The Frequency of Condorcet Winners in Real Non-Political Elections

Andrew C. Myers
Department of Computer Science
Cornell University

Abstract

Condorcet-consistent voting is attractive because it follows the principle that the will of a majority of voters should not be denied. However, it is in general possible that there is no Condorcet winner: a cycle of candidates might exist in which each candidate is preferred to the next. The possibility that such a cycle occurs, and the uncertainty about how to handle it, have been an obstacle to the adoption of Condorcet methods. This paper reports on the experience from CIVS, a long-running open-source voting service that supports Condorcet-consistent voting methods. Over a period of about twenty years, users have conducted tens of thousands of polls using CIVS, including many with real-world consequences. During this time, CIVS has thus accumulated probably the largest existing corpus of data about how Condorcet voting functions in practice. CIVS supports multiple completion methods for handling cycles, but the data show that it usually does not matter which completion method is used, because there is rarely a cycle in polls with a large enough number of voters.

1 Introduction

Condorcet-consistent voting has been promoted as a way to more accurately capture the will of voters while avoiding undesirable phenomena such as vote splitting that distort outcomes. Like Instant Runoff Voting (IRV, also known as Ranked Choice Voting), it is a preferential voting system in which voters rank candidates rather than merely selecting their favorite. However, it differs from IRV in that it guarantees to elect a *Condorcet Winner* when there is one—which our results suggest is usually the case.

A Condorcet winner (CW) is a candidate who wins all possible head-to-head contests against other candidates, where candidate A wins a head-to-head contest against candidate B when the number of ballots ranking A over B exceeds the number ranking B over A. The result of these contest is the overall preference for A over B. When there is a Condorcet winner, to elect any other candidate would require rejecting the overall preference of a majority of voters. A Condorcet-consistent voting system satisfies the *Condorcet criterion* that the CW always wins.

A sticking point for adoption of Condorcet-consistent voting is that in general, there may not be a Condorcet winner, for one of two reasons. A less serious reason is there may be a set of (usually two) top candidates who are tied in the head-to-head comparison. If the number of voters is not large, such ties can be expected fairly often. The possibility of ties makes it useful to consider the notion of a *weak* Condorcet winner, who when compared to each other candidate, is preferred on at least as many ballots as the number on which the other candidate is preferred to them. When a candidate is a weak CW but not a CW, it means that they are essentially tied with some other candidates, although these other candidates might not be weak

CWs themselves. However, for elections with many voters, it is reasonable to expect that such ties will be infrequent.

The more troubling possible reason for the absence of a Condorcet winner is that in any voting system there may exist a *preference cycle* of top candidates in which each candidate is preferred over the next one. Condorcet-consistent voting is not the cause of these cycles but it does reveal them as an issue to be addressed. The Smith set [War61] is the smallest set of candidates such that every candidate in the set is preferred over every candidate not in the set. When there is a CW, the Smith set contains just that single candidate, but in general it may include multiple candidates.

Various *Condorcet completion rules* have been defined to resolve cycles (e.g., [Tid87, Sch03, Bla58, Kem59]), with varying properties. Of course, different completion rules resolve the choice of winner in different ways, and arguments can be made in favor of some rules over others. For example, some completion rules satisfy the *Smith criterion* that they always elect a candidate from the Smith set; others do not.

Completion rules are more difficult to justify than the Condorcet criterion, so it would be useful to know how frequently they might be needed in practice. If they were rarely needed, Condorcet voting would be a more attractive option—and more attractive still if they agreed when they were needed. Unfortunately, there has a lack of empirical evidence based on real polls.

The goal of this paper is to add some empirical evidence regarding Condorcet voting and the choice of completion rules, relying on 20 years of data from the Condorcet Internet Voting Service (CIVS) [MC03], a free, open-source, online voting system that supports preferential voting with Condorcet-consistent voting rules. CIVS has hosted more than 25,000 polls in which voters were able to rank their choices.

The data from the CIVS voting system suggests that it is unusual for a poll with many voters not to result in a CW and even more unusual not to have at least a weak CW. Further, in polls where a completion rule was needed to select a winner, existing completion rules often agree on that winner anyway.

