, and it would stimulate opposing voters to actually vote. *Seventh*, after all that effort and risk, in two of the three maneuvers the participating voters don't even get their favorite candidate, just their second-favorite one. For all these reasons, this example belongs in *Alice in Wonderland*; it does not seem like a maneuver anyone would actually try to execute in real life.

The electoral criteria discussed in Sections 4.1, 4.3, and 4.4 are intended to ensure selection of the best candidates, but are not designed to thwart illicit manipulations. So, I'll call them *optimizing criteria* in contrast to the anti-manipulation criteria of Section 4.2. The anti-manipulation criteria are also intended to ensure selection of the best candidates even in the absence of attempts at illicit manipulation, so they too are optimizing criteria in a sense, and we should consider that possibility. But the whole purpose of optimizing criteria is to promote selection of the best candidates. Thus, when one optimizing criterion conflicts with another (and these all conflict with CC), it seems reasonable to use

computer simulations to see which criteria lead to selection of the best candidates. Section 5 describes a whole series of simulation studies. All found that the best candidates were selected by minimax, which violates all of the criteria in Section 4. And we see in the next two subsections that there are also other reasons for dismissing the criteria discussed there.

4.3 SCC/IIA

Felsenthal calls the criterion of this section the “subset choice criterion” (SCC). Others call it “independence of irrelevant alternatives” (IIA). A voting system violates this criterion if deleting a loser (the “irrelevant alternative”) from the contest can change the winner. But is this actually a flaw? Suppose post-balloting candidate dropouts seemed very likely in a plurality election we were planning. To deal with that, we could allow each voter to rank their top several choices, so vote-counters would know each voter’s top choice even after some candidates drop out. That clearly assesses voter desires better than the usual plurality system. But the plurality system satisfies SCC/IIA, whereas this improvement makes the system violate SCC/IIA. For instance, suppose A would win with no dropouts, but B drops out and most of B’s votes go to C, who then beats A. So, the superior system violates SCC/IIA and the inferior system satisfies it. Perhaps people consider SCC/IIA advantageous merely because that’s what they’re used to.

Another argument also raises doubts about SCC/IIA. Imagine a sports league with three teams in which each pair of teams plays 9 games. Team A has won all 9 of its games against B, and B has won all its games against C, but C has beaten A 5 games to 4. Thus, A has won 13 games, B has won 9, and C has won 5. If we named A the league champion, B would be a loser. But if B were suspended for hazing and became ineligible, and we therefore ignored the results of B’s games, we would be forced to choose C, who had beaten A 5 games to 4. Thus, the initial choice of A violates SCC/IIA. But clearly team B’s ineligibility doesn’t mean that the results of its games are irrelevant to the choice between A and C. The problem with SCC/IIA is even worse, because even if B isn’t suspended or even accused of wrongdoing, we know that B *could* be suspended, so that choosing A would violate SCC/IIA.

We thus see that each team (or each candidate in an election) has two roles: as a potential winner, and as a foil or standard of comparison for other teams or candidates. They may be useful in that second role even if they withdraw from the first. Darlington (2017) reports a wide variety of simulation studies, all showing that voting systems satisfying SCC/IIA select worse candidates than systems violating SCC/IIA, thus confirming that losers can usefully be employed as foils when choosing among the remaining candidates. Since the whole purpose of optimizing criteria is to promote selection of the best candidates, we see again that SCC/IIA is actually counterproductive.

It might be argued that the sports analogy is misleading because the outcome of each game is due partly to chance, so C may have won most of its games against A just by chance. But each vote in an election may also be affected by chance; a voter might happen to hear A’s speech but not B’s. That doesn’t mean votes are meaningless; it just means some chance is involved, as in sports.

4.4 Multiple districts

The multiple-districts criterion is violated if a candidate could win in each of two or more districts but lose if all the districts are merged into one. Darlington (2016, pp. 12-13) describes an artificial-data example found on the internet, in which minimax violates the multiple-districts criterion. In that

example, minimax makes candidate A win a 4-candidate race in each of two districts because within each district, candidates B, C, and D form a Condorcet cycle with one large margin of loss for each of those candidates. But the cycles are in opposite directions in the two districts, so the cycles nearly cancel each other out when the districts are merged. C then becomes the minimax winner, which violates the criterion. Recalling Section 4.2, it seems that examples with two separate Condorcet cycles are a recurring theme in critiques of minimax.

This example has two problems. First, the example is artificial, contrived, and unlikely to arise in real life. Second, the minimax choices actually seem to be the best ones, even in this rare situation. If an election involves just one major issue, define a “fringe candidate” as one supporting a position on that issue opposed by most voters. If that issue has three or more possible conflicting solutions (such as where to locate a municipal facility), with each solution supported by a different candidate, and those circumstances produce a Condorcet paradox, then by definition, all those candidates are fringe candidates. There may be other more centrist candidates, such as one who says the issue needs further study. But in the two-district example the BCD cycle nearly disappears when the two districts are merged, so the fringe candidates become much more centrist. Candidate A was not in the cycle, so merging the districts didn’t affect A’s centrism. But the merger did increase the centrism of the other candidates, making C more centrist than A. Darlington (2016, p. 22) studied the ability of eight voting systems to find the most centrist candidate, and minimax was by far the best; see Table 6 on page 26 of that paper. Or see the results of a similar study in Section 5.4 of the current paper. That property made minimax pick C in the merged district despite picking A in the separate districts. It’s intuitively obvious that a formerly fringe candidate could become centrist if a conservative district merged with a liberal one, but the possibility of Condorcet cycles in opposite directions means that this can happen even if the two districts were similar in liberalism-conservatism or other policy issues. Thus, minimax actually provides a good solution to a complex problem that many people hadn’t even considered. As with SCC/IIA, what appeared to be a fault is actually an advantage.

In summary, Section 4 includes arguments for dismissing all eight of the electoral criteria listed by Felsenthal (2012) as being violated by minimax.

5. Simulation studies on voting systems

5.1. Why our computer simulations use spatial models of voter behavior

This paragraph summarizes Section 5.1 so readers can decide whether they can skip to Section 5.2. Voters may respond to differences among candidates on policy dimensions like liberalism-conservatism, or on general-excellence traits like honesty, intelligence, and experience, or on both. Pure spatial models (explained shortly) recognize differences on policy but not on excellence. But differences among candidates in excellence tend to suppress Condorcet paradoxes while policy differences on multiple issues tend to produce them. Previous sections presented several reasons for dismissing non-CC voting systems: flaws involving burying, incompleteness, vote-splitting and spoiling, and the rejection of non-CC systems by both the general public and many electoral theorists. The obvious major problem faced by CC systems is the Condorcet paradox. Thus we want to understand how Condorcet paradoxes arise, and also need to study how various voting systems behave under Condorcet paradoxes. Using a voting model which suppresses Condorcet paradoxes would interfere with both of those goals. Therefore our simulations should use spatial models.

We now explain these points in more detail. Some models assume that voters all have the same goals, but they disagree on which candidates would pursue those goals most effectively, so the election's primary purpose is to maximize attainment of the agreed-on goals. This case might arise when the members of a nonprofit organization are voting to choose the organization's president, and they all agree on the organization's goals. In such models, the only relevant differences among candidates are on a trait we might call "excellence" or "general attractiveness," and voters disagree with each other only because of random differences in their perceptions of each candidate's true score on that trait.

Other models assume that voters have conflicting goals, so the election's primary purpose is to compromise among those competing goals. This case would presumably arise more often in public elections, where voters may differ on the desired amounts of military spending, business regulation, social-welfare spending, tax policies, regulations on drugs and sexual behavior, and other issues. This case is usually studied with spatial models. In these models we treat each area of disagreement (tax policy, military spending, etc.) as a policy dimension. Each dimension is represented as an axis in space, and we represent each voter and each candidate as a dot in that space, according to their positions on those issues. Each voter is presumed to rank the candidates by their distance from themselves, with the closest candidate ranked highest. In the simplest possible spatial model, voters and candidates are assumed to differ on just one policy dimension, frequently labeled liberal-conservative or left-right. That model should not be confused with the excellence model of the previous paragraph, which has no policy dimensions at all.

If we assume that voters are mutually independent and that we can measure distortions from sampling error, it's quite easy to tell whether voters disagree on policy dimensions. Suppose first that we have at least four candidates. If voters all have the same goals, a voter's choice between two candidates A and B should not correlate significantly with the choice between two other candidates C and D, because all differences among voters are presumed to be produced by random errors of judgment, which are presumed to be mutually independent. But such correlations could easily occur under a spatial model, even a model with just one policy dimension, as when A and C are both more liberal than B and D.

A different test must be used if there are only three candidates. Under an excellence model, differences among voters concerning the excellence of one candidate B will tend to produce a positive correlation between the preference for B over another candidate A and the preference for B over a third candidate C; but such correlations should never be more negative than the laws of chance would allow. But under a spatial model, suppose those candidates fall in the order A B C from political left to political right. Then those on the left will prefer A to B and those on the right will prefer C to B, but few if any will prefer both A and C to B, so the aforementioned correlation will be negative. If any three candidates all differ on a left-right dimension, as in this example, there will always be one candidate between the other two, so negative correlations like this should appear. This is shown by an artificial-data example I ran with 1000 voters, 50 candidates, and two policy dimensions, with both voters and candidates drawn randomly from a bivariate normal distribution. In samples of 1000, all correlations below -0.1 are significant beyond the 0.001 level one-tailed. With 50 candidates the number of ABC correlations of this sort is $50 \cdot 49 \cdot 48 / 2$ or 58,802. In my sample, 4075 of those 58,802 correlations were below -0.1, and the most negative correlation was -0.754.

Tideman and Plassmann (2012) studied real-data elections with several candidates. From each election they formed all possible three-candidate sets, and tested whether those three candidates seemed to differ on policy dimensions. After studying hundreds of these sets, they concluded that the evidence was overwhelming that policy dimensions were important in the elections they studied. Many voting systems, including the famous Borda and Kemeny systems, were specifically developed for excellence models. That suggests these systems will perform poorly in computer simulations using spatial models – a suspicion confirmed by a whole series of computer simulations by Darlington (2016).

It seems plausible *a priori* that the candidates in a set might differ from each other on both excellence and policy dimensions, and the results of any of the aforementioned analyses would not eliminate that possibility. But intuition suggests that the larger the differences are among candidates on general excellence, the less likely a Condorcet paradox is to appear. My own unpublished computer simulations confirm this conclusion. Thus differences in excellence don't produce Condorcet paradoxes, but rather prevent them. Therefore they don't help us understand how the paradoxes arise, and don't help us see which voting systems behave best under those paradoxes. Thus our computer simulations should use pure spatial models.

