Split Cycle: A New Condorcet Consistent Voting Method Independent of Clones and Immune to Spoilers

Wesley H. Holliday† and Eric Pacuit‡

† University of California, Berkeley (wesholliday@berkeley.edu)
‡ University of Maryland (epacuit@umd.edu)

March 23, 2021

Abstract

We propose a Condorcet consistent voting method that we call Split Cycle. Split Cycle belongs to the small family of known voting methods that significantly narrow the choice of winners in the presence of majority cycles while also satisfying *independence of clones*. In this family, only Split Cycle satisfies a new criterion we call *immunity to spoilers*, which concerns adding candidates to elections, as well as the known criteria of *positive involvement* and *negative involvement*, which concern adding voters to elections. Thus, in contrast to other clone-independent methods, Split Cycle mitigates both "spoiler effects" and "strong no show paradoxes."

Contents

1 Introduction

A voting method is Condorcet consistent if in any election in which one candidate is preferred by majorities to each of the other candidates, this candidate—the Condorcet winner—is the unique winner of the election. Condorcet consistent voting methods form an important class of methods in the theory of voting (see, e.g., Fishburn 1977; Brams and Fishburn 2002, § 8; Zwicker 2016, § 2.4; Pacuit 2019, § 3.1.1). Although Condorcet methods are not currently used in government elections, they have been used by several private organizations (see Wikipedia contributors 2020b) and in over 15,000 polls through the Condorcet Internet Voting Service (https://civs.cs.cornell.edu). Recent initiatives in the U.S. to make available Ranked Choice Voting (Kambhampaty 2019), which uses the same ranked ballots needed for Condorcet methods, bring Condorcet methods closer to practical application. Indeed, Eric Maskin and Amartya Sen have recently proposed the use of Condorcet methods in U.S. presidential primaries (Maskin and Sen 2017a,b). In the meantime, Condorcet methods continue to be used by small committees and clubs.

In this paper, we propose a Condorcet consistent voting method that we call Split Cycle, which has a number of attractive axiomatic properties.[1] Split Cycle responds to a concern well expressed by a 2004 letter to the Washington Post sent by a local organizer of the Green Party, as quoted by Miller (2019, p. 119):

> [Electoral engineering] isn't rocket science. Why is it that we can put a man on the moon but can't come up with a way to elect our president that allows voters to vote for their favorite candidate, allows multiple candidates to run and present their issues and... [makes] the 'spoiler' problem... go away?

Starting with the problem of spoilers, Split Cycle satisfies not only the *independence of clones* criterion proposed by Tideman (1987) as an anti-spoiler criterion but also a new criterion we call *immunity to spoilers* that rules out spoiler effects not ruled out by independence of clones. What the Green Party organizer meant by a voting method that "allows voters to vote for their favorite candidate" is open to multiple interpretations; if it means a reasonable voting method that never provides an incentive for strategic voting, as Miller takes it to mean, then such a method is unavailable by well-known theorems on strategic voting (see Taylor 2005). More modestly, one may ask for a voting method such that at the very least, voters will never cause their favorite candidate to lose by going to the polls and expressing that their favorite candidate is their favorite. Understood this way, one is asking for a voting method that satisfies the criterion of *positive involvement* (Saari 1995). Split Cycle satisfies this criterion, as well as a number of other desirable criteria: the Condorcet loser criterion, monotonicity, independence of Smith-dominated alternatives, reversal symmetry, negative involvement, and two criteria concerning the possibility of ties among winners, which we call the narrowing criterion and rejectability criterion. In fact, Split Cycle can be distinguished from all voting methods we know of in any of the following three ways:

- Only Split Cycle satisfies independence of clones, positive involvement, narrowing, and at least one of Condorcet consistency, monotonicity, and immunity to spoilers.[2]
- Only Split Cycle satisfies independence of clones, negative involvement, and narrowing.
- Only Split Cycle satisfies independence of clones, immunity to spoilers, and rejectability.

Split Cycle is an example of a head-to-head (or pairwise) voting method. We compare each pair of candidates a and b in a *head-to-head match*. If more voters rank a above b than rank b above a, then a

[1] After posting a draft of this paper, we learned from Jobst Heitzig of his notion of the "immune set" discussed in a 2004 post on the Election-Methods mailing list (Heitzig 2004), which is equivalent to the set of winners for Split Cycle after replacing 'stronger' with 'at least as strong' in Heitzig's definition in the post. See Remark 3.13 for further connections with Heitzig 2002. As far as we know, the Split Cycle voting method has not been studied in the research literature. In a companion paper, Holliday and Pacuit 2020, we study Split Cycle as what is known as a *collective choice rule* in the social choice theory literature.

[2] Among proposed non-Condorcet methods, we believe only Ranked Choice satisfies both independence of clones and positive involvement, but it fails the monotonicity criterion, which Split Cycle satisfies, as well as immunity to spoilers and negative involvement (see Appendix C.8).

wins the head-to-head match and b *loses* the head-to-head match. If a wins against b, then the number of voters who rank a above b minus the number who rank b above a is a's *margin of victory over* b. If one candidate wins its matches against all other candidates, that candidate is the winner of the election. But there is a chance that every candidate will lose a match to some other candidate. When this happens, there is a *majority cycle*: a list of candidates where each candidate wins against the next in the list, and the last candidate wins against the first. For example, candidates a, b, c form a majority cycle if a wins against b, b wins against c, and c wins against a. There can also be cycles involving more than three candidates.

Split Cycle deals with the problem of majority cycles as follows:[3]

1. In each cycle, identify the head-to-head win(s) with the smallest margin of victory in that cycle.

2. After completing step 1 for all cycles, discard the identified wins. All remaining wins count as *defeats* of the losing candidates.

For example, if a wins against b by 1,000 votes, b wins against c by 2,000 votes, and c wins against a by $3,000$ votes, then a's win against b is discarded. Candidate b's win against c counts as a defeat of c unless it appears in another cycle (involving some other candidates) with the smallest margin of victory in that cycle. The same applies to c's win against a. Crucially, after step 2, there is always an *undefeated* candidate (as we will show). If there is only one, that candidate wins the election. If there is more than one, then a tiebreaker must be used. When an immediate decision is required, the tiebreaker could be, e.g., to randomly choose one of the undefeated candidates. Alternatively, a runoff election could be held between the undefeated candidates, according to the same rules described above, allowing additional voters to participate.

In the rest of this introduction, we provide additional background to the benefits of Split Cycle: we spell out the problem of "spoiler effects" that the independence of clones and immunity to spoiler criteria mitigate (Section 1.1), followed by the "strong no show paradox" that the positive involvement and negative involvement criteria rule out (Section 1.2). We then provide a roadmap of the rest of the paper in Section 1.3.

1.1 The Problem of Spoilers

Let us begin with one of the most famous recent examples of a spoiler effect in a U.S. election.

Example 1.1. In the 2000 U.S. Presidential election in Florida, run using the Plurality voting method, George W. Bush, Al Gore, and Ralph Nader received the following votes:

2,912,790	2,912,253	97,488
Bush	Gore	Nader

It is reasonable to assume that if Nader had dropped out before Election Day, then a sufficiently large number of his 97,488 voters would have voted for Gore so that Gore would have won the election (Magee 2003). It is also reasonable to assume that while for some Gore voters, Nader may have been their favorite but they strategically voted for Gore, still many more voters preferred Gore to Nader than vice versa. So Nader would pose no direct threat to Gore in a two-person election, but by drawing enough votes away from Gore in the three-person election, he handed the election to Bush. Thus, Nader "spoiled" the election for Gore.

[3] There is a more computationally efficient way to calculate the Split Cycle winners (see Footnote 18), but this simple two-step procedure is appropriate for explaining the method to voters.

In elections where voters submit rankings of the candidates, rather than only indicating their favorite, we can give precise content to the claim that one candidate spoiled the election for another. One attempt to do so is given by Tideman's (1987) criterion of *independence of clones*. A set C of two or more candidates is a set of clones in an election with ranked ballots if no candidate outside of C appears in between two candidates from C on any voter's ballot. Suppose, for example, that if we had collected ranked ballots in the 2000 Florida election, the results would have been as follows:

2,912,790	2,912,253	97,488
Bush	Gore	Nader
Gore	Nader	Gore
Nader	Bush	Bush

In this imaginary election with ranked ballots, {Gore, Nader} is a set of clones, because Bush never appears between Gore and Nader on any ballot. The independence of clones criterion says (in part) that a non-clone candidate—in this case, Bush—should win in an election if and only if they would win after the removal of a clone from the election. But if we remove Nader, who is a clone of Gore, we obtain the following election:

2,912,790	2,912,253	97,488
Bush	Gore	Gore
Gore	Bush	Bush

In this election, Gore wins. Thus, this imaginary example shows that the Plurality voting method does not satisfy independence of clones. Of course, independence of clones would not literally account for the sense in which Nader spoiled the 2000 Florida election for Gore, even if we had in fact collected ranked ballots. For surely some ballots would have had Bush in between Gore and Nader.

Next we give an example of a spoiler effect that cannot be captured by independence of clones.

Example 1.2. This example involving Ranked Choice Voting, also known as Instant Runoff Voting (IRV), the Alternative Vote, or the Hare Method, comes from ElectionScience.org (https://www.electionscience.org/library/the-spoiler-effect/) as "a simplified approximation of what happened in the 2009 IRV mayoral election in Burlington, Vermont." Consider the following election for two candidates, a Democrat d and a Progressive p (where we split the voters who prefer d to p into two columns for comparison with the table to follow):

37	29	34
d	d	p
p	p	d

For two candidates, Ranked Choice is simply Majority Voting, so the Ranked Choice winner is d. But now suppose an additional Republican candidate r joins the race:

37	29	34
r	d	p
d	p	d
p	r	r

Ranked Choice works by first removing the candidate who received the fewest first place votes—in this case, candidate d—from all ballots, resulting in the following:

37	29	34
r	p	p
p	r	r

Now p has a majority of first place votes, so p is declared the Ranked Choice winner. Note, however, that in the three-person election, d was the Condorcet winner: a majority of voters (66) prefer d to p, and a majority of voters (63) prefer d to r. Yet the addition of r kicks d out of the winning spot and results in p being the Ranked Choice winner. Thus, r spoiled the election for d.

The independence of clones criterion cannot account for the sense in which r spoiled the election for d, because r is not a clone of any candidate. Moreover, Ranked Choice satisfies independence of clones. Thus, independence of clones does not address all spoiler effects. Below we will propose a criterion of *immunity to spoilers* that does account for cases like that of the Burlington mayoral election.

What kind of "spoiler effects" should we try to prevent? This question mixes the conceptual question of what a "spoiler" is and the normative question of what effects we should prevent.

First consider an obviously flawed definition of a spoiler: b is a "spoiler" for a just in case a would win without b in the election, but when b joins, then b but not a wins. This is of course not the relevant notion, since spoilers are not winners.

Thus, consider a second definition: b is a "spoiler" for a just in case a would win without b in the election, but when b joins, neither a nor b wins. It is clearly necessary, in order for b to be a spoiler for a, that neither a nor b wins after b joins, but is it sufficient? Whether or not it is sufficient according to the ordinary concept of a spoiler, we do not think that we should prevent *all* such effects.[4] Consider the following example, where the diagram on the right indicates that, e.g., the number of voters who prefer a to c is one greater than the number who prefer c to a:

2	3	4
b	a	c
a	c	b
c	b	a

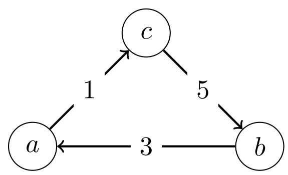

For this election, we agree with proponents of voting methods such as Minimax, Ranked Pairs, and Beat Path (all defined in Appendix C) that c should be the winner. Everyone suffers a majority loss to someone, but while c suffers a slight majority loss to a, a suffers a larger majority loss to b, who suffers an even larger majority loss to c. The electorate is in a sense incoherent, and the fairest way to respond in this case is to elect c.[5] But if b had not been in the election, so we would not have had to account for the majority preferences for b over a and for c over b, then a would have been the appropriate winner in the two-person

[4]Thus, we think it is too strong to require that a voting method satisfy the condition known as the *Aïzerman property* (Laslier 1997, p. 41) or *weak superset property* or $\widehat{\alpha}_{\subseteq}$ (Brandt et al. 2018), which is equivalent to the condition that if a would win were no candidate from a set N in the election, then after the candidates in N (the "newcomers") join the election, if none of the candidates in N wins, then a still wins. For the same reason, we think it is too strong to require that a voting method satisfy the *strong candidate stability* property (studied for resolute voting methods in Dutta et al. 2001 and Ehlers and Weymark 2003 and generalized to irresolute methods in Eraslan and McLellan 2004 and Rodríguez-Álvarez 2006), which implies that if b would not win were b to join the election, then a would win with b in the election if and only if a would win without b in the election (cf. $\widehat{\alpha}$ in Brandt et al. 2018). The problem with these conditions is that they ignore the majority preference relations between a and the new candidates, which our condition of immunity to spoilers takes into account.

[5]At least for a deterministic voting method. A voting method that outputs a probability distribution on the set of candidates (see Brandt 2017) could assign nonzero probabilities to each candidate in this example. But in this paper we do not consider probabilistic voting methods.

election. Since we agree with all of these verdicts, we do not think a voting method should prevent all effects of the kind described in the second definition.

Similar remarks apply to a third definition (from Wikipedia contributors 2020a): b is a "spoiler" for a just in case a would win without b in the election, *and (most of) the voters who prefer b over c also prefer a over c*, but when b joins, neither a nor b wins but rather c wins. Based on the example above, in which *all* voters who prefer b over c also prefer a over c, we do not think a voting method should prevent all such effects.

This bring us to our proposal for the kind of spoiler effects we wish to prevent:

- if a would win without b in the election, *and more voters prefer a to b than prefer b to a*, but when b joins, neither a nor b wins the election, then b is a "spoiler" for a in a way that should be prevented.

This captures Example 1.1 (in the imaginary version with ranked ballots), as a majority of voters prefer Gore to Nader. Unlike independence of clones, it also captures Example 1.2, as a majority of voters prefer the Democrat to the Republican. We will say that a voting method satisfies *immunity to spoilers* if it always prevents the kind of effect just described (for a formal definition, see Section 5.1).

The criterion of immunity to spoilers follows from what we consider a more fundamental criterion, which we call *stability for winners*:

- if a would win without b in the election, *and more voters prefer a to b than prefer b to a*, then when b joins, a should still win.

This criterion can be seen as extending the idea of Condorcet consistency[6] to the variable-candidate setting: a candidate who would be a winner without the newcomers and is majority preferred to all the newcomers remains a winner after the addition of the newcomers. We will show that Split Cycle satisfies immunity to spoilers in virtue of satisfying this more basic criterion of stability for winners.

1.2 The Strong No Show Paradox

The term "no show paradox" was coined by Fishburn and Brams (1983) for violations of what is now called the *negative involvement* criterion (see Pérez 2001). This criterion states that if a candidate x is not among the winners in an initial election scenario, then if we add to that scenario some new voters who rank x as the (unique) last place candidate on their ballots, then the addition of those voters should not make x a winner. Perhaps surprisingly, well-known voting methods such as Ranked Choice, Ranked Pairs, and Beat Path fail to satisfy the negative involvement criterion. In an example of Fishburn and Brams, two voters are unable to make it to a Ranked Choice election, due to their car breaking down. They later realize that had they voted in the election, their *least favorite* candidate would have won. For a simplified version of the Fishburn and Brams example, consider the following example for Ranked Choice (Pacuit 2019, § 3.3):

2	3	1	3
a	b	c	c
b	c	a	b
c	a	b	a

[6] More accurately, the idea that any Condorcet winner ought to be at least tied for winning the election.

Candidate a receives the fewest first place votes, so a is eliminated in the first round. With a eliminated from the ballots, b receives a majority of first place votes and hence wins according to Ranked Choice. But now suppose that two additional voters with the ranking abc (so a is preferred to b and c, and b is preferred to c) make it to the election—their car does not break down—resulting in the following:

4	3	1	3
a	b	c	c
b	c	a	b
c	a	b	a

Now candidate b receives the fewest first place votes, so b is eliminate in the first round. With b eliminated from the ballots, c receives a majority of first place votes and hence wins according to Ranked Choice. Thus, the addition of two voters who rank c last makes c the winner. This is a failure of negative involvement.

The dual of the negative involvement criterion is the *positive involvement* criterion (again see Saari 1995, Pérez 2001).[7] This criterion states that if a candidate x is among the winners in an initial election scenario, then if we add to that scenario some new voters who rank x as the (unique) first place candidate on their ballots, then the addition of these new voters should not make x a loser. Moulin (1988) gives the following example of a failure of positive involvement for the Sequential Elimination voting method in which a faces b in the first round, and then the winner of the first round faces c.[8] In the initial election scenario, we have the following ballots:

2	2	1
a	b	c
b	c	a
c	a	b

In the first round, a beats b (3 voters prefer a to b, and only 2 prefer b to a), and then in the second, c beats a (3 voters prefer c to a, and only 2 prefer a to c). But now suppose two additional voters make it to the election with the ballot cba, so we have:

2	2	1	2
a	b	c	c
b	c	a	b
c	a	b	a

Now in the first round, b beats a (4 voters prefer b to a and only 3 prefer a to b), and then in the second, b beats c (4 voters prefer b to c and only 3 prefer c to b). Thus, adding voters whose favorite candidate is c turns c from being a winner to a loser. This is a failure of positive involvement.