None of the polls run on CIVS appear to be political elections as ordinarily understood. However, many of them clearly have real-world consequences. Important open-source organizations, such as the Linux Foundation, OpenStack, and the Ubuntu community, have used CIVS to make leadership decisions. Several academic institutions have as well.

2 The CIVS voting service

The Condorcet Internet Voting Service is currently accessible at `https://civs1.civs.us` [MC03], and its source code is publicly available as open source [MC03]. CIVS only supports Condorcet-consistent voting, although it allows the user to select among various Condorcet-consistent options. To interpret how results from CIVS apply to other voting systems, it is helpful to understand some specifics of how CIVS works.

2.1 History

CIVS was started in 2003 and has been operational as a free, public service for more than 20 years, making it the longest-running open-source preferential voting system. As of February 2024, more than 25,000 polls have been created on CIVS and more than 870,000 votes have been cast. The source code of CIVS is publicly available [MC03], so some organizations have also used it to set up their own internal voting services that they administrate.

CIVS has been used to make a remarkable variety of decisions, including officers of organizations and members of committees, hiring decisions, invited speakers, award recipients, project names and logos, organization bylaws, course topics, book club selections, movies to watch, restaurants to visit, party menus,

	1	2	3	4	5	6	7	8	9	10
1. Jesse Pinkman	–	3313	3831	3750	4287	4567	5507	5357	5604	5631
2. Walter White	2883	–	3461	3404	4165	4254	5147	5079	5415	5376
3. Saul Goodman	2335	2657	–	3085	3860	3960	5258	5286	5351	5478
4. Mike Ehrmantraut	2347	2674	2871	–	3718	3763	4937	5011	5108	5204
5. Gustavo Fring	1810	1854	2090	2150	–	3081	4609	4656	4908	4994
6. Hank Schrader	1565	1835	2031	2158	2807	–	4744	4828	4945	5177
7. Walter White Jr	501	818	630	837	1165	1053	–	2784	3377	3642
8. Steven Gomez	446	680	395	557	886	723	2529	–	3062	3329
9. Skylar White	412	532	538	654	835	806	2103	2200	–	2977
10. Marie Schrader	311	512	338	491	683	498	1730	1830	2336	–

Figure 1: Preference matrix for favorite characters in the show “Breaking Bad”

and much more. Though probably some CIVS elections have involved some form of politicking, CIVS has not been used for explicitly political elections.

CIVS is largely implemented as a set of scripts written in about 25k lines of code using the Perl programming language. Some additional components are implemented in JavaScript and HTML.

2.2 CIVS voting interface

CIVS has an easy-to-use voting interface that allows users to rank the available choices (candidates) by dragging and dropping them into an ordered list. However, the ballots differ from those in traditional preferential voting in two ways. First, it is not necessary to give each choice a distinct rank. By default, all choices are given the same, lowest rank, and a voter may then change the rank. Voters may also explicitly give a set of choices the same rank, declaring a lack of preference between any two of them.

A second distinctive feature of the CIVS ballot is the ability to entirely avoid expressing any opinion about certain choices. If the poll supervisor has enabled this feature, voters may select the rank “No opinion”, which means their ballot will express no preference regarding the choice. Selecting “No opinion” is very different from selecting the lowest rank. It is particularly helpful in situations where a voter *should* not express an opinion about a given choice, such as when there is a conflict of interest.

CIVS poll results include not just a determination of the winning choice but also a ranking of all the choices. For all voting rules, this ranking is computed by first computing the winner (or winners) of the poll. These choices are then removed from the ballots and the same rule is used to find the next highest ranked choice, and so on.

CIVS can also display in matrix form the results of all of the head-to-head contests, as illustrated in Figure 1. For example, voters had an overall preference for Jesse Pinkman over Walter White of 3313–2883. To make this preference matrix easier to interpret visually, choices are ordered according to the computed ranking. Cells are colored to indicate the winner of each comparison. When all cells above the diagonal are green, it means that each choice is preferred above all lower-ranked choices. This display makes it easy to see when a poll has preference cycles, because pink (loss) or yellow (tie) cells appear above the diagonal, and green or yellow ones appear below.