Before asking *when* a spatial model could produce a Condorcet paradox, we'll see how it could ever do so. For a simple example, suppose there are two policy dimensions X and Y. Suppose three candidates A, B, C are respectively at (1,4), (5,5), and (6,1) on X and Y, and three voters R, S, T are at (2,2), (3,6), and (7,3) respectively. Then the voter-candidate Euclidian distances in the XY space are

	R	S	T
A	2.24	2.83	6.08
B	4.24	2.24	2.83
C	4.12	5.83	2.24

Thus voter R will rank the candidates ACB, S will rank them BAC, and T will rank them CBA, producing a Condorcet paradox. Many similar paradoxes could be created, with either few or many voters.

5.2 Three artifacts that may produce Condorcet paradoxes in spatial models

Under a spatial model, a set of voters is said to have radial symmetry if voters are distributed symmetrically on any axis in that space. Univariate, bivariate, and multivariate normal distributions all have radial symmetry. Plott (1967) showed that a Condorcet paradox can never arise in a spatial model with radial symmetry, because under radial symmetry any candidate will beat any other candidate who is farther from the center of the distribution. Thus if A beats B and B beats C, A must beat C. Thus the presence of a Condorcet paradox means there must be some deviation from radial symmetry. Three artifacts can produce the paradox:

1. Random sampling error. When we draw random samples of artificial voters and candidates from a bivariate or multivariate normal distribution, the fewer voters we draw the more likely the data are to show a Condorcet paradox, even though there can be no such paradox in an infinitely large population with a normal distribution. Consistent with this possibility, Tideman and Plassmann (2012, p. 234) found that real-world Condorcet paradoxes were much more common in smaller sets of voters.

2. Voter misperceptions of candidate policy positions, due either to voter inattention and carelessness or deliberate attempts by candidates to deceive some or all voters. Using artificial data, suppose we use a spatial model to draw a random sample of voters and candidates, and we find it contains a Condorcet winner. Suppose we then add random error to the measured distance between each voter and each candidate, making those random errors independent across voters, candidates, and trials if we conduct many trials. Those errors can represent voter misperceptions of candidate positions. The larger those errors, the more likely it is that the data will show a Condorcet paradox.

3. Asymmetric distributions. Suppose we draw a sample of artificial voters from a symmetric normal distribution and find the sample contains no Condorcet paradox. If we then apply a monotonic nonlinear transformation to those data to make the distribution asymmetric, a paradox may appear.

In Section 5.1 we mentioned models in which candidates differ in overall excellence and also differ on the policy dimensions in spatial models. We saw there that differences in excellence tend to suppress rather than produce Condorcet paradoxes. Thus we might ask here why differences in excellence don't overwhelm and suppress the artifacts just mentioned, thus eliminating Condorcet paradoxes entirely in real data. I'll suggest two reasons. First, in important elections, candidates must pass through a screening process even to get on the ballot; these processes may involve petitions, or endorsements by established bodies. Candidates low on general excellence may be screened out by those processes, so that those on the ballot differ little on general excellence. Second, Condorcet cycles are presumably noticed primarily when they matter – that is, when they involve the very top candidates. But these candidates may differ less in excellence than the most popular candidates differ from the least popular candidates. These factors may explain why Condorcet paradoxes among the top candidates are as common as they are.

5.3. Which voting systems deal best with the artifacts of Section 5.2?

Darlington (2016, pp. 21-28) ran simulation studies comparing minimax to nine other voting systems, comparing systems on their ability to find the “true Condorcet winner” hidden by the three artifacts just mentioned. Minimax far outperformed all other systems in all the studies. Schulze was by far the best of the competing systems, so for brevity we here summarize the results comparing minimax just to Schulze. In a study on the ability of various voting systems to pick the “true Condorcet winner” who had been hidden by voter carelessness and inattention (see the list in Section 5.2), there were 6833 trials (out of 100,000) on which one of the two methods picked the true winner and the other did not. Minimax was the successful method on 76% of those 6833 trials. On 100,000 trials in another study, in which the true Condorcet winner was hidden by random sampling error rather than voter carelessness, there were 5721 trials on which one of those methods picked the true Condorcet winner and the other did not. Minimax was the successful method in 82% of those 5721 trials. When the true Condorcet winner was hidden by transforming a symmetrically distributed policy dimension to an asymmetric one, there were just 1057 trials (out of 100,000) in which only one of those methods picked the true winner. Minimax was the successful method in 78% of those trials. Thus in all three of these very different studies, minimax beat Schulze by margins that can only be called enormous; as just mentioned, minimax's winning percentages were respectively 76, 82, and 78. In each study Schulze and minimax performed equally in well over 90% of all trials, since either both picked correctly or both failed to do so. But large differences in success rates, always favoring minimax, appeared in the trials on which one method picked correctly and the other did not. Minimax is far simpler than Schulze, eliminating one

possible excuse for Schulze's underperformance. As already mentioned, several other systems did far worse still.

5.4 Two studies on “genuine” Condorcet paradoxes

We now turn to “genuine” Condorcet paradoxes – those produced by particular patterns of voter opinion, not by the artifacts of Section 5.2. Genuine paradoxes seem to require two conditions. I have no data on the topic, but I assume the first condition is reasonably common and the second is quite rare. We might call the first condition “multiple minority positions;” in it, there are three or more candidates, each advocating a different position, with none of those mutually conflicting positions being favored by a majority of voters. For instance, suppose voters in a city must choose among several locations for a park, highway, municipal building, or other project. Or the “positions” might involve the candidates themselves, as when the candidates all come from different ethnicities, religions, or economic groups which the voters perceive differently, so the question is which of those groups the winner will come from.

The second condition is that these positions are favored cyclically, as when position A is favored over B by most voters, B is favored over C by most, and C is favored over A by most. For example, suppose a city has three neighborhoods about equal in size, and a sewage treatment plant must be put into one of them. Suppose that for odd geographic reasons, smells from a plant in neighborhood A will be mostly in A but also in B, from a plant in B will be mostly in B but also in C, and from a plant in C will be mostly in C but also in A. If voters rank these choices selfishly, a Condorcet paradox will appear. Or suppose a city's electorate includes three economic or ethnic groups about equal in size. Voters in each group most prefer candidates in their own group. Voters in group A tolerate group B but detest C, those in B tolerate C but detest A, and those in C tolerate A but detest B. If a mayoral election includes one candidate from each group and voters choose based on group membership, again a Condorcet paradox will appear. These examples are artificial and contrived, and “genuine” Condorcet paradoxes may be quite rare. But we must allow their possibility, and we do that here.

There is no reason to believe that the data patterns in “genuine” Condorcet paradoxes will differ from the patterns in paradoxes produced by random sampling error, so the two simulation studies in this section used data samples generated in the same way as in the previous study on random sampling error. **The first of these studies** used all but four of the 17 voting systems emphasized in Section 2. This study excluded the Copeland system because Darlington (2016) showed that when there is a Condorcet paradox, Copeland produces ties whenever there are just three or four candidates; and in a computer simulation in the same paper, it produced ties on over half the trials whenever the number of candidates was between 5 and 10. All that makes Copeland unacceptable in my opinion. The Dodgson, Young, and Kemeny systems were all excluded because they entail prohibitively complex computations. As mentioned in Section 8.4, that also destroys their transparency, which is also important.

This study used 100,000 trials, each with 10 candidates and 100 voters. On each trial, new candidates and new voters were picked randomly from a standardized bivariate normal distribution. Each candidate's “centrality” was defined as 10 minus the candidate's Euclidean distance from the mean of the 100 voters, and each voter's rating of each candidate was defined as 10 minus the Euclidean distance between that voter and that candidate. These ratings were used directly in the range voting and majority judgment systems, so the winners in these systems were the candidates with the highest

mean or median ratings. In all other systems, these ratings were used to find the voter's ranking or other response to that candidate. In approval voting it was assumed that each voter approved any candidate whose rating by that voter exceeded a weighted average of that voter's highest and lowest ratings, with the highest rating getting 3 times the weight of the lowest rating. By this rule the average voter approved about 3 of the 10 candidates. In the plurality runoff system, the same ratings were used for the runoff, and the runoff was a majority-rules vote between the top two candidates.

Table 3. In the last two columns, systems are arranged in their order of performance on the current test. In each row, the number of trials on which the first system outperformed the second is shown first, then the number on which the second outperformed the first. Every system beat all those below it in the list, by margins which increased with the distance between the systems in this list.

First system beats second	Second system beats first	First system	Second system
6835	2433	Range	Minimax
317	252	Minimax	Black
1916	1287	Black	Nanson
4623	460	Nanson	Schulze
6850	6214	Schulze	Coombs
10813	9737	Coombs	Borda
22320	10253	Borda	Majority judgment
19869	18311	Majority judgment	Bucklin
29913	15577	Bucklin	Approval
34441	26481	Approval	Hare
23871	11628	Hare	Plurality runoff
38682	4411	Plurality runoff	Plurality

One system X was considered to outperform another system Y on a given trial if neither X nor Y produced a tie and X picked a more centrist candidate than Y did. X also outperformed Y if Y produced a tie and X picked some candidate other than the least centrist candidate. A square 13 x 13 table was printed, but is not shown here in its entirety. That table shows the number of trials on which each system outperformed every other system. It turned out to be possible to order the 13 systems so that every system beat every system below it in the order, by a statistically significant margin at the 0.01 level two-tailed. Table 3 shows the sizes of differences between *adjacent* systems in that order. E.g., range voting and minimax were the first two systems in that order. Range voting outperformed minimax on 6835 trials while minimax outperformed range voting on 2433 trials. The fact that range voting beat minimax is of little practical interest, since Section 8.3.2 shows that range voting is extremely susceptible to insincere strategic voting. And this study gave range voting an unrealistic advantage, since the calculations assumed that every voter knew exactly the spatial-model distance between himself or herself and each candidate.

Several points don't appear in Table 3 but should be mentioned. When two systems were far apart in the list, the upper system's margin of victory was enormous. E.g., minimax outperformed plurality voting on 68,271 trials while the opposite occurred on only 2118 trials. Thus, they chose

different winners on 70,389 of the 100,000 trials, and minimax outperformed plurality on 97.0% of those trials. That's an enormous difference.