What is wrong with voting methods that fail negative or positive involvement? Our objection to them is *not* that they incentivize a certain kind of strategic (non-)voting. All reasonable voting methods incentivize some kind or other of strategic voting (see Taylor 2005). Suppose we have a group of voters who will definitely cast their ballots and vote sincerely, regardless of the electoral consequences. Thus, the voters in

[7] Also see Kasper et al. 2019, where positive and negative involvement are called the "Top Property" and "Bottom Property", respectively. A closely related criterion for unique winners is given by Richelson (1978) under the name 'voter adaptability'.

[8] Unlike Ranked Choice, this is a Condorcet consistent voting method. We have slightly modified Moulin's example to avoid the use of a tie-breaking rule, at the expense of adding two new voters rather than one in the second election scenario.

the previous example who rank c first will come to the polls and cast their ballots, resulting in c losing the election that otherwise c would have won. Since the voters do not stay home strategically, is the fact that the voting method fails positive involvement unproblematic? Not at all. The problem is that the voting method is responding in the wrong way to additional unequivocal support from a voter for a candidate (c is the voter's unique favorite). As an analogy, a voting method failing the monotonicity criterion (see Section 4.2) also means that it can incentivize a certain kind of strategic voting; but even for a group of always sincere voters, failing monotonicity is a flaw of a voting method because it means that the voting method is responding in the wrong way to voters purely improving a candidate's position relative to other candidates.

The failure of positive or negative involvement is now sometimes called the "strong no show paradox."[9] The reason seems to be that Moulin (1988) changed the meaning of "no show paradox" to stand not for a violation of negative (or positive) involvement but rather for a violation of the *participation* criterion: if a candidate x is the winner in an initial election, then if we add to that scenario some new voters who rank x above y, then the addition of these new voters should not make y the winner. Crucially, it is not required here that x is at the top of the voters' ballots or that y is at the bottom. In our view, this participation criterion is problematic. To see why (as made precise in Appendix B.2), note that if the new voters do not rank x at the top of their ballots and do not rank y at the bottom, then in the presence of majority cycles, new voters having ballots with the ranking $x' x\, y\, y'$ increase the number of people who prefer x' to x, which may result in x' knocking x out of contention, and increase the number of people who prefer y to y', which may result in y' no longer knocking y out of contention. No wonder, then, that the winner may change from x to y. In fact, remarkably, adding new voters who rank x above y may make y's new position vis-à-vis other candidates perfectly symmetrical to x's old position vis-à-vis other candidates, up to a renaming of the candidates (again see Appendix B.2). A certain kind of neutrality then requires that the winner changes from x to y. Of course, a method not satisfying participation will incentivize some strategic non-voting, as the voters in question will have an incentive not to vote (sincerely). But again, all voting methods incentivize strategic behavior. Thus, we are not so troubled by results showing that all Condorcet consistent voting methods fail versions of participation[10] and therefore incentivize some strategic behavior. By contrast, we are troubled by failures of positive or negative involvement, as this shows that the method responds in the wrong way to unequivocal support for (resp. rejection of) a candidate.

Unlike well-known voting methods such Ranked Choice, Ranked Pairs, and Beat Path, the method we propose in this paper, Split Cycle, satisfies positive and negative involvement. Hence it is not only immune to spoilers but also immune to the strong no show paradox. Figure 1 illustrates how solving the spoiler problem and the strong no show paradox leads uniquely to Split Cycle as opposed to standard voting methods.

[9] Perez (2001) calls violations of positive involvement the "positive strong no show paradox" and violations of negative involvement the "negative strong no show paradox." Felsenthal and Tideman (2013) and Felsenthal and Nurmi (2016) call them the "P-TOP" and "P-BOT" paradoxes, respectively.

[10] Note that Moulin (1988) only proves participation failure for Condorcet consistent voting methods that are *resolute*, i.e., always pick a unique winner, which requires imposing an arbitrary tie-breaking rule that violates anonymity or neutrality (see Section 4). Since none of the standard Condorcet consistent voting methods are resolute, one may wonder about the significance of the fact that resolute Condorcet methods all fail participation. For discussion of the irresolute case, see Pérez 2001, Jimeno et al. 2009, and Sanver and Zwicker 2012.

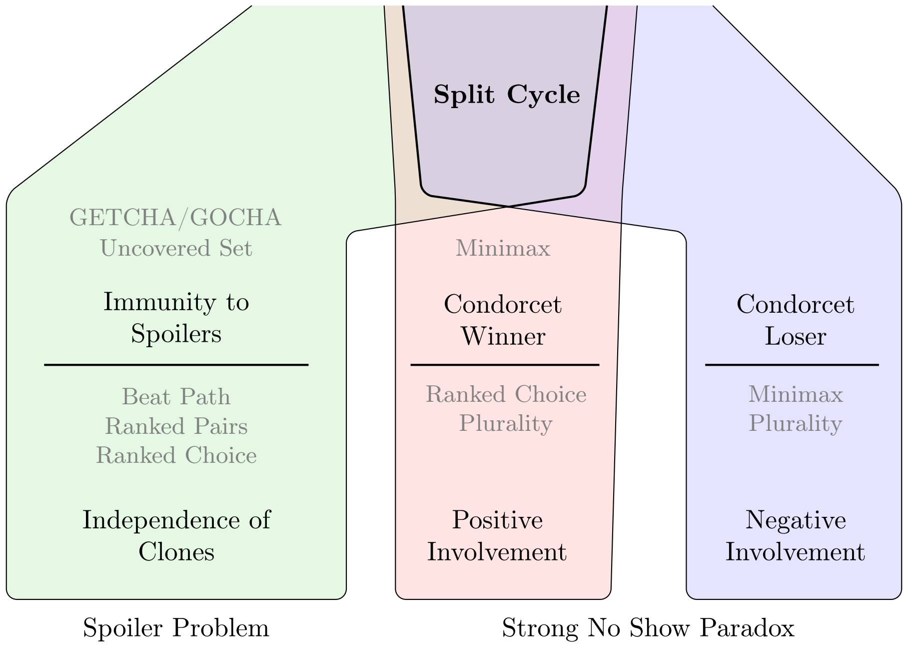

Figure 1: An illustration of how solving the Spoiler Problem and the Strong No Show Paradox leads uniquely to Split Cycle as opposed to standard voting methods, some of which are displayed in gray (and defined in Appendix C). For each of the three "roads" in the diagram and each of the displayed voting methods, a voting method is shown on a road if and only if it satisfies all of the criteria that appear in the road lower in the diagram than the voting method. For example, Minimax satisfies Condorcet Winner and Positive Involvement, explaining its location on the middle road; it satisfies Negative Involvement but not Condorcet Loser, explaining its location on the right road; and it satisfies Immunity to Spoilers but not Independence of Clones, explaining its absence from the left road. All voting methods except Split Cycle are blocked from entering the overlap of the three roads, but some are blocked even earlier by other criteria, as indicated by the horizontal lines. For example, Ranked Choice is blocked by the Condorcet Winner criterion.

1.3 Organization

The rest of the paper is organized as follows. In Section 2, we review some preliminary notions: profiles, margin graphs, and voting methods, as well as operations on profiles. In Section 3, we motivate and define our proposed voting method, Split Cycle. We then assess Split Cycle relative to other standard voting methods, all of which are defined in Appendix C, in a three step process:

- In Section 4, we review the "**core criteria**" that we take for granted in this paper for reducing the class of voting methods to which we compare Split Cycle: the Condorcet winner (and Condorcet loser) criterion, monotonicity, independence of clones, and the narrowing criterion.
- In Section 5, we propose **new criteria**: *immunity to spoilers*, which we strengthen to *stability for winners* and then to a more general criterion of *amalgamation* (closely related to Sen's (1971; 1993) choice-functional condition of *expansion* or γ); and a criterion of *rejectability*, which concerns winnowing down a set of winners to a unique winner. Of the known methods left standing in step one, only Split Cycle satisfies immunity to spoilers and rejectability.

- In Section 6, we test Split Cycle against **other criteria** from the literature: the Smith and Schwartz criteria, independence of Smith-dominated alternatives, reversal symmetry, resolvability, and positive and negative involvement. Split Cycle satisfies all these criteria except the Schwartz criterion and resolvability in either of its two flavors—single-voter resolvability or asymptotic resolvability. However, we argue that neither the Schwartz criterion nor resolvability is normatively compelling. We claim that rejectability is a more compelling resoluteness criterion than single-voter resolvability, and we show that both single-voter and asymptotic resolvability are incompatible with stability for winners (assuming another weak axiom). We also discuss data from computer simulations (given in Appendix D) suggesting that for small numbers of candidates, Split Cycle performs well in terms of the frequency of multiple winners and the average number of multiple winners before tiebreaking. Finally, our proofs that Split Cycle satisfies positive involvement and negative involvement provide two more ways of singling out Split Cycle from the known methods left standing in step one (again see Figure 1).

Our axiomatic analysis of Split Cycle and other methods is summarized in Figure 2. We conclude in Section 7 with a brief summary and directions for further research on Split Cycle.

Remark 1.3. An implementation in Python of the Split Cycle voting method and other methods referenced in this paper is available at https://github.com/epacuit/splitcycle. All of the examples in the paper have been verified in a Jupyter notebook available in the linked repository.

2 Preliminaries

2.1 Profiles, Margin Graphs, and Voting Methods

Fix infinite sets $\mathcal{V}$ and $\mathcal{X}$ of *voters* and *candidates*, respectively. A given election will involve only finite subsets $V \subseteq \mathcal{V}$ and $X \subseteq \mathcal{X}$, but we want no upper bound on the number of voters or candidates who may participate in elections. For $X \subseteq \mathcal{X}$, let $\mathcal{L}(X)$ be the set of all (strict) linear orders on X.

Definition 2.1. A *profile* is a function $\mathbf{P} : V(\mathbf{P}) \to \mathcal{L}(X(\mathbf{P}))$ for some nonempty finite $V(\mathbf{P}) \subseteq \mathcal{V}$ and nonempty finite $X(\mathbf{P}) \subseteq \mathcal{X}$. We call $V(\mathbf{P})$ and $X(\mathbf{P})$ the sets of *voters in* $\mathbf{P}$ and *candidates in* $\mathbf{P}$, respectively. We call $\mathbf{P}(i)$ voter i's *ballot*, and we write '$x\mathbf{P}_i y$' for $(x, y) \in \mathbf{P}(i)$.

As usual, we take $x\mathbf{P}_i y$ to mean that voter i strictly prefers candidate x to candidate y.

Remark 2.2. None of the main results to follow depends on the assumption that voters submit linear orders of all the candidates, as opposed to strict weak orders (allowing ties) or truncated orders (allowing some unranked candidates). We make the linearity assumption only to bypass some less important choice points involving how a voting method with a canonical definition for profiles of linear ballots may have several formulations for ballots allowing ties or truncations. The definition of Split Cycle to be given in Definitions 3.3 and 3.5 can applied without modification to a profile $\mathbf{P}$ of strict weak orders or truncated orders. However, when ties or truncation are allowed, one may consider not only the majority *margin* between candidates, as defined below, but other measures of strength of majority preference, such as ratio, winning votes, and losing votes (see Schulze 2011, § 2.1). We defer discussion of these subtleties to other work.

Next we define the notions of an abstract margin graph and the margin graph of a particular profile.

	Split Cycle	Ranked Pairs	Beat Path	Mini-max	Copeland	GETCHA/ GOCHA	Uncovered Set	Ranked Choice	Plurality
Condorcet Winner	✓	✓	✓	✓	✓	✓	✓	–	–
Condorcet Loser	✓	✓	✓	–	✓	✓	✓	✓	–
Monotonicity	✓	✓	✓	✓	✓	✓	✓	–	✓
Independence of Clones	✓	✓*	✓	–	–	✓	✓†	✓	–
Narrowing	✓	✓	✓	✓	✓	–	–*	✓	✓
Immunity to Spoilers	✓	–	–	✓	✓	✓	✓	–	–
Stability for Winners	✓	–	–	–	–	✓	✓	–	–
Amalgamation (Expansion, γ)	✓	–	–	–	–	✓/–	✓†	–	–
Rejectability	✓	✓	✓	✓	–	–	–	✓	✓
Smith	✓	✓	✓	–	✓	✓	✓	–	–
ISDA	✓	✓	✓	–	✓	✓	✓	–	–
Resolvability	–	✓	✓	✓	–	–	–	✓	✓
Reversal Symmetry	✓	✓	✓	–	✓	✓	✓	–	–
Positive Involvement	✓	–	–	✓	–	–	–	✓	✓
Negative Involvement	✓	–	–	✓	–	–	–	–	✓

Figure 2: Comparison of voting methods in terms of selected voting criteria. A ✓ indicates that the criterion is satisfied, while – indicates that it is not. The ✓* indicates that there are subtleties in how one must define the Ranked Pairs method to ensure full independence of clones (together with anonymity), as discussed in Remark 4.11. For the Uncovered Set column, there are several definitions of the Uncovered Set that are equivalent for an odd number of voters but inequivalent for an even number of voters; the ✓† indicates that while one version of the Uncovered Set (Fishburn 1977) fails to satisfy independence of clones and amalgamation for profiles with an even number of voters, other definitions satisfy both axioms for all profiles, and all definitions do so for profiles with an odd number of voters (see Appendix C.7). The –* indicates that Uncovered Set violates the narrowing criterion for tournaments but whether it satisfies the narrowing criterion for profiles has not been settled (see Section 4.4). For the claims about voting methods other than Split Cycle, see Appendix C.

Definition 2.3. A *margin graph* is a weighted directed graph such that: the edge relation of the graph is asymmetric; either all weights of edges are even positive integers or all weights of edges are odd positive integers; and if there are two nodes with no edge between them, then all weights are even.

An example of a margin graph already appeared in Section 1.1.

Definition 2.4. Let $\mathbf{P}$ be a profile and $a, b \in X(\mathbf{P})$. Then

$$Margin_{\mathbf{P}}(a,b) = |\{i \in V(\mathbf{P}) \mid a\mathbf{P}_i b\}| - |\{i \in V(\mathbf{P}) \mid b\mathbf{P}_i a\}|.$$

The *margin graph of* $\mathbf{P}$, $\mathcal{M}(\mathbf{P})$, is the weighted directed graph whose set of nodes is $X(\mathbf{P})$ with an edge from a to b weighted by $Margin_{\mathbf{P}}(a,b)$ when $Margin_{\mathbf{P}}(a,b) > 0$, in which case we say that a *is majority preferred to* b. We write

$$a \xrightarrow{\alpha}_{\mathbf{P}} b \text{ if } \alpha = Margin_{\mathbf{P}}(a,b) > 0.$$

We may omit the α when the size of the margin is not important and $\mathbf{P}$ when the profile in question is clear. We call the unweighted directed graph underlying $\mathcal{M}(\mathbf{P})$ the *majority graph* of $\mathbf{P}$, denoted $M(\mathbf{P})$.

The key fact about the relation between margin graphs and profiles is given by Debord's Theorem.

Theorem 2.5 (Debord 1987)**.** For any margin graph $\mathcal{M}$, there is a profile $\mathbf{P}$ such that $\mathcal{M}$ is the margin graph of $\mathbf{P}$.

Finally, we define what we mean by a voting method for the purposes of this paper.

Definition 2.6. A *voting method* is a function F on the domain of all profiles such that for any profile $\mathbf{P}$, $\varnothing \neq F(\mathbf{P}) \subseteq X(\mathbf{P})$. We call $F(\mathbf{P})$ the *set of winners* or *winning set for* $\mathbf{P}$ *under* F.

As usual, if $F(\mathbf{P})$ contains multiple winners, we assume that some further tie-breaking process would then apply, though we do not fix the nature of this process (see Schwartz 1986, pp. 14-5 for further discussion). Options include the use of an even-chance lottery on $F(\mathbf{P})$ or a runoff election with the candidates in $F(\mathbf{P})$ in which a partially different set of voters may participate.

2.2 Operations on Profiles

Sometimes we will be interested in combining two profiles for the same set of candidates and disjoint sets of voters, for which we use the following notation.

Definition 2.7. Given profiles $\mathbf{P}$ and $\mathbf{P}'$ such that $X(\mathbf{P}) = X(\mathbf{P}')$ and $V(\mathbf{P}) \cap V(\mathbf{P}') = \varnothing$, we define the profile $\mathbf{P} + \mathbf{P}' : V(\mathbf{P}) \cup V(\mathbf{P}') \to \mathcal{L}(X(\mathbf{P}))$ by

$$(\mathbf{P} + \mathbf{P}')(i) = \begin{cases} \mathbf{P}(i) & \text{if } i \in V(\mathbf{P}) \\ \mathbf{P}'(i) & \text{if } i \in V(\mathbf{P}') \end{cases}.$$

To add $\mathbf{P}$ to itself, we may take $\mathbf{P} + \mathbf{P}^*$ where $\mathbf{P}^*$ is a copy of $\mathbf{P}$ with a disjoint set of voters.[11]

[11] I.e., $X(\mathbf{P}) = X(\mathbf{P}^*)$, $V(\mathbf{P}) \cap V(\mathbf{P}^*) = \varnothing$, and there is a bijection $h : V(\mathbf{P}) \to V(\mathbf{P}^*)$ such that for all $i \in V(\mathbf{P})$ and $x, y \in X(\mathbf{P})$, we have $x\mathbf{P}_i y$ if and only if $x\mathbf{P}^*_{h(i)} y$.

We will also be interested in deleting a given candidate from every ballot in a profile, as follows.

Definition 2.8. Given a profile $\mathbf{P}$ and $x \in X(\mathbf{P})$, let $\mathbf{P}_{-x}$ be the profile with $X(\mathbf{P}_{-x}) = X(\mathbf{P}) \setminus \{x\}$ and $V(\mathbf{P}_{-x}) = V(\mathbf{P})$ such that for all $i \in V(\mathbf{P})$, $\mathbf{P}_{-x}(i)$ is the restriction of $\mathbf{P}(i)$ to $X(\mathbf{P}_{-x})$.