2.3 Completion rules

CIVS implements five different completion rules that come into play when there are preference cycles. The software has standardized interface to modules that implement voting rules, so it is not difficult to extend it with additional rules.

Minimax Minimax [Bla58] (also known as Simpson–Kramer [Sim69]), is the default completion rule used by CIVS. It orders candidates based on their weakest defeats: that is, the strongest preferences against the candidate. Appealingly, it finds the candidate who could become the Condorcet winner with the fewest number of additional ballots. Unlike some other rules that also try to find candidates "close to" being Condorcet winners, Minimax is also inexpensive to compute.

There are several variants of Minimax, with subtly different properties. Two versions of Minimax were implemented for comparison purposes. The first, called Minimax_M here, follows a proposal by Darlington [Dar18]. It orders preferences first by the *margin* of the defeat. Let W stand for the number of ballots ranking candidate 1 over candidate 2, and L the number of ballots with the opposite ranking. Larger margins $W - L$ are considered to be stronger preferences. When margins are equal, defeats are ordered by L, with smaller values of L considered to be stronger preferences. The second version, Minimax_{WV} (for "winning votes") uses the same ordering as the Ranked Pairs and Schulze completion rules: preferences are ordered first by W and secondarily by L as above. Note that the two versions differ only when voters abstain from comparing some candidates.

In both versions, two candidates are compared first by their weakest defeats; then, if tied on those, by their 2nd weakest defeats, and so on, possibly even comparing them on "defeats" that are actually wins.

Ranked Pairs The Ranked Pairs completion rule was introduced by Tideman [Tid87] with the goal of making a voting system unaffected by indistinguishable "clone" candidates. The intuition is to avoid considering preferences that would introduce a preference cycle. The set of preferences is first sorted in order of decreasing strength as defined in Minimax_{WV}. Then, starting from an empty set, preferences are added ("affirmed") in that order until a preference cycle would result. At that point, any undefeated candidates are considered to be winners. No random choices are made in the CIVS implementation of this rule; a set of equal-strength preferences are added all at once or not at all.

Schulze (Beatpath) The Schulze completion rule [Sch03] (also known as Beatpath) orders candidates according to the strongest *beatpath* connecting them. A beatpath is simply a sequence of distinct candidates. The strength of a beatpath is the minimum-strength preference along the beatpath. The effective preference of one candidate over another is the strength of the strongest beatpath connecting them. CIVS uses the same "winning votes" ordering of preferences as in Minimax_{WV}. The Schulze method can be implemented fairly efficiently by using the well-known Floyd–Warshall algorithm over the commutative (min, max) semiring of preferences.

Condorcet–IRV Various completion rules have been proposed that aim to combine the strong strategy resistance of IRV [Tid18] with Condorcet consistency, collectively called Condorcet–Hare rules [GA11]. The approach CIVS implements is a variation on one proposed by Chris Benham. It first constructs a plurality-based seeding of the candidates. If there is no CW, candidates are removed progressively, starting from the lowest-seeded candidate, until a CW exists. When seeding candidates, ties in the number of ballots on which the candidate receives the top rank are broken by using the number of second-rank ballots (and so on with lower ranks used as necessary). When a ballot has multiple candidates at a given rank, they each receive proportional credit for occupying that rank.

Bottom-Two Runoff Bottom-Two Runoff (BTR) [BTR19], proposed by Rob LeGrand, is another hybrid completion rule that incorporates a runoff process. Candidates are first seeded in the same way as for Condorcet–IRV. However, rather than simply eliminating the lowest-seeded candidate, the overall preference

Minimum votes	Count	CW	Weak CW	CW %	Weak CW %	95% CI
10	10354	8600	9857	83.1%	95.2%	± 0.4%
20	5487	4873	5277	88.8%	96.2%	± 0.5%
50	2016	1937	1986	96.1%	98.5%	± 0.5%
100	1117	1093	1108	97.9%	99.2%	± 0.6%
150	653	642	649	98.3%	99.4%	± 0.7%
200	479	472	491	98.5%	99.2%	± 1.0%
300	320	316	317	98.8%	99.1%	± 1.3%

Figure 2: Frequency of CWs and weak CWs with an increasing number of voters

between the two lowest-seeded candidates is used to determine whom to eliminate, with the lower-seeded candidate eliminated in the case of a tie. This process is repeated until the winner is found. Since a CW never loses a head-to-head contest, Bottom-Two Runoff is Condorcet-consistent.