Also noteworthy is the difference between the Condorcet-consistent (CC) systems (minimax, Black, Nanson, and Schulze) and all other systems except for range voting, which must be dismissed for its extreme susceptibility to strategic voting. The best CC system in Table 3 is minimax, and the best non-CC system is Coombs. Table 3 doesn't show it, but minimax outperformed Coombs on 8059 trials while the reverse occurred on only 2558 trials. That too is a very large difference, especially when comparing two systems which are both "best in class."

Also interesting is that the four voting systems which are probably best known (plurality, plurality runoff, Hare, and approval) are all at the very bottom of the list. It seems that all that work of inventing superior systems in the last century was not in vain.

Because the Hare and Coombs systems have been in the news recently, we give more information about the difference between them. Those two systems picked different winners on 51,380 of the 100,000 trials. The Coombs winner was more centrist than the Hare winner in 44,292 of those trials, while the Hare winner was more centrist in the other 6888. Thus, Coombs picked the more centrist winner in over 6 times as many trials as Hare did.

The second of these two simulation studies was suggested by a remark in Young (1977, p. 350) that, for society to function smoothly, an office-holder should be preferred to all other candidates while they're in office as well as at the time of their election. Thus if there is no Condorcet winner in the election, we might want to pick the person most likely to become a Condorcet winner in the future, even though there might be no formal election at that future time to certify the fact. As mentioned in Section 2, studies of real elections by Gehrlein (2006, pp. 31-58) and Tideman and Plassmann (2012) both found that Condorcet paradoxes appeared in only a small percentage of those elections. That raises the possibility that many of those paradoxes might have disappeared all by themselves if it had been practical to repeat the election a few weeks or months later. First I wanted to see if my artificial studies behaved like real data. In a study with 75 voters and 10 candidates in each trial, sampling from a bivariate normal curve, I found that a Condorcet paradox appeared 263,777 times in ten million trials, or about once in every 38 trials. That mimicked the real data fairly well, so I conducted the following study. This study followed the design in Section 5.3, again with 75 voters and 10 candidates in each trial. In this study, only trials with a Condorcet paradox were analyzed. Ten voting systems were applied to each of those trials. The computer program recorded which candidate was picked by each system on each trial. If a system produced a tie on a given trial, the program recorded that the system had picked no winner on that trial.

After all the voting systems had picked their winners, the computer simulated a second election by moving each voter a small amount in the two-dimensional space. This was done by drawing for each voter an observation from the same sort of standardized bivariate normal distribution already described, except that this new distribution had standard deviations of 0.01 instead of 1.0, so the scores found were much smaller than the initial scores. Like the original scores, these values had means of 0, with about half being positive and half negative. These values were added to the previous spatial position of the voter. Thus if both these new values were negative, the voter was moved slightly down and to the left. If both were positive, the voter was moved slightly up and to the right. If one number

was positive and the other negative, the voter was moved appropriately either toward the upper left or lower right. In this study the 10 candidates in each trial were not moved, though that would be possible.

After all 75 voters were moved as just described, they “voted” again. If these votes produced a Condorcet winner, the identity of that winner was recorded. If there was still no Condorcet winner, all 75 voters were moved again by the same rule. If still no Condorcet winner was found, voters were moved again. A trial was abandoned if a Condorcet winner had not been identified after each voter had been moved five times. About two-thirds of all trials were abandoned as a result. Readers troubled by that fact should be sure to note another study mentioned at the end of this section. In the current study, the program ran until 100,000 second-election Condorcet winners had been found. In both of these studies, a voting system scored a “hit” on any trial in which it picked as winner the candidate who emerged as the Condorcet winner after each voter made the small moves just described. A “miss” means that the system picked a different winner, and a “tie” occurred if the system picked no single winner because two or more candidates were tied for that spot. Misses and ties are both considered “failures,” and appear as such in Table 4.

As seen in Table 4, the study included eight voting systems described in Section 2 and two systems (SSMD and SSSMD) described in Section 7. The study used random tie-breakers for the Hare, Coombs, and Black systems. This elevated somewhat the numbers of both hits and misses for those systems. We shall see shortly that this did not affect the study’s most important conclusions.

The ten voting systems just described are arranged in Table 4 by their numbers of hits. In the table, the number of “disagreements” is the number of trials in which minimax and the other method picked different winners, or one method produced a tie and the other didn’t. That column measures the similarity of a voting system’s results to minimax, not its quality. A disagreement between two systems is considered “relevant” to comparing their quality if one of the two methods picked the ultimate winner; most disagreements were “relevant.” The final column shows the percentage of relevant disagreements in which minimax chose the ultimate winner. The SSMD and SSSMD systems are high on those percentage losses despite being high on hits because they disagreed with minimax less often than most other systems did, even though they lost to minimax on over 90% of those disagreements.

Table 4. Results for 10 voting systems in predicting a future Condorcet winner

System	Hits	Misses	Ties	Failures (miss+tie)	Disagree- ments	% loss to minimax
Minimax	73017	26981	2	26983	-	-
Schulze	67965	21878	10157	32035	12354	76.63
SSMD	51734	5201	43047	48266	43067	90.86
SSSMD	51734	5219	43065	48266	43046	90.83
Black	47901	52099	-	52099	46192	73.65
Coombs	46223	53777	-	53777	51543	75.05
Hare	27507	72493	-	72493	74313	74.12
Plurality	19661	70247	10092	80339	82000	75.75
Approval	10005	89976	19	89995	90012	72.29
Copeland	2057	324	97619	97943	97824	92.28

The most important single fact in Table 4 is that all the entries in the final column exceed 72%. That means that when minimax picked a different winner than that picked by any other system tested, the minimax winner was over 2.5 times as likely as the other candidate to be the ultimate Condorcet winner. Schulze is clearly second best, with far more hits, and far fewer disagreements with minimax, than any of the other systems. This repeated the pattern found in the studies of Section 5.3, in which minimax also beat all other systems by substantial margins, with Schulze coming in second. But Schulze loses to minimax in 77% of the relevant disagreements between them in Table 4, fails the SDVC test of Table 1, and is so complex that many voters would have trouble understanding it.

An alternative study looked at a different set of 100,000 trials, in which voters in every trial were moved until a Condorcet winner appeared, rather than stopping after each voter had been moved five times. In that study the median number of moves for each voter was 14, the maximum was 2939, and 12,023 of the trials had over 100 moves per voter. So the two studies were quite different in the number of voter moves. But in that other study, minimax was also the clear winner.

5.5 Might inventing a Schulze tie-breaker make Schulze outperform minimax?

The study reported in Table 4 used minimax-PM, which includes the tie-breaker described in Section 7. I know of no tie-breaker for Schulze. The results in Table 4 raise the question of whether minimax outperformed Schulze merely because it had a tie-breaker and Schulze didn't, and in Table 4 ties were counted as performing poorly. There seem to be two ways to answer this question without actually creating a new tie-breaker for Schulze. One is to ignore all the trials on which Schulze produced a tie, and to see whether minimax-PM still outperforms Schulze on the remaining trials. That method is actually biased in favor of Schulze, because it's ignoring all the trials which Schulze found to be hardest. A second method is biased even more strongly in favor of Schulze. In that method we look only at the trials in which Schulze picked a winner by avoiding a tie, while minimax used its tie-breaker. Thus we're counting only the trials which minimax found to be especially difficult and Schulze didn't.

I used both of these methods to compare minimax to Schulze. The general design of these studies was like that of other studies in Section 5. Every trial had 10 candidates. I ran one million trials with 75 voters per trial, another million with 100 voters per trial, and another million with 500 voters per trial. As in previous studies, we count only trials in which the two systems picked different winners. Together with the other conditions implied above, this meant we're counting only a tiny fraction of the millions of trials in the total study. Mostly the two systems picked the same winner simply because a trial had a Condorcet winner, so both these systems picked that winner.

In the first study, minimax beat Schulze 467 trials to 406 when there were 75 voters per trial, 259 to 219 with 100 voters, and 17 to 8 with 500 voters. When these are all added together, minimax beat Schulze 743 to 633 across the 3 million trials. In the second study, minimax beat Schulze 447 to 364 with 75 voters per trial, 240 to 196 with 100 voters, and 13 to 6 with 500 voters. When these are added together, minimax beat Schulze 700 to 566 across the 3 million trials. I presume these numbers get smaller with more voters simply because there are fewer ties with more voters. But remarkably, minimax beat Schulze in every one of these six comparisons, despite the aforementioned biases in favor of Schulze. Thus it seems clear that no new tie-breaker could make Schulze outperform minimax.

Surely a lot more mental effort went into inventing the Schulze system than into inventing minimax. So why did minimax turn out so much better? I believe it's because the Schulze system was

carefully designed to avoid the three anomalies of Section 4.5: absolute loser, Condorcet loser, and preference inversion. That constraint forced Schulze to create a system lacking SDVC, whose importance we saw in Sections 4.5 and 4.6. The results of this section are further evidence of the need to accept SDVC and ignore the other anomalies.

6. Strategic voting under minimax

Recently Cornell's Andrew Myers created for me a brief example of minimax's susceptibility to strategic voting and the limits of that susceptibility. That example inspired this much longer discussion of the issue. Any errors are my own.

Myers' example involved a scheme with three candidates in active roles. Such a scheme could exist even if there are more candidates total; the extra candidates would simply have no role in the scheme. So we'll imagine there are just three candidates A, B, C. Consider voters whose sincere voting pattern would be BAC. Suppose they are fairly confident that their candidate B will win their two-way race against C, but they fear that A may beat B and perhaps C. Then A will be a Condorcet winner if A beats C, and A would also win under minimax if A's loss to C is smaller than C's loss to B and B's loss to A. The BAC voters could prevent either of those scenarios if enough of them insincerely vote BCA instead of BAC, thus tipping the AC race toward C without changing either the AB race or the BC race. That could produce a Condorcet paradox in which B has the smallest loss, so B would win. There would be no such thing as "too many" of them voting strategically; B would still win. I ran a computer simulation to see how often strategies like this might succeed. I found that if enough BAC voters vote strategically, they could succeed in this plot in a very noticeable fraction of all trials.

One voting system which is invulnerable to this strategy is Condorcet-Hare. In that system, as in minimax, any Condorcet winner wins. If a Condorcet paradox is found, Condorcet-Hare uses the Hare system to select the winner. Under Hare, the candidate with the fewest first-place ranks is eliminated, the ranks of the remaining candidates are recomputed, and that process is repeated until only one candidate is left. That system is invulnerable to the insincere voting scheme just described because in that scheme the insincere voting occurs only in the lower ranks, and Hare looks only at the top ranks.