3 Split Cycle

3.1 Three Main Ideas

The "Paradox of Voting" is the phenomenon that cycles may occur in the margin graph of a profile, e.g., a is majority preferred to b, b is majority preferred to c, and c is majority preferred to a. Recall the formal definition of a cycle.

Definition 3.1. Given a directed graph $\mathcal{G}$ (e.g., a margin graph), a *cycle in* $\mathcal{G}$ is a sequence $\langle x_1, \ldots, x_n \rangle$ of nodes from $\mathcal{G}$ such that $n > 1$, $x_1 = x_n$, and for all $i \in \{1, \ldots, n-1\}$, we have $x_i \to x_{i+1}$, where $\to$ is the edge relation of the graph. The cycle is *simple* if for all distinct $i, j \in \{1, \ldots, n\}$, $x_i = x_j$ only if $i, j \in \{1, n\}$ (i.e., all nodes are distinct except $x_1 = x_n$).

The voting method we propose in this paper, Split Cycle, provides a way of dealing with the problem of majority cycles. It is based on three main ideas:

1. Group incoherence raises the threshold for one candidate to defeat another, but not infinitely. By "group incoherence" we mean cycles in the majority relation. Consider the margin graph on the left again:

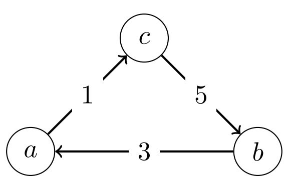

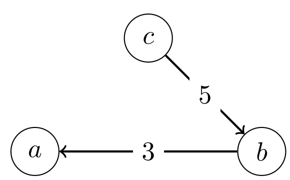

Due to the group incoherence, the margin of 1 for a over c is not sufficient for a to defeat c. But if we raise the threshold for defeat to *winning by more than* 1, and we redraw the graph with an arrow from x to y if and only if $Margin_{\mathbf{P}}(x, y) > 1$, as on the right, then the group is no longer incoherent at this threshold. Since the group is no longer incoherent with respect to the *win by more than* 1 relation, we think it is reasonable to take c to defeat b and b to defeat a, leaving c as the winner. Thus, as suggested, group incoherence does not raise the threshold for b to defeat a *infinitely* but rather only enough to eliminate any incoherence in which b and a are involved. This shows that our proposal differs from the GETCHA and GOCHA methods (Section 6.1), which take all 3-cycles to result in three-way ties regardless of the margins.[12]

2. Incoherence can be localized. Consider the following margin graph:

[12]Cf. Tideman (1987, p. 206): "The GOCHA rule, in a sense, is only half a voting rule. It does not address the issue of what should be done to resolve cycles."

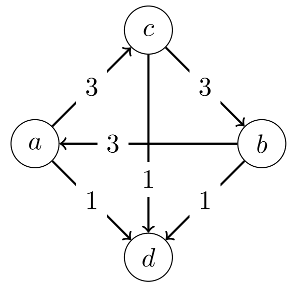

It would be a mistake to think that the margin of 1 for a over d is not sufficient for a to defeat d, due to the incoherence involving a, b, and c, which is only eliminated by raising the threshold to *win by more than 3*. For there is no incoherence with respect to d and the other candidates, all of whom are majority preferred to d, so they all defeat d. The lesson from this example is that when deciding whether the margin of a over d is sufficient for a to defeat d, we set the threshold in terms of *the cycles* (*if any*) *involving* a *and* d. This shows that our proposal differs from the Minimax method, which takes the winner in the example above to be the Condorcet loser d (see Definition 4.3).

3. Defeat is direct. On our view, for a candidate x to defeat a candidate y, so that y is not in the set of winners, x must have a positive margin over y. Consider the following margin graph (note that if there is no edge between two candidates, then the margin of each candidate over the other is 0):

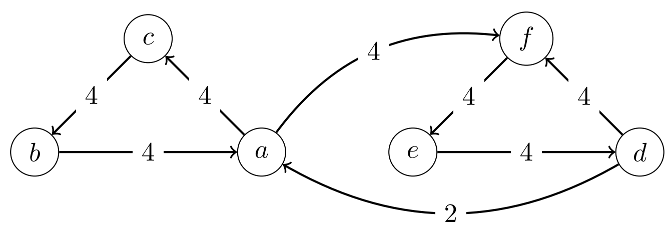

In this case, we think a should defeat f, but a should not defeat d. Some other voting methods, such as Beat Path, commit one to a view that we find dubious: that even though a is not majority preferred to d, nonetheless a should kick d out of the set of winners because of the indirect path from a to f to e to d with margins of 4 at each step. By contrast, we adopt a direct pairwise perspective: *for* a *to kick* d *out of the winning set,* a *must be majority preferred to* d. We find it difficult to try to explain to d's supporters that although a was not majority preferred to d, nonetheless a kicks d out of the winning set because of a's relation to *other candidates*, f and e, *neither of whom defeat* d![13] Of course reasonable definitions of defeat cannot fully satisfy the independence of irrelevant alternatives (IIA) criterion (Arrow 1963),[14] but in our view this seems too flagrant a violation of the idea behind IIA. We endorse the following weakening of IIA, known as *weak IIA* (Baigent 1987): if two profiles are alike with respect to how everyone votes on x vs. y, then it should not be possible that in one profile, x defeats y, while in the other, y defeats x (though it should be possible that in one, x defeats y, while in the other, neither x defeats y nor y defeats x, due to a cycle). Let $\mathbf{P}$ be a profile whose margin graph is shown above, and let $\mathbf{P}'$ be a profile just like $\mathbf{P}$ with respect to how everyone votes on a vs. d but in which all voters have either a followed by d or d followed by a at the top

[13] We assume that e does not defeat d because of the perfect cycle involving d, f, and e.

[14] Here we take IIA to state that if two profiles are alike with respect to how everyone votes on x vs. y, then x defeats y in the one profile if and only if x defeats y in the other.

of their ballots, followed by the linear order $b\mathbf{P}'_i c\mathbf{P}'_i e\mathbf{P}'_i f$. In $\mathbf{P}'$, since d is majority preferred to a by 2 and there are no cycles, surely d should defeat a, kicking a out of the winning set. Then it follows by weak IIA that in $\mathbf{P}$, a does not defeat d. Thus, weak IIA is inconsistent with the indirect notion of defeat according to Beat Path. By contrast, it is satisfied by the direct notion of defeat we will define for Split Cycle.[15]

3.2 Defining Split Cycle

To define Split Cycle, in line with our first idea above, we first measure the degree of incoherence of a cycle by the smallest margin occurring on an edge in the cycle—for if we raise our threshold above that margin, then we split the cycle, restoring coherence at the higher threshold as in the second graph in Section 3.1.

Definition 3.2. Let $\mathbf{P}$ be a profile and ρ a simple cycle in $\mathcal{M}(\mathbf{P})$. The *splitting number* of ρ, $Split\#_{\mathbf{P}}(\rho)$, is the smallest margin between consecutive candidates in ρ (e.g., the splitting number of $a \overset{3}{\to} b \overset{1}{\to} c \overset{5}{\to} a$ is 1). We omit the subscript for $\mathbf{P}$ when the profile is clear from context.

Thus, for example, the splitting number of the cycle in the three-candidate margin graph in Section 3.1 is 1, while the splitting number of the cycle in the four-candidate margin graph in Section 3.1 is 3.

In line with our second idea that incoherence can be localized, when deciding whether a defeats b, we look at all and only the simple cycles containing a and b (not at the other cycles that do not contain a and b); and in line with our third idea about the directness of defeat, for a to defeat b, we require that *the direct margin of a over b* exceeds the splitting number of every simple cycle containing a and b, which means that that direct margin survives after we raise the threshold above those splitting numbers.

Definition 3.3. Let $\mathbf{P}$ be a profile and $a, b \in X(\mathbf{P})$. Then *a defeats b in $\mathbf{P}$* if $Margin_{\mathbf{P}}(a,b) > 0$ and

$$Margin_{\mathbf{P}}(a,b) > Split\#(\rho) \text{ for every simple cycle } \rho \text{ in } \mathcal{M}(\mathbf{P}) \text{ containing } a \text{ and } b.$$

A candidate b is *undefeated in* $\mathbf{P}$ if there is no candidate who defeats b.

Remark 3.4. Just as some sports have a *win by* 2 rule for defeat, Split Cycle says that for a to defeat b, a must win by *more than* n over b, where n is the smallest number such that there are no cycles involving a and b in the *win by more than* n relation. Formally, given a profile $\mathbf{P}$ and $n \in \mathbb{N}$, define a binary relation $W^n_{\mathbf{P}}$ on $X(\mathbf{P})$ by $xW^n_{\mathbf{P}}y$ ("x wins by more than n over y") if $Margin_{\mathbf{P}}(x,y) > n$. Then Definition 3.3 is equivalent to: a defeats b in $\mathbf{P}$ if $aW^n_{\mathbf{P}}b$ for the smallest $n \in \mathbb{N}$ such that there is no simple cycle containing a and b in the graph whose set of nodes is $X(\mathbf{P})$ and whose edge relation is $W^n_{\mathbf{P}}$.

Finally, we can define the voting method we call Split Cycle:

Definition 3.5. For any profile $\mathbf{P}$, the set of Split Cycle winners, $SC(\mathbf{P})$, is the set of candidates who are undefeated in $\mathbf{P}$.

As explained in Section 1, one can determine $SC(\mathbf{P})$ in a simple two-step process (see Footnote 18 for a faster algorithm): 1. For each simple cycle, identify the edges with the smallest margin in that cycle. 2. After completing step 1 for all simple cycles, discard the identified edges. All remaining edges count as defeats.

[15]In fact, in Holliday and Pacuit 2020, we characterize the Split Cycle defeat relation using an axiom of Coherent IIA that is stronger than weak IIA.

Remark 3.6. Since the only information Split Cycle uses about a profile $\mathbf{P}$ is its margin graph, we can also think of Split Cycle as assigning to each margin graph $\mathcal{M}$ a set $SC(\mathcal{M})$ of winners.

Let us consider some examples of calculating the set of Split Cycle winners.

Example 3.7. The Split Cycle winners for the margin graphs illustrating our three main ideas in Section 3.1 are as follows:

1. In the three-candidate example, the unique Split Cycle winner is c;

2. In the four-candidate example, the Split Cycle winners are a, b, and c;

3. In the six-candidate example, the Split Cycle winners are all candidates except f.

Example 3.8. For a more complicated example, consider the following margin graph, repeated three times to highlight the three different simple cycles:

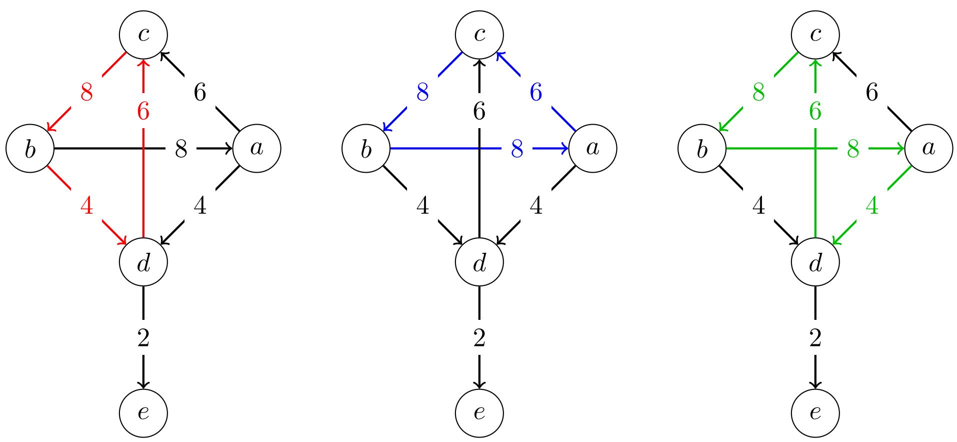

The splitting number of the cycle $b \to d \to c \to b$ is 4; the splitting number of the cycle $b \to a \to c \to b$ is 6; and the splitting number of the cycle $b \to a \to d \to c \to b$ is 4. In each cycle, the edge with the smallest margin in that cycle is not a defeat. After discarding these edges (i.e., the $b \to d$ edge in the red cycle, the $a \to c$ edge in the blue cycle, and the $a \to d$ edge in the green cycle), the remaining edges are defeats:

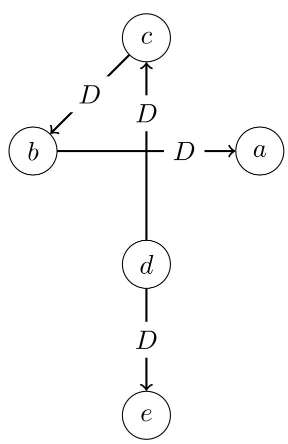

Since d is the only undefeated candidate, d is the unique Split Cycle winner.

Let us now show that the set of Split Cycle winners is always nonempty.

Lemma 3.9. For a profile $\mathbf{P}$, let the *defeat graph of* $\mathbf{P}$ be the directed graph whose set of nodes is $X(\mathbf{P})$ with an edge from a to b when a defeats b in $\mathbf{P}$. Then for any profile $\mathbf{P}$, the defeat graph of $\mathbf{P}$ contains no cycles. Thus, $SC(\mathbf{P}) \neq \varnothing$.

Proof. Suppose there is a cycle $a_1 D a_2 D \dots D a_n D a_1$ in the defeat graph of $\mathbf{P}$, which we may assume is simple (since if there is any cycle, there is a simple one). This yields a simple cycle $\rho = a_1 \xrightarrow{\alpha_1} a_2 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_{n-1}} a_n \xrightarrow{\alpha_n} a_1$ in the margin graph of $\mathbf{P}$ where each margin α_i is greater than the splitting number of any simple cycle containing $a_i, a_{i+1 \bmod n}$ and hence greater than the splitting number of ρ itself, which is impossible. □

Remark 3.10. Like defeat relations in sports tournaments, the Split Cycle defeat relation is not necessarily transitive: it may be, as in Example 3.8, that d defeated c, and c defeated b, while d is not among those who defeated b—nonetheless, b is not among the winners of the tournament, having been defeated by c. *Acyclicity*, as in Lemma 3.9, is sufficient for there always to be a nonempty set of winners—transitivity is not required. That the Split Cycle defeat relation is acyclic but not necessarily transitive explains how it can satisfy weak IIA without contradicting Baigent's (1987) generalization of Arrow's impossibility theorem (cf. Campbell and Kelly 2000), which states that under Arrow's axioms but with IIA weakened to weak IIA, there must be a weak dictator (a voter i such that if i prefers x to y, then y does not defeat x socially). Baigent's theorem requires that the social defeat relation is not only acyclic but a strict weak order (i.e., asymmetric and negatively transitive, which implies transitive).[16] Implicit here is that we can view Split Cycle as a *collective choice rule*, i.e., a function mapping each profile $\mathbf{P}$ to a binary relation on $X(\mathbf{P})$ (cf. Sen 2017, Ch. 2*), by taking the binary relation to be the defeat relation. This is the perspective on Split Cycle adopted in Holliday and Pacuit 2020. However, in this paper we focus on Split Cycle as a voting method (as in Definition 2.6) that maps each profile to a set of winners.

Another useful lemma about Split Cycle is that if a candidate z is not a winner for a profile $\mathbf{P}$, then there is some winner x and a path in the defeat graph of $\mathbf{P}$ from x to z.

Lemma 3.11. For any profile $\mathbf{P}$ and $z \in X(\mathbf{P}) \backslash SC(\mathbf{P})$, there is an $x \in SC(\mathbf{P})$ and distinct $y_1, \dots, y_n \in X(\mathbf{P})$ with $y_1 = x$ and $y_n = z$ such that $y_1 D y_2 D \dots D y_{n-1} D y_n$.

Proof. We first find $w_1, \dots, w_n \in X(\mathbf{P})$ such that $w_n D w_{n-1} D \dots D w_2 D w_1$ and then relabel $w_1, \dots, w_n$ as $y_n, \dots, y_1$, so that $y_1 D y_2 D \dots D y_{n-1} D y_n$. If $z \in X(\mathbf{P}) \setminus SC(\mathbf{P})$, then setting $w_1 = z$, there is a w_2 such that $w_2 D z$. If $w_2 \in SC(\mathbf{P})$, then we are done with $x = w_2$; otherwise, there is a w_3 such that $w_3 D w_2$; and so on. Since $X(\mathbf{P})$ is finite and there are no cycles in the defeat graph of $\mathbf{P}$ by Lemma 3.9, we eventually find the desired $w_n \in SC(\mathbf{P})$. □

Yet another useful lemma about Split Cycle is that to check whether a defeats b, it suffices to check the splitting number of just the simple cycles in which b immediately follows a, rather than all simple cycles containing a and b.

[16] Baigent's theorem also assumes that profiles assign strict weak orders to voters, not just linear orders, but his result also holds for the domain of all linear profiles (also see Campbell and Kelly 2000).