2.4 Proportional Representation

CIVS also offers a proportional representation mode in which multiple winners are chosen while maintaining approximately proportional representation of the voters. This mode is much more expensive, and infrequently used, so it is not evaluated here.

3 Results

As of September 2023, a total of 26,259 polls in which at least one vote had been cast had been run on the CIVS system. Because it is difficult to draw useful conclusions from very small polls, the analysis here focuses on the 10,354 polls in which at least 10 votes were cast, there were no more than 100 candidates to choose from, and CIVS's proportional mode was not used.

The number of polls with any given minimum number of votes cast decreases quickly. Only 1,117 polls had at least 100 votes and only 120 had at least 500 votes. However, even the data from relatively small polls generate useful results that, with some care, appear to extrapolate to larger polls.

3.1 Frequency of Condorcet winners

As the number of voters increases, we would expect the frequency of Condorcet winners and of weak CWs to converge, because exact ties in head-to-head contests become increasingly unlikely.

Figure 2 shows how the frequencies of CWs and weak CWs vary with the size of polls. Each row corresponds to a minimum number of votes cast, and gives statistics regarding all polls with at least that minimum number. The "Count" column gives the number of polls reported on in each row. As expected, the fractions of CWs and of weak CWs become close for larger polls. Notably, the percentage of weak CWs increases to over 99% for the larger poll sizes. The estimate of this percentage naturally becomes less statistically reliable for smaller numbers of polls; the last column reports the 95% confidence interval on the percentage of weak CWs, using the adjusted Wald method [AC98].

Another question of interest is how the frequency of Condorcet winners depends on the number of choices available to voters. Naturally, we would expect Condorcet winners to become less frequent as the

Minimum choices	Count	CW	Weak CW	CW %	Weak CW %	95% CI
3	4398	3797	4188	86.3%	95.2%	± 0.6%
4	3749	3178	3544	84.8%	94.5%	± 0.7%
5	3317	2781	3121	83.8%	94.1%	± 0.8%
7	2270	1919	2137	84.5%	94.1%	± 1.0%
10	1593	1322	1484	83.0%	93.2%	± 1.2%
20	628	508	573	80.9%	91.2%	± 2.2%
40	224	188	210	83.9%	93.8%	± 3.3%

Figure 3: Frequency of CWs and weak CWs with an increasing number of choices

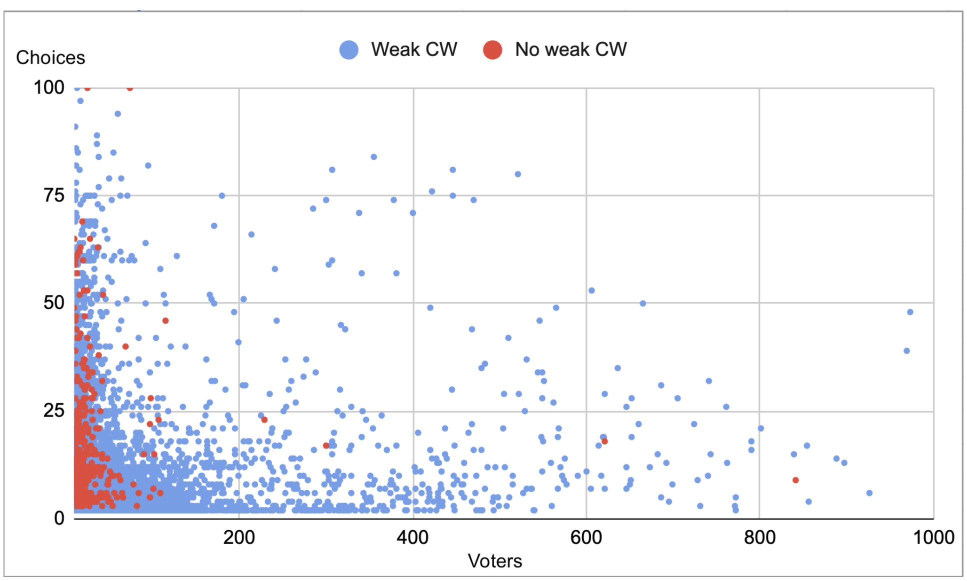

Figure 4: Scatterplot of CIVS polls. Light blue dots are polls with a weak CW; red dots are polls with no weak CW.

number of choices increases. Figure 3 shows that this is true, but that existence of Condorcet winners remains likely even for fairly large sets of choices. Note that only polls with at least 20 voters were considered.