The Condorcet-Hare system does have three noteworthy disadvantages relative to minimax. First, Hare doesn't allow tied ranks, and often doesn't allow missing ranks, because as candidates are eliminated one by one, Hare always needs to know which one candidate is each voter's top choice among those remaining. That can make voting much harder for some voters; see the discussion of ease of voting in Section 3. Minimax works just fine with tied and missing ranks. A voter who fails to rank a candidate is presumed to rank that candidate below all others whom they did rank explicitly.

A second problem with Condorcet-Hare is that it violates the SDVC criterion of Sections 3 and 4.5. Consider the example in Table 1. Candidates A, B, and C in that example are tied for the smallest number of top ranks, so in its first two elimination rounds, Condorcet-Hare would eliminate two of them. Then the remaining one would have more top ranks than D, so D would be eliminated. As mentioned in that section, that would still happen if equal numbers were added to all of the example's six pattern frequencies. So Condorcet-Hare would eliminate D even if every candidate except D would need thousands of vote changes to turn them into a Condorcet winner, while D would need only one change.

A third problem with Condorcet-Hare is demonstrated in a computer simulation which studied sincere voting. I found there that in trials in which minimax and Condorcet-Hare picked different winners, minimax was far better at picking the best candidates. This study included 32 blocks of trials, all under separate conditions. The number of competing candidates was 3 or 5 or 10 or 15. The number of voters was 15 or 35 or 75 or 155. The number of opinion dimensions in the spatial model was either 2 or 3. The study included a block of trials for each possible combination of these three factors, thus producing 4·4·2 or 32 blocks. In each block the computer ran until it had found 1000 trials in which minimax and Condorcet-Hare had chosen different winners. It then recorded the number of those 1000 trials in which minimax had picked a superior (more centrist) winner than Condorcet-Hare. Those 32 numbers ranged from 561 to 876, with a mean of 763 and a median of 770. Thus, minimax beat Condorcet-Hare in every single block, and on average its within-block victories were over 3:1. (Version 4 of this arXiv paper reports a smaller study on this topic. A programming error was later found in that study, so it is now being replaced by the current larger study.)

Thus, we want to see whether minimax is so susceptible to strategic voting that it must be rejected despite these three advantages over Condorcet-Hare. We'll continue with this section's opening example, in which voters favoring B are plotting to defeat A through insincere voting, by voting BCA instead of their sincere vote BAC. They hope to make A lose to C by more than B loses to A, so that B wins by having the smallest margin of loss in a Condorcet paradox, provided C also loses to B by a large enough margin.

If the plotters guess wrong and C actually beats B, and their plot makes C beat A as they planned, then C will be a Condorcet winner so the plot gives the plotters their least preferred outcome. In one subset of this scenario, there is a cycle under sincere voting, with C beating B, who beats A, who beats C, but with B winning because they have the smallest loss. Then the plot changes the outcome from the plotters' most preferred outcome (B wins) to their least preferred one (C wins). Or if it turns out that B beats C as they had expected, and B also beats A, then their favorite B would be a Condorcet winner with no plot. So, their plot will have been unnecessary but would do them no harm.

Now suppose none of those cases arises, and all margins of victory are in the direction they anticipated, with $A > B$ and $B > C$ due to sincere voting and $C > A$ due to their plot. It turns out that their plot might still either hurt them or help them, depending on the relative sizes of the three margins of victory. Each line below lists the three margins in the size order assumed in that line, with largest margin first. Since there are three margins, there are six possible orders, as shown. The winner on each line is the last candidate listed on that line, because each candidate wins one race and loses one, so the minimax winner is the loser with the smallest loss – the last candidate in the line.

$A > B, B > C, C > A$: A wins

$B > C, A > B, C > A$: A wins

$B > C, C > A, A > B$: B wins

$C > A, B > C, A > B$: B wins

$A > B, C > A, B > C$: C wins

$C > A, A > B, B > C$: C wins

Recall that the $C > A$ margin is the only one which the plotters hope to affect through their plot. Study of these lines yields the following summary. If the plotters fail to make C's margin over A larger than at

least one of the other two margins, then they will have no effect, leaving A as winner. But they improve the outcome (from their point of view) only if the $B > C$ margin exceeds the $A > B$ margin; their plot doesn't change either of those margins. If the $A > B$ margin exceeds the $B > C$ margin, they hurt themselves by making C win instead of A. So, their plot requires a fairly accurate assessment of the outcomes.

Another point is worth noting. If the election has many voters, then the number of voters in the plot must also be substantial. If this is a public election there is no way the plot will remain secret, even before the election. It may well end up in the newspapers. This will alienate some of the voters who might have sincerely voted for the candidate benefited by the plot, and energize opposing voters to vote. Early and absentee voting will magnify these effects; the plot will then have to be advertised among participants that much earlier, increasing the chance that it will become generally known before Election Day, increasing the plotters' problems just mentioned.

Also noteworthy is that there may be two or more conflicting plots. If A and B are major candidates and C is generally considered a minor candidate, then those favoring B will try to make C beat A as described above, and those favoring A will try to make C beat B. If both groups succeed in those goals, then C will become a Condorcet winner, and both the A and B voters will have seen their least favorite candidate elected. So, if you believe your major opponent will attempt this plot and may well succeed, then your optimum strategy is to execute no plot so that the winner will be your second-favorite candidate instead of your least-favorite candidate.

Still another noteworthy point is the frequency of last-minute surprises in elections, some of which are intentionally held to the last minute by candidates or their campaigns. Even the weather can offer surprises. One group might carefully plan their strategy, then see it disrupted by a last-minute surprise. The US presidential election of 2016 provides an example of both the inaccuracy of prior guesses of the outcome and of last-minute surprises (some of Hillary Clinton's work-related e-mails found on the laptop of a purported sex offender who was not her employee).

All this suggests that plots under minimax are complex and dangerous enough to discourage most of them. Thus, society should try minimax as an experiment. We have already seen that many of these plots won't remain secret even before voting ends, let alone after that. If it becomes obvious that some of these plots have actually succeeded, and voters everywhere start planning to try to imitate those successes despite all the complexities mentioned here, then we can accept that the experiment failed.

7. Six variants and close relatives of minimax, and a tie-breaker

It seems advisable to allow tied ranks on a ballot, since some voters may genuinely have no preference between two candidates. Let the number of "participants" in each two-way race be the number of voters who expressed some preference between those two candidates. As in most ranked-choice voting systems, let's presume that a voter who fails to rank a candidate would rank that candidate below all others whom the voter did explicitly rank. Thus the "nonparticipants" in any two-way race are those who ranked the two candidates equally or omitted both. Since different two-way races will have different numbers of participants, Darlington (2016, pp. 15-16) defined minimax-P as a voting system in which we express each two-way margin of defeat as a proportion of the number of participants in that two-way race, and the winner is the candidate whose largest proportional margin of defeat is smallest. Darlington (2016, p. 16) found that in a simulation study, minimax-P cut the number

of ties to a small fraction of what it had been with classic minimax, but classic minimax picked noticeably “better” (more centrist) candidates than minimax-P when neither system produced a tie. Neither of these results is surprising, the former for obvious reasons and the latter because races with more participants presumably produce more reliable results, and minimax-P ignores that fact.

Darlington (2016) suggested enjoying the advantages of both these systems by using classic minimax as the main system, with minimax-P as a tie-breaker. If minimax-P fails to break the tie, he suggested using each candidate’s second-largest loss, then the proportional form of that loss, and continuing in that way through the results of all of a candidate’s two-way races. Call this minimax-PM since it uses minimax-P and also has Multiple tie-breaking steps. In this system, wins are considered negative losses and are also included in the analysis. Thus in the later tie-breaking stages, X might beat Y not by having a smaller loss but by having a larger win. PM should break nearly all ties. Simulation studies by Darlington (2016, p. 16) show that both the P and M parts of PM do actually result in the selection of better (more centrist) candidates than random tie-breakers.

Darlington (2016, pp. 15-16, 33-35) also studied two other variants of minimax he called minimax-Z and minimax-L. Both are more complex than minimax-P, but neither outperformed classic minimax in simulations, so he did not end up recommending their use.

More recently, I studied a system I’ll call minimax-WV, where WV stands for “winning votes.” In WV the size of any two-way loss is defined not as a margin, but as the number of voters who voted against the loser. Thus each candidate’s LL value is the largest number of voters who voted against that candidate in any two-way race which they lost. If a Condorcet paradox is found, the winner is the candidate with the smallest LL. Like minimax-P, WV may yield a different winner than margin-based minimax when tied or missing ranks are allowed.

My own intuition finds WV troubling. Suppose that in an election with 1000 voters, A lost to B 499 to 500, with one nonparticipant, while C lost to D 0 to 499, with 501 nonparticipants. WV says that A’s loss to B exceeds C’s loss to D, since 500 voters voted against A while only 499 voted against C. That seems ridiculous to me. In one of the most convincing 1000-voter examples I can imagine *against* defining LL as a margin, 0 vote for A over B and 10 do the opposite, while 494 vote for C over D and 506 do the opposite. Then C’s margin of loss exceeds A’s 12 to 10, although A’s margin is highly statistically significant and C’s doesn’t even come close to significance. But my personal intuition says that the first of these two examples of injustice is far worse than the second, so it prefers using margins.

This conclusion is supported by a 10,000-trial simulation study comparing voting systems on the centrist of the candidates they chose. Minimax-WV lost badly to each of the three systems it was compared to: classic minimax using margins (2556 trials to 1397), minimax-P (4455 to 3120), and minimax-Z (4428 to 3079). Margins in minimax-Z are defined so their size corresponds very closely to their statistical significance. But simulations found minimax-Z to be slightly inferior to simple minimax at picking the best candidates. That suggests the previous paragraph’s second example can be ignored entirely.

Darlington (2016, p. 26) studied two other methods similar to minimax. In the method he named SSMD, the winner is the candidate with the Smallest Sum of Margins of Defeat, and in method SSSMD it is the candidate with the Smallest Sum of Squared Margins of Defeat. Using analyses like those in Section 5, Darlington compared both SSMD and SSSMD to minimax on their ability to find the “true

Condorcet winner” whose identity had been hidden by either (a) random sampling error or (b) voter misperception of candidate positions. In each of these 2-2 or 4 comparisons, minimax’s winning percentages were 96 or higher when either minimax or its competitor, but not both, picked correctly.