Lemma 3.12. Let $\mathbf{P}$ be a profile and $a, b \in X(\mathbf{P})$. Then a defeats b in $\mathbf{P}$ if and only if $Margin_{\mathbf{P}}(a, b) > 0$ and

$$Margin_{\mathbf{P}}(a, b) > Split\#(\rho) \text{ for every simple cycle } \rho \text{ in } \mathcal{M}(\mathbf{P}) \text{ of the form } a \to b \to x_1 \to \cdots \to x_n \to a.$$

Proof. Obviously if $Margin_{\mathbf{P}}(a, b)$ is greater than the splitting number of every simple cycle containing a and b, then it is greater than the splitting number of every simple cycle of the form $a \to b \to x_1 \to \cdots \to x_n \to a$. Conversely, assume $Margin_{\mathbf{P}}(a, b)$ is greater than the splitting number of every simple cycle of the form $a \to b \to x_1 \to \cdots \to x_n \to a$. To show that $Margin_{\mathbf{P}}(a, b)$ is greater than the splitting number of every simple cycle containing a and b, let ρ be a simple cycle containing a and b whose splitting number is maximal among all such cycles. If ρ contains $a \to b$, then we are done. So suppose ρ does not contain $a \to b$. Without loss of generality, we may assume ρ is of the form $b \to x_1 \to \cdots \to x_n \to a \to y_1 \to \cdots \to y_m \to b$. Let ρ' be $a \to b \to x_1 \to \cdots \to x_n \to a$. It follows from our initial assumption that the splitting number of ρ' is not equal to the margin of the $a \to b$ edge. Since ρ has maximal splitting number of any simple cycle containing a and b, it follows that one of the edges in ρ' after the $a \to b$ edge has this splitting number as its margin; for if none of the edges in ρ' after the $a \to b$ edge has this splitting number as its margin, then since the splitting number is defined as a minimum, ρ' has a higher splitting number than ρ, contradicting the fact that ρ has maximal splitting number of any simple cycle containing a and b. Thus, ρ' has splitting number greater than or equal to that of ρ, and by assumption $Margin_{\mathbf{P}}(a, b) > Split\#(\rho')$, so we have $Margin_{\mathbf{P}}(a, b) > Split\#(\rho)$. Thus, $Margin_{\mathbf{P}}(a, b)$ is greater than the splitting number of every simple cycle containing a and b. □

Remark 3.13. After posting a draft of this paper, we learned from Markus Schulze that Lemma 3.12 relates Split Cycle to the notion of *immunity to binary arguments* in Heitzig 2002. In particular, Split Cycle (along with Beat Path and Ranked Pairs) satisfies all of Heitzig's axioms (Im_{M_α}) for $1/2 < \alpha \leq 1$. Although when defining choice rules, Heitzig (2002, Lemma 2 and following) only defines rules based on his notion of *strong* immunity to binary arguments,[17] which includes Beat Path (in his notation, the rule that selects the common optimal elements of the chain $\{\mathrm{tr}_S(M_\alpha) \mid \frac{1}{2} < \alpha \leq 1\}$), not Split Cycle, it would have been natural in his setting to consider the Split Cycle rule formulated as in Lemma 3.12 as well. Heitzig's axioms (Im_{M_α}) are also closely related to the notion of a *stack* from Zavist and Tideman 1989, defined in Appendix C. Subsequently we learned from Jobst Heitzig of his notion of the "immune set" discussed in a 2004 post on the Election-Methods mailing list (Heitzig 2004), which is equivalent to the Split Cycle winning set after replacing 'stronger' with 'at least as strong' in Heitzig's definition in the post.

It will facilitate reasoning about the defeat relation to introduce one more convenient piece of notation.

Definition 3.14. Let $\mathbf{P}$ be a profile and $a, b \in X(\mathbf{P})$. The *cycle number of a and b in $\mathbf{P}$* is

$$Cycle\#_{\mathbf{P}}(a, b) = \max(\{0\} \cup \{Split\#(\rho) \mid \rho \text{ a simple cycle of the form } a \to b \to x_1 \to \cdots \to x_n \to a\}).$$

Then we can equivalently rewrite the definition of the defeat relation as follows.

[17] Compare Heitzig's notion of strong immunity to Schwartz's (1986) characterization of GOCHA in Lemma 6.5 below.

Lemma 3.15. Let $\mathbf{P}$ be a profile and $a, b \in X(\mathbf{P})$. Then a *defeats* b *in* $\mathbf{P}$ if and only if

$$Margin_{\mathbf{P}}(a,b) > Cycle\#_{\mathbf{P}}(a,b).$$

We will often apply Lemmas 3.12 and 3.15 in proofs without comment.

3.3 Comparison with Beat Path and Ranked Pairs

Lemma 3.12 allows us to relate Split Cycle to the Beat Path voting method as follows (see Appendix C.2 for the definition of strengths of paths).

Lemma 3.16. Let $\mathbf{P}$ be a profile and $a, b \in X(\mathbf{P})$. Then a defeats b in $\mathbf{P}$ if and only if $Margin_{\mathbf{P}}(a,b) > 0$ and

$$Margin_{\mathbf{P}}(a,b) > \text{the strength of the strongest path from } b \text{ to } a.$$

Proof. By Lemma 3.12, a defeats b if and only if $Margin_{\mathbf{P}}(a,b)$ is greater than 0 and the splitting number of every simple cycle of the form $a \to b \to x_1 \to \cdots \to x_n \to a$. But this is equivalent to $Margin_{\mathbf{P}}(a,b)$ being greater than 0 and the strength of every path of the form $b \to x_1 \to \cdots \to x_n \to a$, which is equivalent to $Margin_{\mathbf{P}}(a,b)$ being greater than 0 and the strength of the strongest path from b to a. □

By contrast, the condition for a to defeat b according to Beat Path is that the strength of the strongest path from a to b is greater than the strength of the strongest path from b to a.[18]

Lemma 3.17. For any profile $\mathbf{P}$, $BP(\mathbf{P}) \subseteq SC(\mathbf{P})$, where BP is the Beat Path method (Appendix C.2).

Proof. Suppose $a \notin SC(\mathbf{P})$, so there is a $b \in X(\mathbf{P})$ such that b defeats a according to Split Cycle. Hence $Margin_{\mathbf{P}}(b,a)$ is greater than the strength of the strongest path from a to b by Lemma 3.16. Since $b \to a$ is a path from b to a, it follows that the strength of the strongest path from b to a is greater than the strength of the strongest path from a to b. Hence b defeats a according to Bath Path, so $a \notin BP(\mathbf{P})$. □

We can prove an analogous lemma for the Ranked Pairs method RP (Appendix C.1).

Lemma 3.18. For any profile $\mathbf{P}$, $RP(\mathbf{P}) \subseteq SC(\mathbf{P})$.

Proof. Suppose $a \notin SC(\mathbf{P})$, so there is some $b \in X(\mathbf{P})$ such that $Margin_{\mathbf{P}}(b,a) > 0$ and $Margin_{\mathbf{P}}(b,a) > Split\#(\rho)$ for every simple cycle ρ containing b and a. Now suppose for contradiction that $a \in RP(\mathbf{P})$. Then by Lemma C.1 in Appendix C.1, there are distinct $y_1, \dots, y_m \in X(\mathbf{P})$ such that $a \xrightarrow{\alpha_0} y_1 \to \cdots \to y_m \xrightarrow{\alpha_m} b$ with $\alpha_i \geq Margin_{\mathbf{P}}(b,a)$ for each $i \in \{0, \dots, m\}$. But then $\rho := b \to a \xrightarrow{\alpha_0} y_1 \to \cdots \to y_m \xrightarrow{\alpha_m} b$ is a simple cycle such that $Margin_{\mathbf{P}}(b,a) \not> Split\#(\rho)$, which is a contradiction. Hence $a \notin RP(\mathbf{P})$. □

Since we can have $BP(\mathbf{P}) \subsetneq SC(\mathbf{P})$ and $RP(\mathbf{P}) \subsetneq SC(\mathbf{P})$ (as Beat Path and Ranked Pairs allow indirect defeat of candidates, which Split Cycle rejects),[19] it follows that Split Cycle is less deterministically "resolute" than Beat Path and Ranked Pairs. Practically, this means that under whatever tiebreaking process is used,

[18]Since Lemma 3.16 allows us to define the Split Cycle defeat relation in terms of the strength of strongest paths, we can efficiently calculate Split Cycle using a modification of the Floyd-Warshall algorithm used by Schulze (2011) to calculate Beat Path. See the Python implementation at https://github.com/epacuit/splitcycle.

[19]Indeed, there are profiles $\mathbf{P}$ such that $BP(\mathbf{P}) \cap RP(\mathbf{P}) = \varnothing$, which implies that $BP(\mathbf{P}) \subsetneq SC(\mathbf{P})$ and $RP(\mathbf{P}) \subsetneq SC(\mathbf{P})$ by Lemmas 3.17 and 3.18.

whether an even-chance lottery on the set of winners, or a runoff election between the winners allowing new voters, or some other tiebreaker, Split Cycle will give a nonzero chance of being the ultimate single winner to some candidates to whom Beat Path and Ranked Pairs assign zero chance.

By allowing a slightly greater role for chance in the election of the ultimate single winner,[20] Split Cycle satisfies desirable axiomatic properties that Beat Path and Ranked Pairs fail—such as the known variable-voter property of positive involvement (Section 6.5), as well as the new variable-candidate property of stability for winners proposed in Section 1.1 and defined formally in Section 5.1. When these properties—positive involvement and stability for winners—fail, a candidate who ought to stay in the winning set (i.e., to have a nonzero chance of being the ultimate single winner) is kicked out by the addition of new voters or candidates. So Beat Path (resp. Ranked Pairs) is more deterministically resolute but in the wrong way. For example, for positive involvement, after some voters who all rank x as their favorite join the election, x may perversely be kicked out of the Beat Path (resp. Ranked Pairs) winning set, while x stays in the Split Cycle winning set. Here Beat Path (resp. Ranked Pairs) is more deterministically resolute but in the wrong way. The moral is that there can be a tradeoff between being more resolute before tiebreaking and satisfying desirable variable-voter and variable-candidate properties.

We were able to quantify this tradeoff in the case of the variable-candidate property of stability for winners: using computer simulations reported in Figures 14 and 15 of Appendix D, we estimated the percentage of profiles, among those in which $BP(\mathbf{P}) \subsetneq SC(\mathbf{P})$, that witness a violation of stability for winners for Beat Path. A profile $\mathbf{P}$ witnessing such a violation means there are $a, b \in X(\mathbf{P})$ such that $a \in BP(\mathbf{P}_{-b})$ (i.e., a would be a Beat Path winner without b in the election), a is majority preferred to b, and yet $a \notin BP(\mathbf{P})$, so the addition of the majority-dispreferred candidate b kicks the majority-preferred a out of the set of Beat Path winners (though a remains among the Split Cycle winners, as Split Cycle satisfies stability for winners). For 7 candidates and numbers of voters ranging from 9 to 1,001, we found that over 50% of sampled profiles in which $BP(\mathbf{P}) \subsetneq SC(\mathbf{P})$ witness a violation of stability for winners for Beat Path. Thus, over half of the cases in which Beat Path is more resolute than Split Cycle can be attributed to Beat Path violating a specific axiom we find desirable. The results are even more dramatic for a probability model according to which not every profile is equally likely—see Figure 15 for the Mallows-0.8 model. According to this model, if we randomly choose a profile for, say, 7 candidates and 101 voters, then if the profile is such that $BP(\mathbf{P}) \subsetneq SC(\mathbf{P})$, the probability that the profile violates stability for winners for Beat Path is nearly 100%.[21]

As for Ranked Pairs, we were unable to test as many combinations of numbers of candidates and voters, due to the computational difficulty of calculating the Ranked Pairs winners (see Brill and Fischer 2012; Wang et al. 2019). In fact, this computational difficulty is an additional cost of Ranked Pairs, besides the axiomatic costs relative to Split Cycle noted above. We were, however, able to test four combinations of numbers of candidates and voters: either 6 or 7 candidates and either 55 or 101 voters. For at least these combinations, sampling 50,000 profiles according to the impartial culture model described in Appendix D, we found that around 85% of those sampled profiles in which $RP(\mathbf{P}) \subsetneq SC(\mathbf{P})$ witness a violation of stability for winners

[20] If the tiebreaking process is a runoff election between the tied candidates, in which different voters may show up to the polls, we take this to be an example of introducing chance, since it is unpredictable what ballots will be submitted in the runoff.

[21] It would also be desirable to know what percentage of profiles, among those in which $BP(\mathbf{P}) \subsetneq SC(\mathbf{P})$, witness a violation of positive involvement for Beat Path, in the sense that there is some $a \in SC(\mathbf{P}) \setminus BP(\mathbf{P})$ and one or more voters who rank a as their favorite in $\mathbf{P}$ such that removing those voters results in a profile $\mathbf{P}'$ in which $a \in BP(\mathbf{P}')$. (While positive involvement is typically stated for a single voter, one can consider a coalitional version of the criterion—see Lemma 6.25.) A difficulty is that removing only a single voter is unlikely to have this effect when there are a large number of voters, and considering many coalitions of voters to remove is computationally intractable. We leave this investigation for future work.

for Ranked Pairs—an even higher percentage than with the same combinations for Beat Path. For 50,000 profiles sampled according to the Mallows-0.8 model, also described in Appendix D, between 90-100% of sampled profiles in which $RP(\mathbf{P}) \subsetneq SC(\mathbf{P})$ witnessed a violation of stability for winners for Ranked Pairs.

Not only does much of the increased deterministic resoluteness of Beat Path and Ranked Pairs relative to Split Cycle come at the expense of violating a desirable property like stability for winners, but also we think that Split Cycle does not sacrifice too much in resoluteness for its gain in desirable properties. It still satisfies resoluteness criteria, such as the narrowing and rejectability criteria we propose in Sections 4.4 and 5.3, not satisfied by GETCHA or GOCHA, for example. Moreover, other quantitative analyses in Appendix D suggest that Split Cycle is not substantially less resolute than Beat Path and Ranked Pairs for small numbers of candidates. Finally, while Lemmas 3.17 and 3.18 imply that Split Cycle is less resolute, they also imply that Split Cycle can be viewed as a kind of compromise between proponents of the Beat Path winners and of the Ranked Pairs winners, as all such winners are honored as Split Cycle winners.

It also follows from Lemmas 3.17 and 3.18 that for any profile in which there is a *unique* Split Cycle winner, that candidate is also the unique Beat Path winner and Ranked Pairs winner. Just as Schulze (2011, § 1) motivates Beat Path as a voting method that very often agrees with the Minimax method but has axiomatic advantages over Minimax (namely, satisfying independence of clones, ISDA, and reversal symmetry), one can view Split Cycle as a method that very often agrees with Beat Path and Ranked Pairs but has axiomatic advantages over those methods (e.g., satisfying stability for winners, positive involvement, and negative involvement), as well as an arguably more compelling normative justification in terms of the three ideas in Section 3.1. One reason we value axiomatic *guarantees* against pathological phenomena, such as spoiler effects or the strong no show paradox, more than suggestions that such phenomena are improbable for a given voting method, is that in many contexts we are highly uncertain about the "true" probability distribution over profiles that determines how likely such pathological phenomena are to occur for a method that makes them possible.[22] In addition, improbable events do happen. For example, while it is often claimed that violation of the Condorcet criterion by Ranked Choice is improbable, it happened in the Burlington mayoral election, and Ranked Choice was subsequently repealed in Burlington. Finally, regardless of probabilistic considerations, the satisfaction of normatively appealing axioms may be seen as evidence of the reasonableness of a rule for selecting winners, while the violation of such axioms may be a hint that something has gone fundamentally wrong.

4 Core Criteria

One of our aims is to characterize Split Cycle among known voting methods satisfying the following core criteria: the Condorcet criterion, monotonicity, independence of clones, and what we call the narrowing criterion. Among well-known voting methods, Beat Path and Ranked Pairs satisfy all of the core criteria. In this section, we prove that Split Cycle also satisfies the core criteria.

Before stating our core criteria, let us recall some even more basic baseline criteria for general-purpose voting methods:

- F satisfies *anonymity* if for any profiles $\mathbf{P}$ and $\mathbf{P}'$ with $V(\mathbf{P}) = V(\mathbf{P}')$ and $X(\mathbf{P}) = X(\mathbf{P}')$ and any

[22] For studies of the probabilities of violations of various axioms under different probabilistic models, see, e.g., Munie et al. 2009; Plassmann and Tideman 2014; Brandt et al. 2019, 2020.

$i, j \in V(\mathbf{P})$, if $\mathbf{P}'$ is obtained from $\mathbf{P}$ by swapping the ballots assigned to i and j, then $F(\mathbf{P}) = F(\mathbf{P}')$;

- F satisfies *neutrality* if for any profiles $\mathbf{P}$ and $\mathbf{P}'$ with $V(\mathbf{P}) = V(\mathbf{P}')$ and $X(\mathbf{P}) = X(\mathbf{P}')$ and any $a, b \in X(\mathbf{P})$, if $\mathbf{P}'$ is obtained from $\mathbf{P}$ by swapping the place of a and b in every voter's ballot in $\mathbf{P}$, then $a \in F(\mathbf{P})$ if and only if $b \in F(\mathbf{P}')$.

Another baseline criterion for general-purpose voting methods is the well-known Pareto principle (see, e.g., Zwicker 2016, Definition 2.6):

- F satisfies *Pareto* if for any profile $\mathbf{P}$ and $a, b \in X(\mathbf{P})$, if all voters in $\mathbf{P}$ rank a above b, then $b \notin F(\mathbf{P})$.

Proposition 4.1. Split Cycle satisfies anonymity, neutrality, and Pareto.

Proof. Anonymity and neutrality are obvious. For Pareto, suppose all voters in $\mathbf{P}$ rank a above b, so $Margin_{\mathbf{P}}(a, b) = |V(\mathbf{P})|$. Since it is impossible to have a cycle $\rho = a \overset{|V(\mathbf{P})|}{\longrightarrow} b \overset{|V(\mathbf{P})|}{\longrightarrow} x_1 \overset{|V(\mathbf{P})|}{\longrightarrow} \dots \overset{|V(\mathbf{P})|}{\longrightarrow} x_n \overset{|V(\mathbf{P})|}{\longrightarrow} a$, one edge of any such cycle must have weight less than $|V(\mathbf{P})|$, so a defeats b by Lemma 3.12. □

All of the other voting methods considered in this paper satisfy anonymity, neutrality, and Pareto, with the exception of GETCHA and GOCHA, which violate Pareto.