Figure 4 shows a scatterplot of all polls with at least 10 voters. It is visually evident that polls in which there was no weak CW tend to have a relatively large number of choices compared to the number of voters.

3.2 Condorcet vs. plurality

One basic question we might ask is how often Condorcet voting produces a Condorcet winner who would not be the winner if simple plurality voting were used. Recall that in plurality voting, voters cast a vote for a single candidate and the candidate with the most votes wins.

CIVS does not directly support plurality voting. To evaluate the agreement between plurality voting and Condorcet-consistent voting experimentally, voter rankings from CIVS were instead used to simulate plurality voting. For each candidate, the number of ballots in which that candidate received the highest rank was counted. For the small minority of ballots that tied multiple candidates at the highest rank, the count was divided among those candidates. The candidates with the highest top-rank count were then deemed to be plurality winners for that poll.

Plurality voting was deemed to have demonstrably failed the Condorcet criterion in any poll where there was a weak Condorcet winner but none of the simulated plurality winners was a weak CW. Among 9,857 polls with a weak CW and at least 10 votes cast, there were 1,485 in which plurality voting failed. Thus, plurality failed on 14% of these polls—a substantial percentage.

3.3 Condorcet vs. IRV

The Instant Runoff Voting (IRV) system (the single-winner version of Hare or STV) is a popular preferential voting system used around the world. It is not, however, Condorcet-consistent. Candidates are seeded according to the number of ballots on which they were top-ranked. The bottom-seeded candidate is eliminated and the ordering is recomputed as though that candidate did not exist, with the process repeating until only the winning candidate remains.

IRV sometimes fails to elect a Condorcet winner; the 2009 mayoral election in Burlington, Vermont is one prominent example where such a failure occurred [Ols09]. Using the CIVS data set, it is possible to estimate how frequently such a failure occurred by using the CIVS rankings to simulate an IRV election. As with the plurality simulation, lower ranks were used to break top-rank ties, and ballots that placed multiple candidates at a given rank were effectively divided among the candidates at that rank.

As with plurality, IRV was deemed to have failed when there was a weak Condorcet winner but the simulated IRV winner was not a weak CW. Among 9,857 polls with a weak CW and at least 10 votes cast, there were 584 in which IRV failed, or about 6% of the polls.

3.4 Agreement between different completion rules

In the apparently small fraction of elections where there is no CW, a completion rule may be used to select a winner. However, there is little existing data on how often these methods agree on the winner in practice. (Of course, they always agree in the polls for which a CW exists.)

Only polls with at least 20 votes cast were considered, but the results are qualitatively similar for different size thresholds, in the sense that the rates of disagreements were similar. There were 210 such CIVS polls in which there was no weak CW. Two completion rules were considered to disagree if the (usually singleton)

Schulze	14				
Minimax$_{WV}$	26	13			
Minimax$_M$	54	64	64		
Condorcet–IRV	98	101	121	112	
Bottom-2 Runoff	104	107	120	102	65
	Ranked Pairs	Schulze	Minimax$_{WV}$	Minimax$_M$	Condorcet–IRV

Figure 5: Frequency of disagreement among 6 completion rules in 210 polls

sets of winners determined by the two methods were disjoint. In 59 (28%) of these 210 polls, all six completion rules agreed. In the remaining 151 (72%), there was disagreement between at least two methods.

Figure 5 shows how often the completion rules implemented by CIVS disagreed. Some interesting patterns may be noted in the table. Consistent with claims by Tideman [Tid18], the data show that Schulze and Ranked Pairs were in strong agreement, producing different answers in only 14 polls. Further, Minimax$_{WV}$ disagrees with Schulze in only 13 polls, and also usually agrees with Ranked Pairs. Minimax$_M$ still agrees with the above three rules in most cases, but it seems that its use of margins for preference ordering leads more disagreement, even between Minimax$_M$ and Minimax$_{WV}$.