Thus the best system known appears to be minimax-PM, which uses margins and a multi-stage tie-breaker.

8. More on the disadvantages of various systems

8.1 Why “simpler” ballots can actually make voting harder

Plurality elections suffer from vote-splitting and spoilers, and elections with runoffs certainly don’t simplify the voting process. So arguments for “simplicity” tend to focus on approval voting. But approval ballots may make voting easier than ranked-choice ballots for the least-informed voters but more difficult for better-informed ones. If a voter ranks the candidates ABCD, the approval ballot forces them to also choose where to draw the line between “approved” and “unapproved” candidates. To do this rationally, the voter should consider whether the difference in merit between A and B exceeds that between B and C or between C and D. And for maximum impact a voter will also want to draw their line between the two candidates most likely to receive the most approvals from other voters, and therefore must guess who those two are. Thus a voter using an approval ballot must make three types of judgment not needed in a ranked-choice ballot. First, they must assess the relative sizes of differences in merit between various pairs of candidates. Second, they must guess which candidates are most likely to win. Third, they must decide how to combine these two judgments into a final ballot choice.

We should also distinguish between *requiring* and *allowing* voters to express more detail about their preferences. If a ranked-choice ballot allows tied and missing ranks (as most CC systems do), a voter can choose the complexity of their response. A voter could imitate plurality voting by putting one candidate at the top and ignoring all others. Or the voter could imitate approval voting by placing two or more candidates at the top and ignoring all others. Or they could tailor the complexity of their response in other ways, for instance by ranking their top three candidates and ignoring all others. Or they could vote *against* one or more candidates, without expressing preferences among the rest, by putting all the rest at the top but omitting the objects of their special disdain. People can be informed about all these possibilities in schools or even through television entertainment. All these ways of voting, and others, are allowed by the Condorcet Internet Voting Service mentioned in Section 4.5. Darlington (2016, p. 14) shows how paper ranked-choice ballots can also allow a single voter to add one or more write-ins, while still allowing ballots without write-ins to be read by machine.

8.2 Problems with positional voting systems

The Borda positional system allows tied ranks but the basic Hare and Coombs systems do not, because ties could prevent us from determining each voter’s first or last choice. Other positional systems also differ from each other in this respect. There is no need to elaborate here on that point, since we argue here that all positional systems are inferior to CC systems. For simplicity we assume throughout Section 8.2 that there are no tied ranks.

8.2.1 Examples of strategic voting under Borda and Coombs

As illustrated by the Hare, Coombs, and Borda systems respectively, positional systems typically

emphasize a candidate's number of high ranks, or their number of low ranks, or some statistic similar to the candidate's mean rank. The discussion of Burlington VT in Section 3 illustrated the problem with emphasizing the number of high ranks: this number is excessively influenced by the vote-splitting problem, in which even a Condorcet winner could get few top ranks because other candidates are so similar to them. But we now show that systems which emphasize average or low ranks are very susceptible to strategic maneuvers. To "bury" a candidate is to try to defeat them by insincerely giving them very low ranks. I'll call the opposite maneuver "strategic elevation." That name is my own, but the concept is very familiar to voters who use plurality voting. If a plurality voter regards candidate X highly but not at the very top, but thinks X has the best chance to beat candidates the voter really dislikes, they might strategically elevate X to the very top to increase the chance of defeating their least favorite candidates.

Section 6 showed that burying is rarely a real problem under minimax. But under Borda, suppose a voter would sincerely rank four candidates in the order ABCD, but they believe that B and C are by far the two candidates most likely to win. They prefer B to C, so they might choose to rank them as BADC on their ballot, thus strategically elevating B and burying C. Let m_{XY} denote the difference between the Borda counts of candidates X and Y, with m_{XY} being negative if X loses to Y. Then each step by which a voter elevates X or buries Y increases m_{XY} by 1. Thus the more candidates there are, the more effective strategic voting might be. If c is the number of candidates, a single voter voting strategically might increase m_{XY} by the same amount as adding $(c - 2)$ voters who put X one step before Y.

Thus even a landslide could be tipped by just a few strategic voters if c is high enough. For instance, suppose everyone agrees A and B are the two best candidates, and 55% prefer A to B while 45% do the opposite. Thus A wins by a landslide under a common definition of that term. But suppose there are 4 minor candidates so $c - 2 = 4$, and 6% of the B-voters bury A. That's 2.7% of all the voters, so it's as if new B-voters had appeared in numbers equal to $4 \cdot 2.7\%$ or 10.8% of the original voters, thus tipping the election to B. The possibility of strategic voting also makes it far more difficult for every voter to plan how to vote most effectively, since they must try to guess each candidate's chance of winning. They will feel obligated to make those guesses, since they will assume the opposing voters are doing so.

The Coombs system is most susceptible to insincere strategic voting when there are several minor candidates who are all about equally popular because they're all little known. However, major-party strategists could arrange for those candidates to appear, in order to make their strategies work. Consider the case in which there are two major candidates A and B, plus ten minor candidates little known to many voters. Suppose for simplicity that there are 1000 voters, and both A and B are sincerely preferred to all minor candidates by every voter. Voters favoring A over B will correctly perceive B as the major threat to their candidate, and those favoring B will do the same for A. Suppose first that A is preferred to B by 900 voters, and B to A by the other 100, but A's voters all vote sincerely while B's all vote strategically by ranking A at the very bottom, so A gets 100 last-place ranks. If the 900 sincere voters all put the minor candidates in random order, then each minor candidate will get about $900/10$ or 90 last-place ranks. Then A will likely be the first candidate eliminated, making B win despite A's enormous lead over B in sincere preferences.

Now suppose that things are as above, except that at least 100 voters prefer each of the major candidates, and at least 100 voters in each group vote strategically by ranking their main opponent at the very bottom. Each minor candidate will almost certainly get fewer than 100 last-place ranks since

there are now just 800 or fewer sincere voters. Thus, A and B will likely be the first two candidates eliminated by Coombs, making the winner be someone who was not sincerely preferred to either A or B by even one voter.

8.2.2 *Is strategic voting a real danger?*

Some electoral theorists have argued that there is little danger of substantial insincere voting. But these examples show that elections can sometimes be tipped even by small numbers of insincere voters, and Section 8.3.2 gives even worse examples. In a world in which newspapers regularly discuss charges of gerrymandering, “fake news” designed to change votes, and legislatures dominated by one party changing the voting rules to make it harder for opposing voters to vote, it seems naive to design or choose voting systems on the assumption that strategic voting will never be a problem. Campaign workers could even encourage their voters to vote strategically, claiming that other parties could do it or have done it. Suppose a magazine article were to show that around the world there were several instances in which later analysis suggested that strategic voting might well have tipped an election, since a noncontroversial and reasonably qualified candidate had lost under Borda because they got many ranks at the very bottom. That article could lead to a worldwide increase in the frequency of strategic voting, which would lead to more such articles, etc. Widespread strategic voting could lead to an even broader feeling that society and government are corrupt and need to be overthrown, violently if necessary. All that can be avoided by using voting systems as resistant as possible to strategic voting.

It might be argued that strategic voting could be moderately common, but still wouldn't matter because all sides would do it equally. But voters for one candidate might be more inclined than others to vote strategically if (a) they fear they'll come in second if they vote sincerely, or (b) their candidate offers them some major benefit such as a tax break, or (c) their major opponent favors idealistic social programs opposed by the most selfish and Machiavellian voters, who are most likely to vote strategically, or (d) all three of those.

Borda would also encourage major political parties to employ a strategic maneuver which seems odd today because it would actually be counterproductive under plurality voting. Again suppose the two major candidates are A and B. B's party could identify political aspirants much like B ideologically but far less experienced and less known. The party could encourage them to run as independents, perhaps even supporting them financially. Then nearly all sincere voters in both parties would rank each of those candidates below B, but many voters in B's party would sincerely rank them ahead of A because of their ideology. That would change the Borda-count A-B difference in the direction favoring B. Thus there are actually two ways to bury a candidate: persuade your voters to insincerely put them at the bottom, or add minor candidates in such a way that your voters will do that simply by voting sincerely. If this were done often, it would make elections more burdensome for everyone by adding minor candidates.

8.2.3 *Two Borda-based attacks on CC systems*

Are there arguments for positional systems? An anonymous reviewer for a very prestigious academic journal used two examples to attack CC systems for conflicting with positional reasoning. But closer examination shows that both examples actually support CC.

In the first example, 1001 voters rank 6 candidates in the order ABCDEF while the other 1000 voters rank them BCDEFA. Thus A is the Condorcet winner, winning every two-way race by just one vote.

The reviewer considered it obvious that B should win, because every voter ranks B first or second, while nearly half the voters rank A last. B would win under Borda, or under a positional system in which the winner is the candidate placed in the top two ranks by the most voters. But the votes in the example are exactly what we would expect if C, D, E, and F were minor candidates sincerely ranked below A and B by all voters, and the B-voters all chose to bury A while the A-voters voted sincerely. Or even without strategic voting, this example could appear if votes were determined by a liberal-conservative scale like that mentioned in the discussion of Burlington VT in Section 3, with voters distributed symmetrically on that scale, and with every voter ranking the candidates in the order of their closeness to themselves. Suppose candidates B-F are all close together, in that order with B closest to the mean, with A on the opposite side of the mean and just a bit closer to the mean than any of the others. For instance, suppose the mean is 5, A is at 4, and B-F are at 6.01, 6.02, 6.03, 6.04, and 6.05 respectively. This would produce the central features of the current example, with A beating every other candidate by very small margins but being ranked last by almost half the voters. But A should win since they're the most centrist candidate. This is actually the type of situation described in the last paragraph of Section 8.2.2.

Anyone who finds these arguments unconvincing should show how else the example might arise. Also, A is the absolute majority winner, with over half of the first-place ranks. As mentioned by Felsenthal (2012, p. 21), failure to make any such person a winner is widely considered to be a serious flaw in a voting system. A also wins under the Hare and Coombs positional systems. That makes five reasons to deny that B "should" win: strategic voting, the liberal-conservative model, the absolute majority rule, and the Hare and Coombs systems.