4.1 Condorcet Criteria

The first of our core criteria are the *Condorcet criterion*—a candidate who is majority preferred to every other candidate must be the unique winner—and the *Condorcet loser criterion*—a candidate who is majority dispreferred to every other candidate must not be among the winners.

Definition 4.2. For a profile $\mathbf{P}$ and $x \in X(\mathbf{P})$, we say that x is a *Condorcet winner* (resp. *Condorcet loser*) *in* $\mathbf{P}$ if for every $y \in X(\mathbf{P}) \setminus \{x\}$, we have $Margin(x, y) > 0$ (resp. $Margin(y, x) > 0$ and $X(\mathbf{P}) \neq \{x\}$).

Definition 4.3. A voting method F satisfies the *Condorcet criterion* (resp. *Condorcet loser criterion*) if for every profile $\mathbf{P}$ and $x \in X(\mathbf{P})$, if x is the Condorcet winner (resp. Condorcet loser), then $F(\mathbf{P}) = \{x\}$ (resp. $x \notin F(\mathbf{P})$). If F satisfies the Condorcet criterion, we say that F is *Condorcet consistent*.

Proposition 4.4. Split Cycle satisfies the Condorcet criterion and the Condorcet loser criterion.

Proof. If x is the Condorcet winner (resp. loser), then for every $y \in X(\mathbf{P}) \setminus \{x\}$, we have $Margin_{\mathbf{P}}(x, y) > 0$ (resp. $Margin_{\mathbf{P}}(y, x) > 0$). It follows that x is not involved in any cycles, so for every $y \in X(\mathbf{P}) \setminus \{x\}$, we have $Margin_{\mathbf{P}}(x, y) > Cycle\#_{\mathbf{P}}(x, y) = 0$ (resp. $Margin_{\mathbf{P}}(y, x) > Cycle\#_{\mathbf{P}}(y, x) = 0$). Hence x defeats every other candidate, so $SC(\mathbf{P}) = \{x\}$ (resp. is defeated by every other candidate, so $x \notin SC(\mathbf{P})$). □

4.2 Monotonicity

The next of our core criteria is the *monotonity criterion* (also called *non-negative responsiveness* in Tideman 1987): lifting the position of a winner x on voters' ballots cannot result in x becoming a non-winner.

Definition 4.5. For any profiles $\mathbf{P}$ and $\mathbf{P}'$ with $V(\mathbf{P}) = V(\mathbf{P}')$ and $x \in X(\mathbf{P}) = X(\mathbf{P}')$, we say that $\mathbf{P}'$ *is obtained from* $\mathbf{P}$ *by a simple lift of* x if the following conditions hold:

1. for all $a, b \in X(\mathbf{P}) \setminus \{x\}$ and $i \in V(\mathbf{P})$, $a\mathbf{P}_i b$ if and only if $a\mathbf{P}'_i b$;
2. for all $a \in X(\mathbf{P})$ and $i \in V(\mathbf{P})$, if $x\mathbf{P}_i a$ then $x\mathbf{P}'_i a$.

Definition 4.6. A voting method F satisfies *monotonicity* if for every profile $\mathbf{P}$ and $x \in X(\mathbf{P})$, if $x \in F(\mathbf{P})$ and $\mathbf{P}'$ is obtained from $\mathbf{P}$ by a simple lift of x, then $x \in F(\mathbf{P}')$.

Proposition 4.7. Split Cycle satisfies monotonicity.

Proof. Suppose $x \in SC(\mathbf{P})$ and $\mathbf{P}'$ is obtained from $\mathbf{P}$ by a simple lift of x. Since $x \in SC(\mathbf{P})$, for all $y \in X(\mathbf{P})$, y does not defeat x in $\mathbf{P}$, so $Margin_{\mathbf{P}}(y, x) \leq Cycle\#_{\mathbf{P}}(y, x)$. We claim that y does not defeat x in $\mathbf{P}'$ either. Since $\mathbf{P}'$ is obtained from $\mathbf{P}$ by a simple lift of x, we have $Margin_{\mathbf{P}'}(y, x) \leq Margin_{\mathbf{P}}(y, x)$. If $Margin_{\mathbf{P}'}(y, x) \leq 0$, then y does not defeat x in $\mathbf{P}'$, so suppose $Margin_{\mathbf{P}'}(y, x) > 0$. We claim that

$$Cycle\#_{\mathbf{P}'}(y, x) \geq Cycle\#_{\mathbf{P}}(y, x) - (Margin_{\mathbf{P}}(y, x) - Margin_{\mathbf{P}'}(y, x)). \quad (1)$$

For given any simple cycle $\rho = y \overset{\alpha}{\to} x \overset{\beta}{\to} z_1 \overset{\gamma_1}{\to} \dots \overset{\gamma_{n-1}}{\to} z_n \overset{\gamma_n}{\to} y$ in $\mathcal{M}(\mathbf{P})$, by Definition 4.5 we have that $\rho' = y \overset{\alpha'}{\to} x \overset{\beta}{\to} z_1 \overset{\gamma_1}{\to} \dots \overset{\gamma_{n-1}}{\to} z_n \overset{\gamma_n}{\to} y$ is a simple cycle in $\mathcal{M}(\mathbf{P}')$ where $\alpha = Margin_{\mathbf{P}}(y, x)$ and $\alpha' = Margin_{\mathbf{P}'}(y, x)$. Hence $Split\#(\rho') \geq Split\#(\rho) - (Margin_{\mathbf{P}}(y, x) - Margin_{\mathbf{P}'}(y, x))$. This establishes (1), which with $Margin_{\mathbf{P}}(y, x) \leq Cycle\#_{\mathbf{P}}(y, x)$ implies $Margin_{\mathbf{P}'}(y, x) \leq Cycle\#_{\mathbf{P}'}(y, x)$. Hence y does not defeat x in $\mathbf{P}'$. Since y was arbitrary, we conclude that $x \in SC(\mathbf{P}')$. ☐

4.3 Independence of Clones

Another of our core criteria is the *independence of clones* criterion (Tideman 1987) mentioned in Section 1.1. Recall that a set C of two or more candidates is a set of clones if no candidate outside of C appears in between two candidates from C on any voter's ballot.

Definition 4.8. Given a profile $\mathbf{P}$, a set $C \subseteq X(\mathbf{P})$ is a *set of clones in* $\mathbf{P}$ if $2 \leq |C| < |X(\mathbf{P})|$ and for all $c, c' \in C$, $x \in X(\mathbf{P}) \setminus C$, and $i \in V(\mathbf{P})$, if $c\mathbf{P}_i x$ then $c'\mathbf{P}_i x$, and if $x\mathbf{P}_i c$ then $x\mathbf{P}_i c'$.

The independence of clone criterion states that (i) removing a clone from a profile does not change which non-clones belong to the winning set, and (ii) a clone wins in a profile if and only if after removing that clone from the profile, one of the other clones wins in the resulting profile.

Definition 4.9. A voting method F is such that *non-clone choice is independent of clones* if for every profile $\mathbf{P}$, set C of clones in $\mathbf{P}$, $c \in C$, and $a \in X(\mathbf{P}) \setminus C$,

$$a \in F(\mathbf{P}) \text{ if and only if } a \in F(\mathbf{P}_{-c}).$$

F is such that *clone choice is independent of clones* if for every profile $\mathbf{P}$, set C of clones in $\mathbf{P}$, and $c \in C$,

$$C \cap F(\mathbf{P}) \neq \varnothing \text{ if and only if } C \setminus \{c\} \cap F(\mathbf{P}_{-c}) \neq \varnothing.$$

Finally, F satisfies *independence of clones* if F is such that non-clone choice is independent of clones and clone choice is independent of clones.

We prove the following in Appendix A.

Theorem 4.10. Split Cycle satisfies independence of clones.

Remark 4.11. Tideman (1987) shows that the version of Ranked Pairs defined in Appendix C.1 satisfies the condition of independence of clones for all profiles $\mathbf{P}$ such that for all $a, b, x, y \in X(\mathbf{P})$, $Margin_{\mathbf{P}}(a, b) = 0$ only if $a = b$, and $Margin_{\mathbf{P}}(a, b) = Margin_{\mathbf{P}}(x, y)$ only if (i) $a = x$ or a and x belong to a set of clones, and (ii) $b = y$ or b and y belong to a set of clones. Zavist and Tideman (1989) show that the same version of Ranked Pairs does not satisfy independence of clones for all profiles. They propose a modified version of Ranked Pairs that satisfies independence of clones at the expense of violating anonymity. However, as suggested by Tideman (p.c.), one can obtain an anonymous and fully clone-independent version of Ranked Pairs by declaring a candidate x a winner in profile $\mathbf{P}$ if there exists a voter $i \in V(\mathbf{P})$ such that the Zavist and Tideman version of Ranked Pairs declares x a winer in $\mathbf{P}$ when i is the designated voter used to generate their tie-breaking ranking of candidates (TBRC).[23]

4.4 Narrowing

The aim of a voting method is to narrow down the choice of winning candidates. The last of our core criteria is related to Fey's (2008) work on voting methods that "almost never narrow the set of social choices" (p. 302) in large elections (also see Scott and Fey 2012 and Saile and Suksompong 2020). Fey concentrated on *tournament solutions*, which assign to every tournament (i.e., asymmetric directed graph with an edge between each pair of distinct nodes) a nonempty set of winners. We say that a tournament solution fails the *narrowing criterion* (*for tournaments*) if the proportion of tournaments T with $|T| = k$ for which the tournament solution selects all candidates as winners approaches 1 in the limit as k goes to infinity. Regarded as tournament solutions, GETCHA and Uncovered Set fail the narrowing criterion for tournaments (Moon and Moser 1962, Fey 2008).

We can also consider the following narrowing criterion for profiles.

Definition 4.12. A voting method F fails the *narrowing criterion* (*for profiles*) if the proportion of profiles $\mathbf{P}$ with $|V(\mathbf{P})| = n$ and $|X(\mathbf{P})| = k$ for which $F(\mathbf{P}) = X(\mathbf{P})$ approaches 1 in the limit as n goes to infinity and then k goes to infinity. Otherwise F satisfies the narrowing criterion (for profiles).

Bell (1981) showed that GETCHA fails the narrowing criterion for profiles. Of course, in reality the number candidates does not go to infinity,[24] but GETCHA's failure of narrowing for profiles is matched by high levels of irresoluteness—relative to other Condorcet methods satisfying narrowing, such as Copeland and Split Cycle—for realistic numbers of candidates (see Figure 3 and Appendix D). The narrowing criterion for profiles provides a precise way of classifying voting methods—without arbitrary choices of numbers of voters, candidates, and winners—and we conjecture that natural voting methods failing the criterion also fall into the vague category of methods with relatively high levels of irresoluteness.

Split Cycle satisfies an even stronger property than narrowing: the proportion of profiles for which $F(\mathbf{P}) \subsetneq X(\mathbf{P})$ approaches 1 in the limit as the number of voters and then the number of candidates goes to

[23]This formulation takes advantage of our assumption that voters submit linear orders of the candidates. Zavist and Tideman allow ties in voter ballots and use a randomizing device to generate the linear TBRC from the designated voter's ballot, in case it contains ties. Thus, Zavist and Tideman define a probabilistic voting method.

[24]Note that taking the number of voters to infinity only makes the voting method have smaller winning sets, so this idealization works for rather than against the methods we are testing for narrowing.

Voting Method	3	4	5	6	7	8	9	10	20	30
Split Cycle	1	1.01	1.03	1.06	1.08	1.11	1.14	1.16	1.42	1.62
Copeland	1.17	1.26	1.29	1.3	1.31	1.31	1.31	1.31	1.28	1.25
Uncovered Set	1.17	1.35	1.53	1.71	1.9	2.09	2.26	2.46	4.56	6.82
GETCHA	1.17	1.44	1.8	2.21	2.72	3.31	3.94	4.68	13.55	22.94

Figure 3: Estimated average sizes of winning sets for profiles with a given number of candidates (top row) in the limit as the number of voters goes to infinity.

infinity. This is a consequence of the following fact, for which we say that a profile $\mathbf{P}$ is *uniquely weighted* if for all $x, y, x', y' \in X(\mathbf{P})$, if $x \neq y$, $x' \neq y'$, and $(x, y) \neq (x', y')$, then $Margin_{\mathbf{P}}(x, y) \neq Margin_{\mathbf{P}}(x', y')$.

Proposition 4.13. Let $\mathbf{P}$ be a uniquely weighted profile such that $|X(\mathbf{P})| \geq 3$. Then $|SC(\mathbf{P})| \leq |X(\mathbf{P})| - 2$.

Proof. Pick $a, b \in X(\mathbf{P})$ such that $Margin_{\mathbf{P}}(a, b)$ is the highest margin of any edge in $\mathcal{M}(\mathbf{P})$, which implies that $Margin_{\mathbf{P}}(a, b) > 0$. Then clearly a defeats b in $\mathbf{P}$. Now pick $c, d \in X(\mathbf{P})$ with $d \neq b$ such that $Margin_{\mathbf{P}}(c, d)$ is the highest margin in $\mathcal{M}(\mathbf{P})$ of any edge not going to b. Suppose for contradiction that c does not defeat d in $\mathbf{P}$. Then there is a simple cycle ρ containing c and d such that $Margin_{\mathbf{P}}(c, d)$ is strictly less than the other margins along the cycle, at least one of which is a margin of an edge not going to b. But this contradicts the fact that $Margin_{\mathbf{P}}(c, d)$ is the highest margin in $\mathcal{M}(\mathbf{P})$ of any edge not going to b. Thus, c defeats d. Hence $SC(\mathbf{P})$ contains neither b nor d, so $|SC(\mathbf{P})| \leq |X(\mathbf{P})| - 2$. □

Since the proportion of profiles that are uniquely weighted approaches 1 as the number of voters approaches infinity, Proposition 4.13 yields the following.

Corollary 4.14. For each $k \geq 3$, the proportion of profiles $\mathbf{P}$ with $|X(\mathbf{P})| = k$ and $|V(\mathbf{P})| = n$ for which $|SC(\mathbf{P})| \leq |X(\mathbf{P})| - 2$ approaches 1 as n approaches infinity. As a consequence, Split Cycle satisfies the narrowing criterion.

Figure 3 shows estimates for the average sizes of winning sets in the limit as the number of voters goes to infinity for the methods discussed in this section. Estimates were obtained using the Monte Carlo simulation technique described in Harrison-Trainor 2020, § 9, sampling 50,000 profiles for each number of candidates.

5 New Criteria

In the previous section, we saw that Split Cycle belongs to the small family of known voting methods that satisfy the Condorcet criterion, monotonicity, independence of clones, and narrowing. In this section, we propose new voting method criteria that are satisfied by Split Cycle but not by the other known methods satisfying the core criteria. Thus, Split Cycle is distinguished from all other known methods by satisfying both the core criteria above and the new criteria proposed below. In addition, we will see in Section 6.5 that Split Cycle is also distinguished from known methods satisfying the core criteria by satisfying the criteria of positive and negative involvement.

5.1 Immunity to Spoilers and Stability for Winners

As proposed in Section 1.1, the criterion of *immunity to spoilers* states that if a is a winner in an election $\mathbf{P}$, and we change the election to $\mathbf{P}'$ only by adding a candidate b such that more voters prefer a to b than prefer b to a, and b does not become a winner in $\mathbf{P}'$, then a should still be a winner in $\mathbf{P}'$.

Definition 5.1. A voting method F satisfies *immunity to spoilers* if for every profile $\mathbf{P}$ and $a, b \in X(\mathbf{P})$, if $a \in F(\mathbf{P}_{-b})$, $Margin_{\mathbf{P}}(a, b) > 0$, and $b \notin F(\mathbf{P})$, then $a \in F(\mathbf{P})$.

It is easy to check that GETCHA, GOCHA, Uncovered Set, and Minimax satisfy immunity to spoilers. However, Beat Path and Ranked Pairs do not.

Proposition 5.2. Beat Path and Ranked Pairs do not satisfy immunity to spoilers.

Proof. By Debord's Theorem, there is a profile $\mathbf{P}$ whose margin graph is shown on the right below:

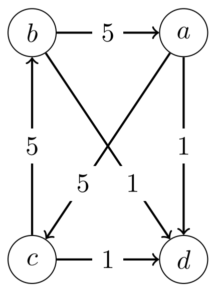

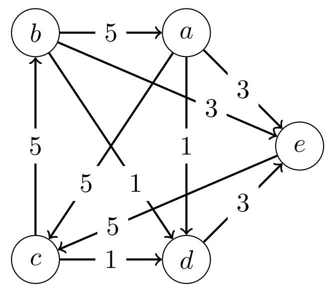

On the left above, we show the margin graph for $\mathbf{P}_{-e}$. Observe that $BP(\mathbf{P}_{-e}) = \{a, b, c\}$ but $BP(\mathbf{P}) = \{d\}$. Thus, the addition of e kicks a and b out of the winning set even though each of a and b are majority preferred to e and e is not among the winners after it is added.