The two rules that incorporate some form of runoff based on rank counts, Condorcet–IRV and Bottom-Two Runoff, agree with each other more often than with the other three rules. These rules aim to reduce the impact of strategic voting techniques such as ballot truncation or burying, but it is difficult to infer from these data whether voters were deliberately voting strategically. The strongest disagreement in the table is between Minimax$_{WV}$ and Condorcet–IRV, which agreed less than half the time.

4 Conclusions

The author is not aware of any similarly large-scale analysis of Condorcet-consistent, real-world voting, so this study helps answer some of the questions around the use of Condorcet-consistent voting. As some have predicted, Condorcet winners seem to be very likely for elections with many voters. Perhaps surprisingly, they seem to remain likely even when there are many candidates. The data also sheds light on how frequently various completion rules agree on the winner in the minority of elections where no Condorcet winner exists.

It must be acknowledged that there are reasons why the results reported here, collected from non-political polls, might not apply to real political elections. For example, CIVS has some features that political elections might avoid, such as allowing a large number of candidates on the ballot, and giving the ability to tie candidates in the ranking or to avoid ranking them entirely. However, these features would seem only to decrease the likelihood of Condorcet winners. Perhaps a greater threat to validity is that candidates in political elections spend far more effort optimizing their strategy and their positions to the voting system in use. With Condorcet-consistent voting, we may expect coordinated strategic voting aimed at creating apparent preference cycles. In that case, Condorcet completion rules with resistance to strategic voting seem likely to be helpful; the present study offers some insight about how often the choice of completion rule affects election results. No doubt more can be learned from further analysis of the CIVS data; an anonymized version of the CIVS dataset is available at the CIVS source repository [MC03].

5 Acknowledgments

The author thanks Michael Clarkson for his help implementing CIVS, Chris Benham and Robert Bristow-Johnson for useful background, Drew Zagieboylo, Yulun Yao, and Nicolaus Tideman for suggestions on experiments, and Robbie Robinette and Pamela MacDougal for suggestions on presentation.

References

[AC98] Alan Agresti and Brent A. Coull. Approximate is better than ‘exact’ for interval estimation of binomial proportions. *The American Statistician*, 52:119–126, 1998.

[Bla58] Duncan Black. *Theory of committees and elections*. Cambridge Univ. Press, 1958.

[BTR19] Bottom-two-runoff IRV. https://electowiki.org/wiki/Bottom-Two-Runoff_IRV, 2019. Version of February, 2024.

[Dar18] Richard B. Darlington. Are Condorcet and minimax voting systems the best? Arxiv preprint 1807.01366v1, 2018.

[GA11] James Green-Armytage. Four Condorcet-Hare hybrid methods for single-winner elections. *Voting Matters*, 29, 2011.

[Kem59] John G. Kemeny. Mathematics without numbers. *Daedalus*, 88(4):577–591, 1959.

[MC03] Andrew C. Myers and Michael Clarkson. CIVS: Condorcet Internet Voting System. `https://github.com/andrewcmyers/civs`, 2003. Service at `https://civs1.civs.us`.

[Ols09] Brian Olson. IRV failure in the real world. `https://bolson.org/~bolson/2009/20090303_burlington_vt_mayor.html`, 2009.

[Sch03] Markus Schulze. A new monotonic and clone-independent single-winner election method. *Voting Matters*, 17(1):9–19, 2003.

[Sim69] Paul B. Simpson. On defining areas of voter choice: Professor Tullock on stable voting. *Quarterly Journal of Economics*, 83(3):478–490, August 1969.

[Tid87] Nicolaus Tideman. Independence of clones as a criterion for voting rules. *Social Choice and Welfare*, 4:185–206, 1987.

[Tid18] Nicolaus Tideman. *The Oxford Handbook of Public Choice (Vol. 2)*, chapter 1, pages 5–23. December 2018.

[War61] Benjamin Ward. Majority rule and allocation. *J. of Conflict Resolution*, 5(4):379–389, 1961.