In the reviewer's second example, 240 voters rank three candidates as ABC, 261 give ACB, 300 give BAC, and 200 give CBA. Calculations show that among these 1001 voters, A beats B by one vote, A beats C by 601, and B beats C by 79. Thus A is the Condorcet winner. A also wins under the Borda, Hare, and Coombs positional systems. Then 6 more voters appear, with two showing ACB, two showing CBA, and two showing BAC. By positional thinking, the three candidates are all tied among these six new voters, since each candidate appears twice in each of the three ranks. But adding these 6 voters to the previous 1001 changes the Condorcet winner from A to B, since B now beats A by one vote. A remains the Borda winner by a wide margin. The reviewer felt this exposes CC as ridiculous, since a set of votes showing a tie shouldn't change an election's result. That reasoning would also dismiss the Hare and Coombs positional systems, since under those systems the six new voters also switch the winner from A to B. My own guess is that if a neutral judge were asked to use common sense to choose a winner for this one case, without trying to specify a rule for all cases, they would say that C is an obvious loser, since C loses to both A and B by large margins. Thus we should choose between A and B by majority rule, so the six new voters should make B win by one vote although A should win without them. It's hard to see why that is ridiculous.

Since a reviewer for a very prominent journal presented these two examples as their main arguments for preferring positional systems over CC systems, and they were accepted as sound arguments by the journal's expert editor, I infer that there is a severe shortage of convincing arguments for that viewpoint. In a computer simulation study, Darlington (2016, p. 28) did find that Borda outperformed minimax in one voting model, but this effect was reversed under other models. And the arguments here against positional systems seem quite telling.

8.3. Problems with rating-scale voting systems

8.3.1 Introduction and summary on rating-scale systems

Section 8.3 is quite long, so I begin with a summary. Busy readers convinced by Section 8.3.2 might choose to skip Sections 8.3.3 and 8.3.4, though I believe they are also quite convincing.

The two best-known rating-scale voting systems are majority judgment (MJ) and range voting (RV). I criticize all rating-scale systems but focus primarily on these two. Insincere strategic voting is the biggest problem with all these systems. The strategic-voting example of Section 8.3.2 uses artificial data, but data which was designed to simulate real votes rather than to make MJ and RV look as bad as possible. Nevertheless, the example shows landslide elections being tipped if only 1.4% of voters vote strategically under MJ or only 1% under RV. Sections 8.3.3 and 8.3.4 show that under both these systems but especially under MJ, there are persuasive arguments which campaign workers can use to persuade their voters to vote insincerely by giving every candidate either the highest or lowest possible rating. I'll call that the *max-and-min* voting strategy. Even without strategic voting, Section 8.3.3 (on MJ) and Section 8.3.4 (on RV) both give examples in which B wins even though 98 of the 99 voters preferred A. I argue that cases much like these could arise in real life.

RV uses a numeric scale, often from 0 to 100, and the candidate with the highest mean wins. MJ uses a scale with non-numeric verbal labels such as "excellent" or "poor," and the candidate with the highest median wins. If the number of voters is even, the highest rating in the lower half is used as the median. If there is a tie between candidates for highest median, the ratings for each tied candidate are sorted from high to low. These sorted columns can be assembled into a matrix. In that matrix we find the row nearest the median row in which one candidate beats all others, and that candidate is the winner. If two such rows are equidistant from the median row, the one below the median row is used. These rules will break all ties unless two candidates have absolutely identical distributions of ratings.

Arguments for RV appear most fully at rangevoting.org. MJ was first published in 2007, but it is described and defended most fully by its creators Balinski and Laraki (2010). The book jacket can be seen online via Google Scholar, and is covered with effusive praise from Nobel laureates and other prominent electoral theorists. Aside from MJ's many flaws, that's the main reason Section 8.3 is so long.

Rating-scale systems are often defended as the only systems satisfying the subset choice criterion (SCC), also known as independence of irrelevant alternatives (IIA). Section 4.1 described SCC/IIA and Section 4.3 showed why it can be dismissed; it is not mentioned again in Section 8.

MJ and RV are the only well-known voting systems which do not reduce to majority rule (MR) in two-candidate races. Those races are free of most of the paradoxes and anomalies afflicting multi-candidate races. Therefore two-candidate races offer the simplest and clearest illustrations of the faults of MJ and RV. Thus we focus primarily on those races. After all, if a voting system doesn't behave reasonably even in two-candidate races, surely we shouldn't trust it to do so with more candidates – especially since a multi-candidate race can be thought of as a set of two-candidate races.

8.3.2 Using MJ or RV, how many insincere voters are needed to tip a landslide two-candidate election?

I used a one-dimensional spatial model to study this question. These models are described in Sections 1 and 5.1. I divided a standard normal distribution z into 1000 equal-area sections, and I placed

an artificial voter at each of the 999 borders between adjacent sections, thus simulating a normal distribution as closely as possible with 999 voters. I placed candidate A at the distribution's mean, with a score of 0. Thus A is the "perfect" candidate in this model since they're the most centrist possible candidate. I placed candidate B at $z = 0.257$. That puts 601 of the 999 voters to the left of B, so B is noticeably but not extremely off center. These placements made 551 of the 999 voters closer to A than to B, with the other 448 closer to B. A spatial model assumes every voter prefers the candidate closest to themselves, so the model has A's margin of victory under sincere MR voting at $551 - 448$ or 103. That's over 10% of the voters, so with sincere voting the election is a landslide under a common definition of that term. I defined each voter's rating of each candidate as 10 minus the absolute distance between that voter and that candidate. A would win under either RV or MJ as well as under MR, since A's mean and median ratings both exceed B's.

But then, starting with the most extreme voters at the end of the distribution in which all voters favored B, I changed the rating of each voter, one at a time, to max-and-min ratings. Specifically, I had each of these voters now give A the lowest rating sincerely given to either candidate by any voter, and each give B the highest such rating. I applied the changes to the voters at the far end of the normal distribution because I judged that in real life these would be the voters who could most easily be persuaded to vote insincerely to increase B's chance of victory. Those at the extremes are often alienated from society, and thus could probably be easily persuaded to lie a bit to advance their own personal agendas. After making these changes for each voter (one voter at a time), I recomputed the two median ratings to study MJ, and stopped when B's median rose above A's. This took just 14 voters. That's just 3.1% of the 448 voters favoring B, and 1.4% of all voters. It's hard to imagine that B's campaign workers couldn't persuade that number of voters to vote strategically. When I did the same analysis using means rather than medians, thus assessing RV instead of MJ, it took only 9 insincere voters to tip the election. That's about 2% of the 448 voters preferring B, and under 1% of all voters. Thus RV is even more susceptible to strategic voting than MJ, though both are unacceptably susceptible.

Note that in both MJ and RV, all these voters were counted as rating B above A even under sincere voting; the only switch was to insincere ratings. The MR result would not be affected at all by the insincere votes. That is, if vote-counters were given MJ or RV ratings and then applied MR after noting merely which candidate each voter preferred, they would find the same winner regardless of whether voters had voted sincerely or insincerely, and the same winner as in a poll with ordinary MR ballots.

For those interested, here is more detail concerning this example. If a voter changes their rating of candidate X, it doesn't change X's median at all if the rating stays on the same side of the median. Of the 448 voters favoring B, the most centrist 198 rate both A and B above their own medians. If some of these 198 voters switch to max-and-min, under MJ it will help B by lowering A's median, but it will not raise B's. Under MJ the next most centrist 80 of those 448 voters cannot help B at all through insincere voting, since they already rate A below A's median and rate B above theirs. The 170 least centrist of the 448 B-voters rate both A and B below their own medians. If some of these 170 voters switch to max-and-min, under MJ it will help B by raising B's median, but it will not lower A's. Thus under MJ there are "only" $198 + 170$ or 368 voters favoring B who can actually help B win through insincere voting. It turns out that the 198 most centrist of these 368 voters have slightly less power to tip the election than the more extreme 170, since the tip would require 16 of those 198, not 14. I consider these to be details which weaken the example's main point only slightly. A similar analysis for RV shows that the earlier

paragraphs actually slightly *understate* the power of insincere voting; under RV any 9 of the 448 voters favoring B could tip the election by using max-and-min, and some sets of 8 voters could do so.

MJ's defenders might point out that MJ typically uses a rating scale with only 5 to 8 points, while this example used a scale with infinitely fine gradations. Thus in this example there is more "room" than in a typical MJ for an insincere voter to assign a rating at the very top or bottom of all ratings. But that doesn't matter under MJ, since any rating above a median will raise that median as much as a rating at the very top of the scale, and a similar effect applies at the bottom of the scale. Thus the example's criticism of MJ applies to any form of MJ as long as the extreme points on the MJ scale were above and below both candidate medians – a condition which is presumably nearly always met. When this example was modified to use a six-point rating scale, it still required only 14 insincere voters to tip the election.

This example focused on MJ and RV, but any two-candidate voting system not equivalent to majority rule (MR) must be susceptible to strategic voting. MR counts equally all voters favoring a particular candidate, so any system not equivalent to MR must count some of them unequally. But when voters understand that some votes count more than others, they will have every incentive to insincerely give the response which they know will maximize their vote's impact.

It also seems almost inevitable that voters will ultimately respond to these incentives. Suppose a two-candidate election is run under some system other than MR, but for general interest or because some people demanded it, authorities also report the number of voters who rated each candidate above the other. Suppose voters then see that B won even though more voters favored A. That means that those favoring B had in some sense voted more effectively than the others. Those favoring A will immediately doubt the sincerity of the B-voters, so in the next election the A-voters will almost certainly vote insincerely themselves. Soon everyone will be using the insincere max-and-min strategy. Thus all two-candidate elections will reduce to MR. And the need for insincere voting – a need created by the failure to use MR – will promote a general attitude of cynicism about politics and distrust of the government. Nationwide or worldwide publicity could make these adverse effects widespread even if they were triggered by just a few local elections.

It seems clear that any system comparing measures of central tendency, as MJ and RV do, will be at least as vulnerable to insincere voting as MJ is. Strategic voters use extreme scores, and measures of central tendency differ in the amount they are influenced by those scores. E.g., a group's midrange is the mean of the highest and lowest scores in the group; it's obviously heavily influenced by extreme scores. Of all measures of central tendency, the median is the one least influenced by extreme scores. As already mentioned, if we change any score other than the original median score, while keeping the changed score on the same side of the median, it will not change the median at all. Thus it seems likely that MJ will be less vulnerable to insincere voting than systems which use almost any other measure of central tendency. But we have already seen that MJ itself is unacceptably vulnerable to insincere voting.