For Ranked Pairs, by Debord's Theorem there is a profile $\mathbf{P}$ whose margin graph is shown on the right below (we make all the margins but two distinct to make the calculation of Ranked Pairs winners easier):

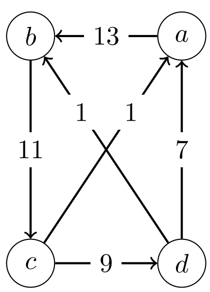

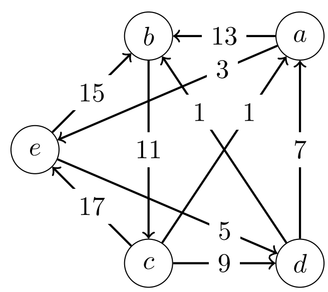

On the left above, we show the margin graph for $\mathbf{P}_{-e}$. In $\mathbf{P}_{-e}$, Ranked Pairs locks in (a, b), then (b, c), then (c, d), and then the remaining pairs are ignored as being inconsistent with those previously locked in, resulting in the linear order *abcd*, so $RP(\mathbf{P}_{-e}) = \{a\}$. However in $\mathbf{P}$, Ranked Pairs locks in (c, e), then (e, b), then (a, b), then ignores (b, c) as inconsistent, then locks in (c, d), then (d, a), then (e, d), then ignores (a, e) as inconsistent, and then locks (d, b) and (c, a), resulting in the linear order *cedab*. Hence $RP(\mathbf{P}) = \{c\}$, so the addition of e kicks a out of the winning set even though a is majority preferred to e and e is not among the winners after it is added. ☐

By contrast, we will show that Split Cycle satisfies immunity to spoilers by showing that it satisfies the following *stronger* and more fundamental property: if a is a winner without b in the election, and a is majority preferred to b, then a is still a winner after the addition of b.

Definition 5.3. A voting method F satisfies *stability for winners* (resp. *strong stability for winners*) if for every profile $\mathbf{P}$ and $a, b \in X(\mathbf{P})$, if $a \in F(\mathbf{P}_{-b})$ and $Margin_{\mathbf{P}}(a, b) > 0$ (resp. $Margin_{\mathbf{P}}(a, b) \geq 0$), then $a \in F(\mathbf{P})$.

Strong stability for winners can be motivated by two of our main ideas from Section 3: since *defeat is direct*, if $Margin_{\mathbf{P}}(a, b) \geq 0$, then b does not defeat a; and since *incoherence raises* (*not lowers!*) *the threshold for defeat*, and adding a candidate can increase incoherence[25] but cannot decrease incoherence in the initial set of candidates, if a candidate x did not defeat a before the addition of b, then it does not defeat a after.

Since stability for winners implies immunity to spoilers, neither Beat Path nor Ranked Pairs satisfy stability for winners. As we will now show, neither does Minimax.

Proposition 5.4. Minimax does not satisfy stability for winners.

Proof. By Debord's Theorem, there is a profile $\mathbf{P}$ whose margin graph is shown on the right below:

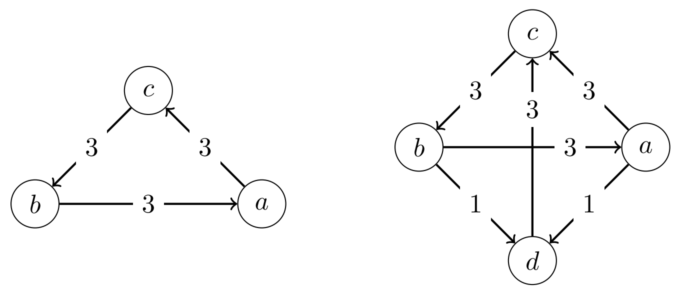

On the left above, we show the margin graph for $\mathbf{P}_{-d}$. Observe that $Minimax(\mathbf{P}_{-d}) = \{a, b, c\}$ but $Minimax(\mathbf{P}) = \{d\}$. Thus, the addition of d kicks a and b out of the winning set even though each of a and b are majority preferred to d. ☐

Note that the proof of Proposition 5.4 provides another demonstration that Beat Path does not satisfy stability for winners, as $BP(\mathbf{P}) = \{d\}$. GOCHA does satisfy stability for winners but not strong stability for winners (see the proof of Proposition 6.7), while GETCHA satisfies both.[26] Uncovered Set satisfies stability for winners, as well as strong stability for winners under some definitions (see Appendix C.7).

As promised, we now prove that Split Cycle satisfies the desired properties.

Proposition 5.5. Split Cycle satisfies strong stability for winners and hence immunity to spoilers.

Proof. If $a \in SC(\mathbf{P}_{-b})$, then for all $c \in X(\mathbf{P}_{-b})$, $Margin_{\mathbf{P}_{-b}}(c, a) \leq Cycle\#_{\mathbf{P}_{-b}}(c, a)$. As $Margin_{\mathbf{P}_{-b}}(c, a) = Margin_{\mathbf{P}}(c, a)$ and $Cycle\#_{\mathbf{P}_{-b}}(c, a) \leq Cycle\#_{\mathbf{P}}(c, a)$, we have that (i) for all $c \in X(\mathbf{P}_{-b})$, $Margin_{\mathbf{P}}(c, a) \leq Cycle\#_{\mathbf{P}}(c, a)$. By the assumption that $Margin_{\mathbf{P}}(a, b) \geq 0$, we have $Margin_{\mathbf{P}}(b, a) \leq 0$ and hence (ii) $Margin_{\mathbf{P}}(b, a) \leq Cycle\#_{\mathbf{P}}(b, a)$. By (i) and (ii), $a \in SC(\mathbf{P})$. ☐

[25] This is why Split Cycle may allow a candidate x who is not among the winners in $\mathbf{P}_{-b}$ to become a winner in $\mathbf{P}$.

[26] To see that GETCHA satisfies strong stability for winners, using the definition of GETCHA in Definition 6.1, suppose $a \in GETCHA(\mathbf{P}_{-b})$ and $Margin_{\mathbf{P}}(a, b) \geq 0$. Further suppose for contradiction that $a \notin GETCHA(\mathbf{P})$, which with $Margin_{\mathbf{P}}(a, b) \geq 0$ implies $b \notin GETCHA(\mathbf{P})$, so $GETCHA(\mathbf{P}) \subseteq X(\mathbf{P}_{-b})$. It follows that $GETCHA(\mathbf{P})$ is $\to_{\mathbf{P}_{-b}}$-dominant, so $GETCHA(\mathbf{P}_{-b}) \subseteq GETCHA(\mathbf{P})$, which contradicts the facts that $a \in GETCHA(\mathbf{P}_{-b})$ and $a \notin GETCHA(\mathbf{P})$.

5.2 Amalgamation

Stability for winners (and immunity to spoilers) concerns when a winner a in a profile $\mathbf{P}_{-b}$ remains a winner in a profile $\mathbf{P}$ with one new candidate b added. The profile $\mathbf{P}$ may be viewed as what we call an *amalgamation* of the profile $\mathbf{P}_{-b}$ without b and the two-candidate profile $\mathbf{P}_{ab}$ with $X(\mathbf{P}_{ab}) = \{a, b\}$ that agrees with $\mathbf{P}$ on the ranking of a vs. b assigned to each voter. As $Margin_{\mathbf{P}}(a, b) \geq 0$ is equivalent to a being a winner in the two-candidate profile $\mathbf{P}_{ab}$, strong stability for winners can be restated as follows: if a is a winner in both $\mathbf{P}_{-b}$ and $\mathbf{P}_{ab}$, then a is a winner in their amalgamation $\mathbf{P}$. We will generalize this idea to apply not only to profiles of the form $\mathbf{P}_{-b}$ and $\mathbf{P}_{ab}$ but to any two "amalgamable" profiles.

Definition 5.6. Two profiles $\mathbf{P}$ and $\mathbf{P}'$ are *amalgamable* if $V(\mathbf{P}) = V(\mathbf{P}')$ and for every $x, y \in X(\mathbf{P}) \cap X(\mathbf{P}')$ and $i \in V(\mathbf{P})$, $x\mathbf{P}_i y$ if and only if $x\mathbf{P}'_i y$. An *amalgamation of* $\mathbf{P}$ *and* $\mathbf{P}'$ is any profile $\mathbf{Q} : V(\mathbf{P}) \to \mathcal{L}(X(\mathbf{P}) \cup X(\mathbf{P}'))$ such that for all $i \in V(\mathbf{P})$:

1. if $x, y \in X(\mathbf{P})$, then $x\mathbf{Q}_i y$ if and only if $x\mathbf{P}_i y$;
2. if $x, y \in X(\mathbf{P}')$, then $x\mathbf{Q}_i y$ if and only if $x\mathbf{P}'_i y$.

Example 5.7. Consider the following profiles $\mathbf{P}$ (top) and $\mathbf{P}'$ (bottom):

4	2	1	1	4
d	d	c	c	b
a	c	b	b	a
c	b	d	a	c
b	a	a	d	d

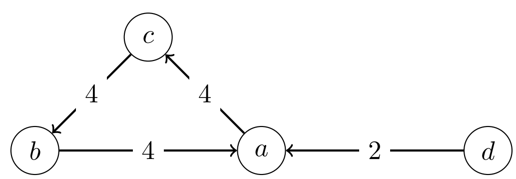

4	2	1	1	4
d	f	b	b	b
f	e	f	f	e
e	d	e	e	d
b	b	d	d	f

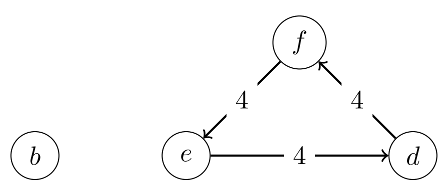

The only candidates common to both profiles are b and d, and every voter has the same ranking of b vs. d in both profiles. Hence the profiles are amalgamable. The following is an amalgamation $\mathbf{Q}$ of $\mathbf{P}$ and $\mathbf{P}'$:

1	3	2	1	1	3	1
d	d	f	c	c	b	b
f	a	e	b	b	e	a
a	f	d	f	a	a	e
c	c	c	e	f	c	c
e	e	b	d	e	d	d
b	b	a	a	d	f	f

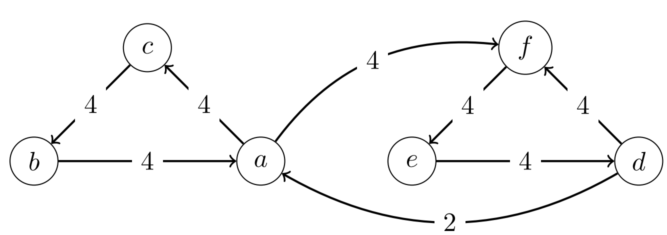

The amalgamation criterion states that if a is a winner in two amalgamable profiles, then a remains a winner in any amalgamation of the profiles ("if a can win the battles separately, then a can win the war").

Definition 5.8. A voting method F satisfies *amalgamation* if for any two amalgamable profiles $\mathbf{P}$ and $\mathbf{P}'$, if $a \in F(\mathbf{P}) \cap F(\mathbf{P}')$, then for any amalgamation $\mathbf{Q}$ of $\mathbf{P}$ and $\mathbf{P}'$, $a \in F(\mathbf{Q})$.

Amalgamation can be motivated as follows. First, when $a \in F(\mathbf{P}) \cap F(\mathbf{P}')$ and $\mathbf{Q}$ is an amalgamation of $\mathbf{P}$ and $\mathbf{P}'$, then by definition of amalgamation, the margins for each candidate x over a do not change. How then could a suddenly be defeated by a candidate x in the amalgamation $\mathbf{Q}$? By our first and third main ideas in Section 3, this would require the margin for x over a to exceed the threshold for defeat determined by incoherence; but incoherence can only *increase* by amalgamation, not decrease, so if x's margin over a was not sufficient to defeat a in $\mathbf{P}$ (resp. $\mathbf{P}'$), then it is not sufficient to defeat a in $\mathbf{Q}$.

Proposition 5.9. Split Cycle satisfies amalgamation.

Proof. The proof is similar to that of Proposition 5.5. If $a \in SC(\mathbf{P}) \cap SC(\mathbf{P}')$, then for all $c \in X(\mathbf{P})$, $Margin_{\mathbf{P}}(c,a) \leq Cycle\#_{\mathbf{P}}(c,a)$, and for all $c \in X(\mathbf{P}')$, $Margin_{\mathbf{P}'}(c,a) \leq Cycle\#_{\mathbf{P}'}(c,a)$. In addition, since $\mathbf{Q}$ is an amalgamation of $\mathbf{P}$ and $\mathbf{P}'$, for all $c \in X(\mathbf{P})$, $Margin_{\mathbf{Q}}(c,a) = Margin_{\mathbf{P}}(c,a)$, and for all $c \in X(\mathbf{P}')$, $Margin_{\mathbf{Q}}(c,a) = Margin_{\mathbf{P}'}(c,a)$. Also since $\mathbf{Q}$ is an amalgamation of $\mathbf{P}$ and $\mathbf{P}'$, for all $c \in X(\mathbf{P})$, $Cycle\#_{\mathbf{P}}(c,a) \leq Cycle\#_{\mathbf{Q}}(c,a)$, and for all $c \in X(\mathbf{P}')$, $Cycle\#_{\mathbf{P}'}(c,a) \leq Cycle\#_{\mathbf{Q}}(c,a)$. Putting the previous facts together, we have that for all $c \in X(\mathbf{P}) \cup X(\mathbf{P}') = X(\mathbf{Q})$, $Margin_{\mathbf{Q}}(c,a) \leq Cycle\#_{\mathbf{Q}}(c,a)$. Thus, $a \in SC(\mathbf{Q})$. □

Note that in Example 5.7, since $d \in SC(\mathbf{P}) \cap SC(\mathbf{P}')$, we have $d \in SC(\mathbf{Q})$ by amalgamation, in line with our third main idea in Section 3.

Remark 5.10. Instead of starting with $\mathbf{P}, \mathbf{P}'$ and considering an amalgamation $\mathbf{Q}$ of $\mathbf{P}, \mathbf{P}'$, we could start with a profile $\mathbf{Q}$ and consider *subprofiles* $\mathbf{P}, \mathbf{P}'$ of $\mathbf{Q}$. For $S \subseteq X(\mathbf{Q})$, let $\mathbf{Q}_{|S}$ be the profile with $X(\mathbf{Q}_{|S}) = S$ such that $\mathbf{Q}_{|S}(i)$ is the restriction of $\mathbf{Q}(i)$ to S. Then amalgamation is equivalent to: for any profile $\mathbf{Q}$ and $S, S' \in \wp(X(\mathbf{Q})) \setminus \{\varnothing\}$, $F(\mathbf{Q}_{|S}) \cap F(\mathbf{Q}_{|S'}) \subseteq F(\mathbf{Q}_{|S \cup S'})$. This condition is closely related to Sen's (1971, p. 314) choice-functional condition of *expansion* (γ). A *choice function* on a nonempty set X is a function $C : \wp(X) \setminus \{\varnothing\} \to \wp(X) \setminus \{\varnothing\}$ such that for all nonempty $S \subseteq X$, $\varnothing \neq C(S) \subseteq S$. Then C satisfies expansion if for all nonempty $S, S' \subseteq X$, $C(S) \cap C(S') \subseteq C(S \cup S')$. To relate this condition to social choice, Sen (1993) defines, for some fixed nonempty sets X and V, a *functional collective choice rule* (FCCR) for (X, V) to be a function f mapping each profile $\mathbf{P}$ with $X(\mathbf{P}) = X$ and $V(\mathbf{P}) = V$ to a choice function $f(\mathbf{P}, \cdot)$ on $X(\mathbf{P})$. Let a *variable-election FCCR* (VFCCR) be a function f mapping each profile $\mathbf{P}$ (i.e., allowing $X(\mathbf{P})$ and $V(\mathbf{P})$ to vary with $\mathbf{P}$) to a choice function $f(\mathbf{P}, \cdot)$ on $X(\mathbf{P})$. For voting theory, we take $X(\mathbf{P})$ to be the set of *candidates who appeared on the ballot* in the election scenario modeled by $\mathbf{P}$—since there is no practical point in considering voting procedures that have access to ranking information not on submitted ballots—but after the ballots are collected, some candidates may withdraw from consideration, be rejected by higher authorities, become incapacitated, etc., leaving us to choose winners from some "feasible set" $S \subsetneq X(\mathbf{P})$.

We say that a VFCCR f satisfies expansion if for all profiles $\mathbf{P}$, $f(\mathbf{P}, \cdot)$ satisfies expansion. We say that a VFCCR satisfies *amalgamation* if for any amalgamable profiles $\mathbf{P}, \mathbf{P}'$ and amalgamation $\mathbf{Q}$ of $\mathbf{P}, \mathbf{P}'$, $S \subseteq X(\mathbf{P})$, and $S' \subseteq X(\mathbf{P}')$, $f(\mathbf{P}, S) \cap f(\mathbf{P}', S') \subseteq f(\mathbf{Q}, S \cup S')$, which implies expansion, taking $\mathbf{P} = \mathbf{P}' = \mathbf{Q}$.

Thus, amalgamation constrains the relation between sets of winners for *different election scenarios* involving rankings of different (but possibly overlapping or extended) sets of candidates, modeled by different profiles; by contrast, expansion constrains the choices of winners from feasible sets based on a fixed profile of rankings. Expansion and amalgamation are equivalent for VFCCRs satisfying the independence condition that for all profiles $\mathbf{P}$ and nonempty $S \subseteq X(\mathbf{P})$, $f(\mathbf{P}, S) = f(\mathbf{P}_{|S}, S)$. However, expansion and amalgamation are not equivalent conditions on VFCCRs in general. For example, consider the *global Borda count* VFCCR (cf. Kelly 1988, p. 71): $f(\mathbf{P}, S)$ is the set of all $x \in S$ such that for all $y \in S$, the Borda score of x calculated with respect to the full profile $\mathbf{P}$ is at least that of y.[27] Global Borda count satisfies expansion but not amalgamation. By contrast, both local and global VFCCR versions of Split Cycle[28] satisfy amalgamation.

Remark 5.11. Amalgamation may remind one of the *reinforcement* criterion (see Pivato 2013b), but that criterion concerns combining profiles for disjoint sets of voters voting on the same set of candidates, whereas amalgamation concerns profiles for the same set of voters voting on different sets of candidates. Reinforcement states that for any profiles $\mathbf{P}$ and $\mathbf{P}'$ with $V(\mathbf{P}) \cap V(\mathbf{P}') = \varnothing$ and $X(\mathbf{P}) = X(\mathbf{P}')$, if $F(\mathbf{P}) \cap F(\mathbf{P}') \neq \varnothing$, then $F(\mathbf{P} + \mathbf{P}') = F(\mathbf{P}) \cap F(\mathbf{P}')$. No Condorcet consistent voting method satisfies reinforcement (see, e.g., Zwicker 2016, Proposition 2.5); we know of no non-trivial voting method that satisfies both amalgamation and reinforcement;[29] and we do not find reinforcement normatively compelling for all voting contexts.[30] Why when the voting method is given the full information in $\mathbf{P} + \mathbf{P}'$ should it be constrained by what it outputs when given only the limited information of $\mathbf{P}$ and only the limited information in $\mathbf{P}'$? For example, for three candidates a, b, c, suppose $\mathbf{P}$ is the classic Condorcet paradox profile with 6 voters such that $Margin_{\mathbf{P}}(a, b) = 2$, $Margin_{\mathbf{P}}(b, c) = 2$, and $Margin_{\mathbf{P}}(c, a) = 2$, so $F(\mathbf{P}) = \{a, b, c\}$ by fairness considerations (i.e., for any anonymous and neutral method F), while $\mathbf{P}'$ is a profile with 3 voters such that $Margin_{\mathbf{P}'}(b, a) = 1$, $Margin_{\mathbf{P}'}(a, c) = 1$, and $Margin_{\mathbf{P}'}(b, c) = 1$, so $F(\mathbf{P}') = \{b\}$ by Condorcet consistency. When we look at all the information in $\mathbf{P}+\mathbf{P}'$, we see that a is majority preferred to every other candidate—a is the Condorcet winner, so $F(\mathbf{P} + \mathbf{P}') = \{a\}$. Note that b is only majority preferred to a by a small margin in $\mathbf{P}'$, whereas a is majority preferred to b by a larger margin in $\mathbf{P}$. Due to c's poor performance in $\mathbf{P}'$, there is no cycle in the full profile $\mathbf{P} + \mathbf{P}'$, so although a's margin over b failed to make a the winner in $\mathbf{P}$ due to fairness considerations, a's margin over b in the full cycle-free profile makes a the winner, as it should. Note this also shows it is a mistake to think that given a profile like $\mathbf{P} + \mathbf{P}'$, one should assume that the "Condorcet component" $\mathbf{P}$ can be deleted from the profile, as if it contains no information, while not changing the winning set (cf. the property of *cancelling properly* in Balinski and Laraki 2010, p. 77). The Condorcet component $\mathbf{P}$ contains the important information that $Margin_{\mathbf{P}}(a, b) = 2$.

[27] The *local Borda count* VFCCR (cf. Kelly 1988, p. 74) takes $f(\mathbf{P}, S)$ to be the set of all $x \in S$ such that for all $y \in S$, the Borda score of x calculated with respect to the restricted profile $\mathbf{P}_{|S}$ is at least that of y.

[28] Global Split Cycle takes $f(\mathbf{P}, S)$ to be the set of all $x \in S$ such that for all $y \in S$, y does not defeat x according to the Split Cycle defeat relation calculated with respect to the full profile $\mathbf{P}$, whereas local Split Cycle uses the defeat relation calculated with respect to $\mathbf{P}_{|S}$.

[29] Amalgamation implies a weakening of Condorcet consistency—it implies that if there is a Condorcet winner, that candidate must be *among* the winners. But with no other axioms, amalgamation and reinforcement are consistent, as they are both satisfied by the trivial voting methods that always picks all candidates as winners.

[30] In the context where there is a "true" ranking of the candidates of which voters have noisy perceptions, reinforcement is satisfied by any voting method that can be rationalized as the maximum likelihood estimator for some noise model with i.i.d. votes (see Conitzer and Sandholm 2005; Pivato 2013a).

5.3 Rejectability

The next criterion we propose concerns winnowing a set of winners down to a single winner. The *rejectability* criterion states that if in a profile $\mathbf{P}$, candidate x is *among* the winners, then we should be able to make x the *unique* winner in a profile $\mathbf{P}^+$ obtained from $\mathbf{P}$ by adding voters who sufficiently strengthen the rejection of other candidates, i.e., sufficiently increase what were already positive margins against other candidates, so as to defeat the others (recall our idea in Section 3 that incoherence does not raise the threshold for defeat *infinitely*). Thus, if candidate a is majority preferred to b in $\mathbf{P}$, then this still holds in $\mathbf{P}^+$ with a margin that is at least as large and possibly larger than in $\mathbf{P}$. No majority preferences are reversed from $\mathbf{P}$ to $\mathbf{P}^+$, for if we were to allow that, then we could simply make x the Condorcet winner in $\mathbf{P}^+$, trivializing the criterion.

Definition 5.12. A voting method F satisfies *rejectability* if for any profile $\mathbf{P}$ such that $|F(\mathbf{P})| > 1$ and $x \in F(\mathbf{P})$, there is a profile $\mathbf{P}^+$ with $X(\mathbf{P}) = X(\mathbf{P}^+)$ and $V(\mathbf{P}) \subseteq V(\mathbf{P}^+)$ such that for all $a, b \in X(\mathbf{P})$, if $Margin_{\mathbf{P}}(a,b) \geq 0$, then $Margin_{\mathbf{P}^+}(a,b) \geq Margin_{\mathbf{P}}(a,b)$, and $F(\mathbf{P}^+) = \{x\}$.

Thus, if a method *fails* rejectability, then for some $\mathbf{P}$ and $x \in F(\mathbf{P})$, no matter how extremely we turn majority preferences against other candidates into enormous landslides, we cannot make x the unique winner.

Rejectability is a strong criterion insofar as it rules out all irresolute C1 voting methods (as does the resolvability criterion of Section 6.4). Recall that a voting method F is C1 (Fishburn 1977) if for any profiles $\mathbf{P}$ and $\mathbf{P}'$, if their majority graphs (Definition 2.4) are the same—$M(\mathbf{P}) = M(\mathbf{P}')$—then their winners are also the same—$F(\mathbf{P}) = F(\mathbf{P}')$. Copeland, GETCHA/GOCHA, and Uncovered Set are all C1.

Proposition 5.13. No anonymous and neutral C1 voting method satisfies rejectability.

Proof. Given a profile $\mathbf{P}$ with $X(\mathbf{P}) = \{a, b, c\}$ and whose margin graph contains the cycle $a \to b \to c \to a$, no matter the margins, an anonymous and neutral C1 method F must have $F(\mathbf{P}) = \{a, b, c\}$; hence we can never increase any margins in such a way that one candidate becomes the unique winner. □

In our proof that Split Cycle satisfies rejectability, we use the following lemma.

Lemma 5.14. Split Cycle satisfies the *overwhelming majority*[31] criterion: for all profiles $\mathbf{P}$ and $\mathbf{P}'$ with $X(\mathbf{P}) = X(\mathbf{P}')$ and $V(\mathbf{P}) \cap V(\mathbf{P}') = \varnothing$, there is an $n \in \mathbb{N}$ such that for all $m \in \mathbb{N}$ with $m \geq n$, we have $SC(\mathbf{P} + m\mathbf{P}') \subseteq SC(\mathbf{P}')$, where $m\mathbf{P}' = \gamma_1(\mathbf{P}') + \cdots + \gamma_m(\mathbf{P}')$ with $\gamma_1(\mathbf{P}'), \ldots, \gamma_m(\mathbf{P}')$ being copies of $\mathbf{P}'$ with pairwise disjoint sets of voters (recall Definition 2.7).

Proof. Let $n = 2|V(\mathbf{P})|$. To show that $SC(\mathbf{P} + m\mathbf{P}') \subseteq SC(\mathbf{P}')$, it suffices to show that for any $a, b \in X(\mathbf{P}')$, if a defeats b in $\mathbf{P}'$, then a defeats b in $\mathbf{P} + m\mathbf{P}'$. Assume a defeats b in $\mathbf{P}'$, so $Margin_{\mathbf{P}'}(a,b) > Cycle\#_{\mathbf{P}'}(a,b)$, which implies $Margin_{\mathbf{P}'}(a,b) - Cycle\#_{\mathbf{P}'}(a,b) > 1$ as all margins have the same parity. Then for all $m \geq n$, since

$$Margin_{m\mathbf{P}'}(a,b) = m \times Margin_{\mathbf{P}'}(a,b) \text{ and } Cycle\#_{m\mathbf{P}'}(a,b) = m \times Cycle\#_{\mathbf{P}'}(a,b),$$

we have $Margin_{m\mathbf{P}'}(a,b) - Cycle\#_{m\mathbf{P}'}(a,b) > 2|V(\mathbf{P})|$ by our choice of n. Also note that

$$Margin_{\mathbf{P}+m\mathbf{P}'}(a,b) \geq Margin_{m\mathbf{P}'}(a,b) - |V(\mathbf{P})| \text{ and } Cycle\#_{\mathbf{P}+m\mathbf{P}'}(a,b) \leq Cycle\#_{m\mathbf{P}'}(a,b) + |V(\mathbf{P})|.$$

It follows that $Margin_{\mathbf{P}+m\mathbf{P}'}(a,b) > Cycle\#_{\mathbf{P}+m\mathbf{P}'}(a,b)$, so a defeats b in $\mathbf{P} + m\mathbf{P}'$. □

[31] This is the terminology from Myerson 1995. Cf. Smith's (1973) "Archimedean property" and Young's (1975) "continuity."

Proposition 5.15. Split Cycle satisfies rejectability.

Proof. We claim that to establish rejectability, it suffices to show that for any profile $\mathbf{P}$ such that $|SC(\mathbf{P})| > 1$ and $x \in SC(\mathbf{P})$, there is a profile $\mathbf{P}'$ with $X(\mathbf{P}) = X(\mathbf{P}')$ such that for all $a, b \in X(\mathbf{P})$, if $Margin_{\mathbf{P}}(a,b) \geq 0$, then $Margin_{\mathbf{P}'}(a,b) \geq Margin_{\mathbf{P}}(a,b)$, and $SC(\mathbf{P}') = \{x\}$. For then by Lemma 5.14, there is an $m \in \mathbb{N}$ such that $SC(\mathbf{P} + m\mathbf{P}') = \{x\}$, and for all $a, b \in X(\mathbf{P})$, if $Margin_{\mathbf{P}}(a,b) \geq 0$, then $Margin_{\mathbf{P}+m\mathbf{P}'}(a,b) \geq Margin_{\mathbf{P}}(a,b)$. As $V(\mathbf{P}) \subseteq V(\mathbf{P} + m\mathbf{P}')$, we may take $\mathbf{P}^+ = \mathbf{P} + m\mathbf{P}'$ for rejectability.

Suppose $|SC(\mathbf{P})| > 1$ and $x \in SC(\mathbf{P})$. We show how to modify the margin graph $\mathcal{M}(\mathbf{P})$ to a margin graph $\mathcal{M}'$ on $X(\mathbf{P})$ such that (i) all edges between nodes are preserved from $\mathcal{M}(\mathbf{P})$ to $\mathcal{M}'$, (ii) no weights on edges decrease from $\mathcal{M}(\mathbf{P})$ to $\mathcal{M}'$, and (iii) $SC(\mathcal{M}') = \{x\}$ (recall Remark 3.6). Then Debord's Theorem yields a profile $\mathbf{P}'$ whose margin graph is $\mathcal{M}'$. By (i)-(ii), we have that for all $a, b \in X(\mathbf{P})$, $Margin_{\mathbf{P}'}(a,b) \geq Margin_{\mathbf{P}}(a,b)$. By (iii), $SC(\mathbf{P}') = \{x\}$.

Let the set of edges in $\mathcal{M}'$ be the set of all edges in $\mathcal{M}(\mathbf{P})$ plus an edge from x to any y such that $Margin_{\mathbf{P}}(x,y) = 0$. Let n be the largest margin in $\mathcal{M}(\mathbf{P})$. Each edge (a,b) in $\mathcal{M}'$ has weight either $n+1$ or $n+3$ according to the following rules (we use $Margin_{\mathcal{M}'}$ and $Cycle\#_{\mathcal{M}'}$ with their obvious meanings):

1. if the edge (a,b) occurs on a shortest simple path[32] from x to b in $\mathcal{M}'$, set $Margin_{\mathcal{M}'}(a,b) = n+3$;
2. otherwise, set $Margin_{\mathcal{M}'}(a,b) = n+1$.

We now claim that every $y \in X(\mathbf{P}) \setminus \{x\}$ is defeated in $\mathcal{M}'$.

Case 1: $Margin_{\mathbf{P}}(x,y) \geq 0$. Then $Margin_{\mathcal{M}'}(x,y) = n+3$ by rule 1. Moreover, for any simple cycle ρ of the form $x \to y \to z_1 \to \dots \to z_k \to x$ in $\mathcal{M}'$, we have $Margin_{\mathbf{P}}(z_k, x) > 0$ by the construction of $\mathcal{M}'$ from $\mathcal{M}(\mathbf{P})$ and hence $Margin_{\mathcal{M}'}(z_k, x) = n+1$ by rule 2, so $Split\#(\rho) = n+1$. Hence $Cycle\#_{\mathcal{M}'}(x,y) = n+1$. Thus, $Margin_{\mathcal{M}'}(x,y) > Cycle\#_{\mathcal{M}'}(x,y)$, so x defeats y in $\mathcal{M}'$.

Case 2: $Margin_{\mathbf{P}}(y,x) > 0$. Then since $x \in SC(\mathbf{P})$, it follows by Lemma 3.12 that there is a simple cycle of the form $y \to x \to z_1 \to \dots \to z_k \to y$ in $\mathbf{P}$ where $x \to z_1 \to \dots \to z_k \to y$ is a shortest simple path from x to y. Hence $Margin_{\mathcal{M}'}(z_k, y) = n+3$ by rule 1. We claim that z_k defeats y in $\mathcal{M}'$. If there is no simple cycle of the form $z_k \to w_1 \to \dots \to w_\ell$ with $w_1 = y$ and $w_\ell = z_k$ in $\mathcal{M}'$, then z_k defeats y in $\mathcal{M}'$. If there is such a simple cycle ρ, then we claim that one of the edges $w_i \to w_{i+1}$ in ρ has weight $n+1$. If there is no simple path from x to any of $w_2, \dots, w_\ell$, this follows from rule 2 above. So suppose there is a simple path from x to one of $w_2, \dots, w_\ell$. Then there is a w_i such that (i) the shortest path p from x to w_i is no longer than the shortest path from x to any w_j. This setup is shown in Figure 4. Now we claim that the edge $w_{i-1} \to w_i$ in ρ has weight $n+1$; for it to have weight $n+3$, the edge $w_{i-1} \to w_i$ must occur on a shortest path from x to w_i, which is impossible. For suppose p' is a path from x to w_i including the edge $w_{i-1} \to w_i$. By (i), the initial segment of p' from x to w_{i-1} has length at least that of p, by our choice of p; so the length of p' is at least the length of p plus one; hence p' is not a shortest path from x to w_i. Thus, we have proved that one of the edges $w_i \to w_{i+1}$ has weight $n+1$. Thus, $Split\#(\rho) = n+1$. It follows that $Cycle\#_{\mathcal{M}'}(z_k, y) = n+1$, which with $Margin_{\mathcal{M}'}(z_k, y) = n+3$ implies that z_k defeats y in $\mathcal{M}'$. ☐

Corollary 5.16. Beat Path and Ranked Pairs satisfy rejectability.

[32] A *simple path* in a graph is a sequence $\langle x_1, \dots, x_n \rangle$ of *distinct* nodes with $x_i \to x_{i+1}$ for each $i \in \{1, \dots, n-1\}$. The *length* of a path is the number of nodes in the path minus one.

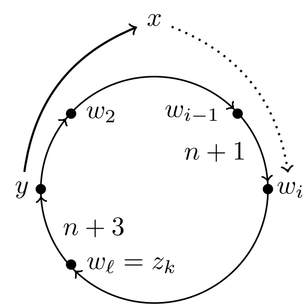

Figure 4: diagram for the proof of Proposition 5.15.

Proof. Let $F \in \{BP, RP\}$ and $\mathbf{P}$ be a profile such that $|F(\mathbf{P})| > 1$ and $x \in F(\mathbf{P})$. Then by Lemmas 3.17 and 3.18, $|SC(\mathbf{P})| > 1$ and $x \in SC(\mathbf{P})$. Hence by Proposition 5.15, there is a $\mathbf{P}'$ as in the definition of rejectability such that $SC(\mathbf{P}') = \{x\}$, which implies $F(\mathbf{P}') = \{x\}$ given Lemmas 3.17 and 3.18 and $F(\mathbf{P}') \neq \varnothing$. ☐

Example 5.17. If we pick any candidate x in the majority graph shown on the left below, the proof of Proposition 5.15 give us an algorithm to weight the edges of the majority graph such that in the resulting margin graph x is the unique Split Cycle winner. For example, we can make a the unique winner with the weighting on the middle graph and d the unique winner with the weighting on the right graph.