8.3.3 Why MJ almost forces voters to vote insincerely

I just argued that MJ may require more people voting insincerely to tip an election than any other rating-scale system. But this section shows why under MJ, it may be especially easy to persuade voters to do that. For simplicity this section continues to focus on two-candidate elections. Any candidate will prefer that all their voters use the max-and-min strategy, so campaign workers will try to

persuade voters to do so. This section describes three arguments which seem particularly persuasive for MJ elections.

The “post-election impact” argument is aimed at voters who favor a particular candidate because they favor that candidate’s policy positions rather than their good looks or magnetic personalities. Whether the office won is legislative or executive, a “big win” will increase the respect an office-holder receives, and it will thus help them implement their policies better than a smaller win. Therefore voters will want not merely to elect their favored candidate, but to make that candidate’s win as impressive as possible. Under MJ, news media will presumably publish the number of voters rating each candidate at each level of the rating scale, and politicians will study those figures carefully. Thus the bigger the difference between the winner’s figures and the loser’s, the easier the winner will find it to implement their policies. In fact, if A beat B without voter X, and X rated A just slightly above B, adding X’s ratings to the data set might actually lower the difference between the two sets of ratings and thus make it harder for A to implement their policies once in office, much as if X had preferred B. X doesn’t want that, and the surest and easiest way for X to prevent it is to use max-and-min.

The “no impact” and “reverse impact” arguments for using max-and-min define a “single-sided” voter as one who rates both candidates above the medians of both, or rates both candidates below both medians. In a spatial model, single-sided voters will be especially numerous if the two candidates are close together in space, since voters close to both will then rate both above both medians, and voters far from both will rate both below. In a simulation study in which the spatial positions of both voters and candidates were picked randomly from a bivariate normal distribution, on the average about 40% of all voters were single-sided, and that frequency approached 99% in trials in which the two candidates were especially close to each other. That study is described more fully at the end of this subsection.

The “no impact” argument points out that if a single-sided voter were to switch their ratings, giving candidate B the rating they had previously given A and vice versa, it could not possibly change the winner because neither of the two medians would change at all. So their preference between candidates has no effect at all on who wins. The “reverse impact” argument notes that a single-sided voter’s decision to vote could actually make their preferred candidate lose. For instance, suppose 4 voters voted before X did, and their ratings were 1 2 5 6 for candidate A and 1 3 4 6 for B. When the number of voters is even, MJ uses as a candidate’s median the lower of their two central ratings. Thus in this example, B was winning because the calculated medians were 2 for A and 3 for B. Suppose X then votes, rating A at 5 and B at 6. This moves the medians up to 5 and 4 respectively, making A win even though X prefers B and B was winning before X voted. This illustrates MJ’s nonmonotonicity, which was mentioned in Section 3.

In one sense the “reverse impact” argument is weaker than the other two, because it requires an exceptionally large gap between adjacent ratings for one candidate. But the real question is whether these arguments could be used to persuade voters to use max-and-min, and I think all three could be so used. Thus a campaign worker might say, “If you don’t use max-and-min, your vote might not really be counted, or might lower the published win size of your favored candidate, or might even make them lose. And max-and-min actually requires less hard thinking than the alternative.”

The anomalies of single-sidedness can be pushed to an extreme. For instance, suppose ratings are on a scale from 1 to 6, with 6 high. Suppose 49 voters rate A at 2 and B at 1, 49 others rate A at 6 and B at 5, and one voter rates A at 3 and B at 4. Then A’s median is 3 and B’s is 4, so B wins under MJ

even though 98 of the 99 voters prefer A to B. Those 98 voters are all single-sided, so their preferences don't count. Balinski and Laraki (2010, pp. 328-329) dismiss examples like this as contrived and unlikely to occur in real life. But in the simulations mentioned above and described more fully in the next few paragraphs, the number of voters whose preferences didn't count did approach 99% in some trials designed to simulate real elections. And this paragraph's 99-voter example merely extends to extremes the problems for single-sided voters which affect virtually every use of MJ. Note that the anomaly in this paragraph does not require insincere voting.

The aforementioned computer simulation was as follows. I used a two-dimensional spatial model with a bivariate normal distribution of voters. I drew a random sample of 10,000 voters and used that same sample for all of the 100,000 trials about to be described, since I assumed that single set of 10,000 voters provided an adequate approximation to a bivariate normal population. Then I randomly drew 100,000 pairs of candidates from the same bivariate distribution; each such pair constituted a new trial. In each trial I computed the Euclidian distance between each candidate and each voter, and defined the voter's rating of that candidate as 10 minus that distance. I then computed each candidate's median rating, and used those medians to label each voter as single-sided or not. The mean number of single-sided voters per trial (out of 10,000 voters in the trial) was 4004.7, and the median was 3929. Thus we can say that on average, about 40% of voters were single-sided. But as mentioned four paragraphs ago, those numbers were higher when two candidates were similar and lower when they were dissimilar. Thus the number of single-sided voters in a trial ranged from 0 to 9889, so there were trials in which about 99% of voters were single-sided.

I also defined a voter's "preference intensity" as the absolute value of that voter's difference between the ratings of the two candidates. On average the intensity of single-sided voters was about 66% of the intensity of other voters. Across all trials, about 32% of the total voter intensity was found in single-sided voters. But there were trials in which over 99% of the total voter intensity was in single-sided voters, since that percentage correlated -0.81 with the Euclidian distance between the two candidates and those distances ranged from under 0.01 to 6.87.

In the 100,000 trials, there were just 1128 trials in which MJ and MR (majority rule) picked different winners. Each system picked the more centrist of the two candidates on over 90% of the trials. But when the two systems picked different winners, MR picked the more centrist candidate on 840 trials and MJ did so only on 288. Thus MR "won" on 74% of these trials. That difference is enough to persuade a neutral outsider that MR is superior, even assuming everyone votes sincerely. It also turned out that the more single-sided voters there were in a trial, the more likely MJ was to pick the "wrong" (less centrist) candidate in that trial. And the results in this subsection seem more than enough to persuade many voters to vote insincerely in order to avoid being classified as single-sided by MJ.

Darlington (2017) describes several other simulation studies showing how poor MJ is at picking the best candidates, even under sincere voting.

8.3.4 Like MJ, RV can totally ignore majority rule

Suppose we have a rating scale from 0 to 100. Suppose there are 99 voters, and 98 of them rate A one point above B, but the 99th voter rates A at 0 and B at 100. Then B's rating total across voters exceeds A's by two points. Thus B wins by RV although 98 of the 99 voters preferred A to B. B might seem like a reasonable winner in this example until we consider the possibility of max-and-min strategic

voting. A single strategic voter who favors B would vote exactly as in this illustration, thus tipping an election in which every other voter prefers A.

Campaign workers would not find it as easy under RV as under MJ to persuade their voters that they *should* vote strategically, since there is no obvious sense in which RV ignores anyone's sincere vote. But other arguments for strategic voting still apply, such as the desire to make your candidate's win as big as possible, and the fear that opposing voters may vote strategically and may have done so in the past. And as we saw in Section 8.3.2, it usually takes fewer insincere voters to tip an election under RV than under MJ. One reason for this is that if a voter favoring B sincerely rates B above B's median and sincerely rates A below A's median, strategic voting by that voter will not help B at all under MJ. But under RV every strategic vote helps the favored candidate, except for the rare voter who sincerely rates their favored candidate at the very top and their opponent at the very bottom. Thus overall, it's hard to argue that RV is more resistant than MJ to strategic voting.

8.4 Problems with the Dodgson and Young systems

Aside from minimax, Dodgson and Young are the only well-known voting systems which satisfy both CC and SDVC. This section describes their limitations. I know of no countervailing advantages of these systems over minimax.

Oxford mathematics lecturer Charles Dodgson is better known as Lewis Carroll, the author of *Alice's Adventures in Wonderland*. In a pamphlet printed in 1876 and reprinted in Black (1958), Dodgson considered a voter who interchanges two candidates whom the voter had previously ranked adjacently, as when the ABCD ranking is changed to ACBD or ABDC. Reversing ABCD to DCBA would require six such interchanges. Dodgson suggested naming as winner the candidate who would need the smallest number of such interchanges, summed across voters, to become a Condorcet winner. In a different approach, Young (1977) considered voters who help defeat a candidate by ranking them poorly. He suggested counting for each candidate X the smallest number of such voters who would have to be deleted to turn X into a Condorcet winner, and naming as winner the candidate for whom this number was smallest.

Perez (2001) showed that both the Dodgson and Young systems lack monotonicity. The important writers Felsenthal and Nurmi (2017) consider monotonicity essential, and I agree. Nonmonotonicity also makes a system susceptible to strategic voting, and Green-Armytage (2014) found minimax to be one of the systems most resistant to strategic voting.

The conceptual simplicity of the Dodgson and Young systems also masks substantial computational difficulty. Hemaspaandra, Hemaspaandra, and Rothe (1997) show that the Dodgson system is so difficult it can be a challenge even for modern computers. Rothe, Spakowski, and Vogel (2003) reach a similar conclusion for the Young system. Even if those difficulties can be overcome, they raise an important issue of transparency. A voting system is transparent if the average citizen can see that it would be difficult or impossible for corrupt insiders to alter the election's results. But under Dodgson or Young, there is no way an average citizen could understand or check the necessary computer programs, meaning they would have to accept their conclusions on faith. That can be a major political problem when citizens don't trust governmental institutions, as is widely true. In contrast, minimax would use a computer program which counts the number of voters who rank one candidate X above some other candidate Y, and does that for every pair of candidates. That program would be so simple that the program writers would have no plausible claim that the program must

remain confidential to protect trade secrets, and even a small town would contain people who could check the honesty of such a program. Once that step is done, the remaining calculations in minimax are so simple they could easily be done by hand by almost anyone: we compute all the margins of victory and defeat, we find each candidate's largest margin of defeat, and name as winner the candidate for whom this margin is smallest.

Some European governments have attempted to maximize transparency by publishing replicas of all ballots. But according to Naish (2013), the Italian mafia appears to have devised a way to use this information to force many voters to vote as the mafia wishes, at least in elections with many candidates. They can direct the voter to put the mafia candidate first, but direct them to use the next few ranks to endorse a set of candidates so opposed to each other that no sincere voter would ever choose that set. Those ranks constitute the voter's "signature" which allows the mafia to verify that the voter put their candidate first. We have just seen that minimax can achieve a very high level of transparency without sacrificing voter privacy. But that's not true for the Dodgson and Young systems, since an outsider couldn't really check every step without seeing each voter's actual ranking of candidates, thus running into the Naish problem. The minimax-PM tie-breaker described in Section 7 uses the same vote totals used by simple minimax, and thus constitutes no greater threat to voter privacy.

The Dodgson system can also make it easy to "bury" a candidate – defeat them by having voters insincerely put them at or near the bottom of their rankings. For instance, suppose we have 25 voters and 6 candidates A-F, with the following circumstances:

1. All voters perceive A, B, and C as the best candidates and the ones most likely to win. Under sincere voting everyone would put D, E, and F at the bottom in that order.
2. It's generally agreed that the most likely winners are A, B, and C in that order. Thus voters favoring C will try to bury A and B, and voters favoring B will try to bury A.
3. A Condorcet paradox appears among these three candidates, with A beating B in their two-way race, B beating C, and C beating A. A Condorcet paradox is an essential feature of any example comparing CC systems, since without the paradox, all those systems would pick the same winner.

Consistent with all these specifications, we have:

Pattern	Frequency
ABCDEF	10
BCDEFA	8
CDEFAB	7

The numbers of Dodgson interchanges needed to make A, B, or C a Condorcet winner are 12, 5, and 6 respectively. Thus A loses big, with at least twice as many interchanges needed as for either B or C. But under minimax the largest two-way losses of the six candidates, in alphabetical order, are 5, 9, 11, 25, 25, and 25, so A beats even B and C by wide margins relative to the number of voters. This difference in outcomes is caused entirely by the presence of candidates D, E, and F and the burying strategies used against A and B. But D, E, and F are all losers by any reasonable voting system, since they are all ranked below C by every single voter. If we remove them from the race, the largest two-way losses of A, B, and C remain the same, so A still wins under minimax. But their Dodgson values become 3, 5, and 7 respectively, so A now wins under Dodgson as well. With D, E, and F gone, A also wins under the Borda,

Hare, Coombs, plurality, and many other systems, as well as under minimax and Dodgson. It thus seems clear that A is the best candidate, though Dodgson put A far behind B and C when all 6 candidates were included. And this is with just three minor candidates (candidates who would lose big by virtually any rule). The more minor candidates there are, the more “dirt” each insincere voter has to pile onto any candidate they wish to bury, and the fewer such voters it takes to bury a candidate.

As mentioned above, Young (1977) proposed naming as winner the candidate with the smallest number of voters who would have to be removed to turn them into a Condorcet winner. He commented (p. 350) that the Young and minimax systems are probably very similar, and I agree. But if anything, minimax is superior to Young when we think more carefully about how votes would need to change to turn a candidate into a Condorcet winner. Since a Condorcet winner is defined in terms of two-way races, consider the three ways a voter could act more favorably toward B in a two-way race between A and B which A had won. (1) The voter could abstain when they had previously favored A. (2) They could favor B when they had previously abstained. (3) They could favor B when they had previously favored A. There are three parallel ways they could act more favorably toward A than they had before, making 6 ways a voter could change. We would ideally like to consider the effects of all of those possible changes. If we did so, we would presumably want to add them up, presumably giving (3) twice the weight given to (1) and (2). The three parallel ways of changing toward A would be given negative weights, with the third again being given twice the weight of the others. The Young system considers just the first of those six types of possible change. One’s first thought is that it would be wonderful but completely impractical to consider all six of these types of change. But surprisingly, it’s actually much easier to consider all six of these together, using the additive rule just described, than to implement the Young system. That’s because the additive rule simply yields the change in B’s margin of defeat against A. Thus we can effectively implement that rule by using minimax, in which the winner is the candidate whose margins of defeat have the smallest maximum.

We can also say that the minimax winner is the candidate X who would need the fewest new voters, all putting X first, to transform X into a Condorcet winner, since that number is always just one greater than X’s largest two-way margin of defeat. But that statement about new voters doesn’t mean that minimax is ignoring the other types of voting change just mentioned; it just means that those other types of change can be transformed into numbers of new voters needed by the additive rule of the previous paragraph.

As mentioned in Section 4.5, Darlington (2016, Section 8) found in a simulation study that minimax picked the same winners as the CMO system on about 99.99% of over 5 million trials. As mentioned there, CMO would presumably be the ideal system if all Condorcet paradoxes were caused by random sampling error. Thus by that measure, there would be little room for Dodgson, Young, or any other system to surpass minimax even if all the other problems with those systems were to disappear.

8.5 Problems with the Hare system (instant runoff voting or single transferable vote)

The Hare system is by far the best-known ranked-choice voting system, and has been used for public elections in at least eight English-speaking countries, though in few others. A serious problem with that system is the frequency with which a Condorcet winner can lose despite winning all their two-way races by majority rule. This occurred one of the first times the Hare system was used for a major public election in the US. That was in 2009, when it was used to elect a mayor for Burlington, VT. The Democratic candidate won all five of his two-way races, but was eliminated by Hare before either the

Republican or the Progressive candidate, who won. In the two-way race between them, the Democrat actually beat the Progressive by a 7.8% margin; a 10% margin is sometimes called a landslide. The voters of Burlington found the Hare choice bizarre, and voted shortly thereafter to discard that system. This anomaly can never occur under minimax.

To see how this can happen under Hare, suppose voters and candidates differ primarily on a liberal-conservative or left-right dimension, A is slightly left of center, B is very centrist, and C is slightly to B's right. Then in a standard spatial model, A would receive top ranks from all voters to the left of the midpoint between A and B, and C would receive top ranks from all those to the right of the midpoint between B and C, while B receives top ranks only from the few remaining voters. Thus, B might well receive the fewest top ranks even though B would beat either A or C in a two-way race.

I ran a computer simulation to estimate how often this might happen. The number of candidates in a trial was 3, 5, 10, or 15, and the number of voters was 15, 35, 75, 155, or 501. A spatial model was assumed to have 1, 2, or 3 dimensions. Each possible combination of these three factors was studied, producing 4·5·3 or 60 different conditions, with 10,000 trials in each condition. Less than 2.5% of these 600,000 trials had Condorcet paradoxes, so nearly all had Condorcet winners. Call it an HC anomaly if the Hare and Condorcet winners are different people. It turned out that the number of voters per trial had little effect on the frequency of these anomalies, so I'll report the results for just the largest number of voters per trial (501). Table 5 shows these frequencies. We see that the frequency of HC anomalies depends enormously on both the number of candidates and the number of opinion dimensions assumed. In real life both these factors will vary greatly across elections.

I'll discuss the results in Table 5 after mentioning another feature of this study. As mentioned above, Table 5 gives results for just 12 of the 60 conditions studied, because each of the 12 conditions in Table 5 was also replicated with 15, 35, 75, or 155 voters per trial. In each of these 60 conditions I also computed the proportion of times that the Hare winner was inferior to (less centrist than) the Condorcet winner when those two were different people. Those 60 proportions ranged from 0.808 to 0.979, with a mean of 0.921 and median of 0.929. That supports a conclusion which I believe most people would accept intuitively – when the Hare winner and Condorcet winner are different people, the Condorcet winner should clearly win.

Table 5. Number of HC anomalies (trials in which the Hare winner and Condorcet winner were different people), in 10,000 artificial-data trials in each of 12 conditions, with 501 voters in each trial. Here c is the number of candidates in each trial, and d is the number of opinion dimensions assumed.

$c \backslash d$	1	2	3
3	1506	409	213
5	4090	1621	691
10	6863	4797	2379
15	7769	6780	4302

I'll use that fact in interpreting the results of Table 5. We see that in one of the conditions there, the HC anomaly appears in over three-fourths of all trials. And even in the condition in which these anomalies are rarest, they would be expected to occur about once in every 47 elections (213 times in

every 10,000 elections). If we assume (as I do) that even a single occurrence of this anomaly would be a substantial embarrassment for proponents of Hare, these results offer a strong reason to avoid Hare.

References

- Balinski M, Laraki R (2010) *Majority judgment: measuring, ranking, and electing*. MIT Press, Cambridge MA
- Black D (1958) *Theory of committees and elections* (Cambridge U Press, Cambridge MA)
- Colomer JM (2013) Ramon Llull: from 'Ars electonis' to social choice theory. *Soc Choice Welf* 40: 317-328
- Darlington RB (2016) Minimax is the best electoral system after all. <https://arxiv.org/ftp/arxiv/papers/1606/1606.04371.pdf>
- Darlington RB (2017). In elections, irrelevant alternatives provide relevant data. <http://arxiv.org/abs/1706.01083>
- Felsenthal DS (2012) Review of paradoxes afflicting procedures for electing a single candidate. In: Felsenthal DS, Machover M (eds) *Electoral systems: paradoxes, assumptions, and procedures* (Springer Verlag, New York) pp. 19-92
- Gehrlein WV (2006) *Condorcet's paradox* (Springer Verlag, New York)
- Green-Armytage J (2014) Strategic voting and nomination. *Soc Choice Welf* 42: 111-138
- Hemaspaandra E, Hemaspaandra LA, Rothe J (1997) Exact analysis of Dodgson elections: Lewis Carroll's 1876 voting system is complete for parallel access to NP. *Journal of the ACM* 44: 806-825
- Kramer GH (1977) A dynamical model of political equilibrium. *J Econ Theory* 16: 310-334
- Naish L (2013) Partial disclosure of votes in STV elections. *Voting Matters* 30:9-13
- Perez J (2001) The strong no show paradoxes are a common flaw in Condorcet voting correspondences. *Soc Choice Welf* 18: 601-616
- Plott CR (1967) A notion of equilibrium and its possibility under majority rule. *American economic review* 57: 787-806
- Rothe J, Spakowski H, Vogel J (2003) Exact complexity of the winner problem for Young elections. *Theory Comput. Systems* (2003) 36: 375-386
- Saari DG (2001) *Chaotic elections!* (American Mathematical Society, Providence RI)
- Simpson PB (1969) On defining areas of voter choice: Professor Tullock on stable voting. *Quarterly journal of economics* 83: 478-490
- Tideman TN (2006) *Collective decisions and voting* (Ashgate, Burlington VT)

Tideman TN, Plassmann F (2012) Modeling the outcomes of vote-casting in actual elections. In: Felsenthal DS, Machover M (eds) *Electoral systems* (Springer Verlag, New York)

Young HP (1977) Extending Condorcet's rule. *Journal of economic theory* 16: 335-353