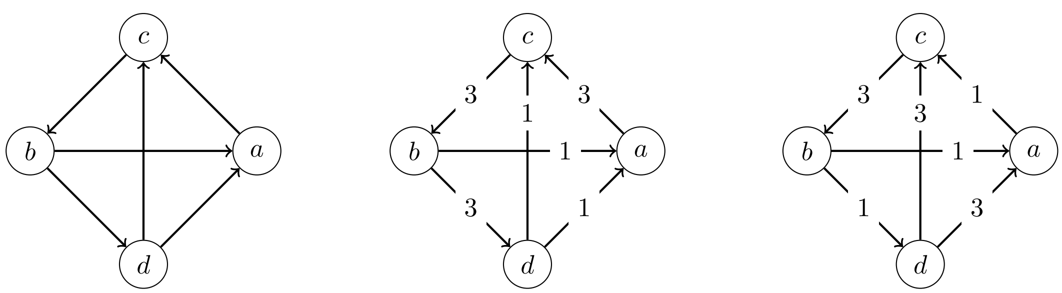

In fact, from the proof of Proposition 5.15 we can extract a proof of the following proposition about when it is possible, starting from an arbitrary graph, to turn the graph into a margin graph in which a given candidate is a (unique) winner for Split Cycle.

Proposition 5.18. For any asymmetric directed graph $\mathcal{G} = (G, \rightarrow)$ and $a \in G$, the following are equivalent:

1. there is a margin graph $\mathcal{M}$ based on $\mathcal{G}$ such that $a \in SC(\mathcal{M})$ (recall Remark 3.6);
2. there is a margin graph $\mathcal{M}$ based on $\mathcal{G}$ such that $SC(\mathcal{M}) = \{a\}$;
3. for every $x \in G \setminus \{a\}$, if $x \rightarrow a$, then there is a simple cycle of the form $x \rightarrow a \rightarrow y_1 \rightarrow \cdots \rightarrow y_n \rightarrow x$ in $\mathcal{G}$.

6 Other Criteria

In this section, we consider how Split Cycle fairs with respect to several other standard criteria for voting methods from the literature.

6.1 Smith and Schwartz Criteria

The first standard criterion we consider is the *Smith criterion* (Smith 1973), according to which the set of winners must be a subset of the *Smith set*—the smallest set of candidates such that every candidate inside the set is majority preferred to every candidate outside the set. Following the terminology of Schwartz (1986), we also call the Smith set the GETCHA set ("GETCHA" stands for "generalized top-choice assumption").

Definition 6.1. Let $\mathbf{P}$ be a profile and $S \subseteq X(\mathbf{P})$. Then S is $\rightarrow_{\mathbf{P}}$*-dominant* if $S \neq \varnothing$ and for all $x \in S$ and $y \in X(\mathbf{P}) \setminus S$, we have $x \rightarrow_{\mathbf{P}} y$. Define

$$GETCHA(\mathbf{P}) = \bigcap \{S \subseteq X(\mathbf{P}) \mid S \text{ is } \rightarrow_{\mathbf{P}}\text{-dominant}\}.$$

Definition 6.2. A voting method F satisfies the *Smith criterion* if for any profile $\mathbf{P}$, $F(\mathbf{P}) \subseteq GETCHA(\mathbf{P})$.

Proposition 6.3. Split Cycle satisfies the Smith criterion.

Proof. Suppose $b \in SC(\mathbf{P}) \setminus GETCHA(\mathbf{P})$. Since $b \notin GETCHA(\mathbf{P})$, there is an $a \in GETCHA(\mathbf{P})$ such that $a \rightarrow_{\mathbf{P}} b$. Then since $b \in SC(\mathbf{P})$, it follows by Lemma 3.12 that there is a simple cycle ρ of the form $a \rightarrow b \rightarrow x_1 \rightarrow \cdots \rightarrow x_n \rightarrow a$. Hence one of the edges in ρ goes from a candidate outside $GETCHA(\mathbf{P})$ to a candidate inside $GETCHA(\mathbf{P})$, which is a contradiction. □

Next we consider a strengthening of the Smith criterion, based on the idea of the Schwartz set, or in Schwartz's (1986) terminology, the GOCHA set ("GOCHA" stands for "generalized optimal choice axiom").

Definition 6.4. Let $\mathbf{P}$ be a profile and $S \subseteq X(\mathbf{P})$. Then S is $\rightarrow_{\mathbf{P}}$*-undominated* if for all $x \in S$ and $y \in X(\mathbf{P}) \setminus S$, we have $y \not\rightarrow_{\mathbf{P}} x$. Define

$$GOCHA(\mathbf{P}) = \bigcup \{S \subseteq X(\mathbf{P}) \mid S \text{ is } \rightarrow_{\mathbf{P}}\text{-undominated and no } S' \subsetneq S \text{ is } \rightarrow_{\mathbf{P}}\text{-undominated}\}.$$

Another useful characterization of the GOCHA set is given by the following lemma.

Lemma 6.5 (Schwartz 1986, Corollary 6.2.2)**.** Let $\mathbf{P}$ be any profile, and let $\rightarrow^*_{\mathbf{P}}$ be the transitive closure of $\rightarrow_{\mathbf{P}}$, i.e., $a \rightarrow^*_{\mathbf{P}} b$ if and only if there are $x_1, \dots, x_n \in X(\mathbf{P})$ with $a = x_1$ and $b = x_n$ such that $x_1 \rightarrow_{\mathbf{P}} \cdots \rightarrow_{\mathbf{P}} x_n$. Then

$$GOCHA(\mathbf{P}) = \{x \in X(\mathbf{P}) \mid \text{there is no } y \in X(\mathbf{P}) : y \rightarrow^*_{\mathbf{P}} x \text{ and } x \not\rightarrow^*_{\mathbf{P}} y\}.$$

Just as the Smith criterion states that the set of winners should always be a subset of the Smith set, the Schwartz criterion states that the set of winners should always be a subset of the Schwartz set.

Definition 6.6. A voting method F satisfies the *Schwartz criterion* if for any profile $\mathbf{P}$, $F(\mathbf{P}) \subseteq GOCHA(\mathbf{P})$.

In contrast to Proposition 6.3, Split Cycle does not satisfy the Schwartz criterion. After the proof, we will explain why we do not find the Schwartz criterion normatively plausible.

Proposition 6.7. Split Cycle does not satisfy the Schwartz criterion.

Proof. By Debord's Theorem, there is a profile $\mathbf{P}$ with the following margin graph (simplifying our example for the third idea of Section 3.1):

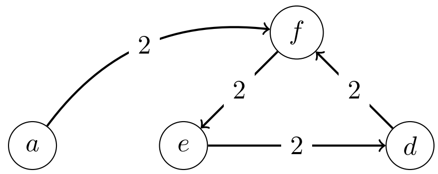

First, note that $d \in SC(\mathbf{P})$ (indeed, $SC(\mathbf{P}) = \{a, d, e\}$), because the only candidate with a positive margin over d is e, but $Margin_{\mathbf{P}}(e, d) \not> Cycle\#_{\mathbf{P}}(e, d)$. Yet $d \notin GOCHA(\mathbf{P})$, because $a \to^*_{\mathbf{P}} d$ and $d \not\to^*_{\mathbf{P}} a$. $\square$

For the reasons explained in Section 3.1 for the idea that *defeat is direct*, we think that d should not be kicked out of the winning set by a in the profile $\mathbf{P}$ used in the proof of Proposition 6.7. Thus, we do not accept the Schwartz criterion. The profile used in the proof of Proposition 6.7 also shows the following.

Proposition 6.8. There is no voting method satisfying anonymity, neutrality, strong stability for winners, and the Schwartz criterion.

Proof. Where $\mathbf{P}$ is the profile used in the proof of Proposition 6.7, by anonymity and neutrality, $F(\mathbf{P}_{-a}) = \{d, e, f\}$. Then since $Margin_{\mathbf{P}}(a, d) = 0$, it follows by strong stability for winners that $d \in F(\mathbf{P})$, which contradicts the Schwartz criterion as in the proof of Proposition 6.7. $\square$

6.2 Independence of Smith-Dominated Alternatives

The Smith criterion of the previous section can be strengthened to the criterion that deletion of candidates outside the Smith set should not change the set of winners.

Definition 6.9. A voting method F satisfies *independence of Smith-dominated alternatives* (ISDA) if for any profile $\mathbf{P}$ and $x \in X(\mathbf{P}) \setminus GETCHA(\mathbf{P})$, we have $F(\mathbf{P}) = F(\mathbf{P}_{-x})$.

Remark 6.10. ISDA implies the Smith criterion, since if $x \in F(\mathbf{P}) \setminus GETCHA(\mathbf{P})$, then $F(\mathbf{P}) \neq F(\mathbf{P}_{-x})$.

Remark 6.11. Given a profile $\mathbf{P}$ and $S \subseteq X(\mathbf{P})$, define the restricted profile $\mathbf{P}_{|S} : V(\mathbf{P}) \to \mathcal{L}(S)$ such that $\mathbf{P}_{|S}(i)$ is the restriction of the linear order $\mathbf{P}(i)$ to S. Then observe that if F satisfies ISDA, then F satisfies the following inter-profile condition: for any profiles $\mathbf{P}$ and $\mathbf{P}'$, where $S = GETCHA(\mathbf{P})$ and $S' = GETCHA(\mathbf{P}')$, if $\mathbf{P}_{|S} = \mathbf{P}'_{|S'}$, then $F(\mathbf{P}) = F(\mathbf{P}')$. This inter-profile condition may be viewed as a weakening of the independence of irrelevant alternatives (IIA).

Proposition 6.12. Split Cycle satisfies ISDA.

Proof. Suppose $x \in X(\mathbf{P}) \setminus GETCHA(\mathbf{P})$. It follows that $GETCHA(\mathbf{P}) = GETCHA(\mathbf{P}_{-x})$.[33] Toward showing that $SC(\mathbf{P}) = SC(\mathbf{P}_{-x})$, let $\mathbf{Q}, \mathbf{Q}' \in \{\mathbf{P}, \mathbf{P}_{-x}\}$ with $\mathbf{Q} \neq \mathbf{Q}'$. Suppose $y \in SC(\mathbf{Q})$. We will show $y \in SC(\mathbf{Q}')$. For any $z \in X(\mathbf{P}_{-x})$, we have

$$Margin_{\mathbf{P}}(z, y) = Margin_{\mathbf{P}_{-x}}(z, y). \tag{2}$$

[33] Clearly $GETCHA(\mathbf{P})$ is $\to_{\mathbf{P}_{-x}}$-dominant, so $GETCHA(\mathbf{P}_{-x}) \subseteq GETCHA(\mathbf{P})$ by Definition 6.1. To see that $GETCHA(\mathbf{P}_{-x})$ is $\to_{\mathbf{P}}$-dominant, consider an $a \in GETCHA(\mathbf{P}_{-x})$ and $b \in X(\mathbf{P}) \setminus GETCHA(\mathbf{P}_{-x})$. If $b \neq x$, then $b \in X(\mathbf{P}_{-x}) \setminus GETCHA(\mathbf{P}_{-x})$ and hence $a \to_{\mathbf{P}_{-x}} b$ because $GETCHA(\mathbf{P}_{-x})$ is $\to_{\mathbf{P}_{-x}}$-dominant, which implies $a \to_{\mathbf{P}} b$. If $b = x$, then since $GETCHA(\mathbf{P}_{-x}) \subseteq GETCHA(\mathbf{P})$ and $x \in X(\mathbf{P}) \setminus GETCHA(\mathbf{P})$, again we have $a \to_{\mathbf{P}} b$. Thus, $GETCHA(\mathbf{P}_{-x})$ is $\to_{\mathbf{P}}$-dominant, so $GETCHA(\mathbf{P}) \subseteq GETCHA(\mathbf{P}_{-x})$ by Definition 6.1.

Hence if $Margin_{\mathbf{Q}}(z,y)=0$, then $Margin_{\mathbf{Q}'}(z,y)=0$, so z does not defeat y in $\mathbf{Q}'$. Suppose instead that $Margin_{\mathbf{Q}}(z,y)>0$. Since $y\in SC(\mathbf{Q})$, z does not defeat y in $\mathbf{Q}$, so we have

$$Margin_{\mathbf{Q}}(z,y) < Cycle\#_{\mathbf{Q}}(z,y). \tag{3}$$

Now we claim that

$$Cycle\#_{\mathbf{P}}(z,y) = Cycle\#_{\mathbf{P}_{-x}}(z,y). \tag{4}$$

Since $y\in SC(\mathbf{Q})$, $y\in GETCHA(\mathbf{Q})$ by Proposition 6.3. Then from $Margin_{\mathbf{Q}}(z,y)>0$, it follows that $z\in GETCHA(\mathbf{Q})$. Since $x\in X(\mathbf{P})\setminus GETCHA(\mathbf{P})$ and $z\in GETCHA(\mathbf{Q}) = GETCHA(\mathbf{P})$, there is no simple cycle in $\mathcal{M}(\mathbf{P})$ of the form $z\to y\to w_1\to\cdots\to w_n\to z$ with $x\in\{w_1,\dots,w_n\}$, since there is no path from a candidate outside $GETCHA(\mathbf{P})$, like x, to a candidate inside $GETCHA(\mathbf{P})$, like z. This establishes (4). Then together (2), (3), and (4) entail $Margin_{\mathbf{Q}'}(z,y) < Cycle\#_{\mathbf{Q}'}(z,y)$. Hence z does not defeat y in $\mathbf{Q}'$. Finally, if $\mathbf{Q}'=\mathbf{P}$, then x does not defeat y in $\mathbf{Q}'$, since $y\in GETCHA(\mathbf{Q}')$ while $x\notin GETCHA(\mathbf{Q}')$. Thus, no candidate defeats y in $\mathbf{Q}'$, so $y\in SC(\mathbf{Q}')$. ☐

6.3 Reversal Symmetry

Next we consider a criterion due to Saari (Saari 1994, § 3.1.3; Saari 1999, § 7.1) that can be seen as an extension of the neutrality criterion. Recall that neutrality states that if we swap the places of two candidates a and b on every voter's ballot, then if a won the election before the swap, b should win the election after the swap. Reversal symmetry extends this idea from pairwise swaps to full reversals of voters' ballots.

Definition 6.13. A voting method F satisfies *reversal symmetry* if for any profile $\mathbf{P}$ with $|X(\mathbf{P})|>1$, if $F(\mathbf{P})=\{x\}$, then $x\notin F(\mathbf{P}^r)$, where $\mathbf{P}^r$ is the profile obtained by reversing all the ballots in $\mathbf{P}$.

Proposition 6.14. Split Cycle satisfies reversal symmetry.

Proof. Suppose $\mathbf{P}$ is such that $|X(\mathbf{P})|>1$ and $SC(\mathbf{P})=\{x\}$. It follows by Lemma 3.11 that there is a $y\in X(\mathbf{P})$ such that x defeats y in $\mathbf{P}$, i.e., such that $Margin_{\mathbf{P}}(x,y) > Cycle\#_{\mathbf{P}}(x,y)$. But then since $Margin_{\mathbf{P}}(x,y) = Margin_{\mathbf{P}^r}(y,x)$ and $Cycle\#_{\mathbf{P}}(x,y) = Cycle\#_{\mathbf{P}^r}(y,x)$, we have $Margin_{\mathbf{P}^r}(y,x) > Cycle\#_{\mathbf{P}^r}(x,y)$, so $x\notin SC(\mathbf{P}^r)$. ☐

6.4 Resolvability

Like the rejectability criterion of Section 5.3, the criteria considered in this section concern winnowing a set of winners down to a unique winner. The first criterion, *single-voter resolvability*, says that this can always be done by the addition of *just one new voter*. We see no justification for requiring that one voter is always sufficient, and no arguments for the normative necessity of this criterion are given in the literature. Tideman (1987) uses single-voter resolvability to rule out the GOCHA method, but this can be accomplished by rejectability or narrowing instead. Indeed, we suspect that some intuitions about winnowing sets of winners down to a unique winner are better captured by rejectability than by single-voter resolvability.

Definition 6.15. A voting method F satisfies *single-voter resolvability* if for any profile $\mathbf{P}$, if $|F(\mathbf{P})|>1$, then for any $x\in F(\mathbf{P})$, there is a profile $\mathbf{P}'$ with $V(\mathbf{P})\cap V(\mathbf{P}')=\varnothing$ and $|V(\mathbf{P}')|=1$ such that $F(\mathbf{P}+\mathbf{P}')=\{x\}$.

Proposition 6.16. Split Cycle does not satisfy single-voter resolvability.

Proof. Recall the margin graph of the profile $\mathbf{P}$ from the proof of Proposition 5.4 showing that Minimax and Beat Path do not satisfy stability for winners:

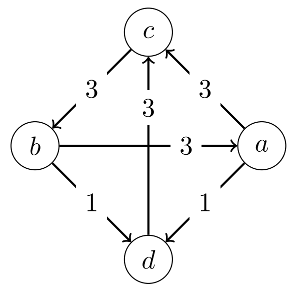

Here $SC(\mathbf{P}) = \{a, b, d\}$, but there is no one-voter profile $\mathbf{P}'$ with $SC(\mathbf{P}+\mathbf{P}') = \{a\}$ or $SC(\mathbf{P} + \mathbf{P}') = \{b\}$. □

Below we will show a deep tension between single-voter resolvability and stability for winners.

Resolvability and rejectability can be related using the following additional criterion from Smith 1973.

Definition 6.17. A voting method F satisfies *homogeneity* if for any profile $\mathbf{P}$, if $\mathbf{P}^*$ is a copy of $\mathbf{P}$ with a disjoint set of voters (recall Definition 2.7), then $F(\mathbf{P}) = F(\mathbf{P} + \mathbf{P}^*)$.

Lemma 6.18. If a voting method satisfies homogeneity and single-voter resolvability, then it satisfies rejectability.

Proof. Let $\mathbf{P}$ be such that $|F(\mathbf{P})| > 1$ and $x \in F(\mathbf{P})$. Let $\mathbf{P}^*$ be a copy of $\mathbf{P}$ with a disjoint set of voters. Then by homogeneity, $F(\mathbf{P}) = F(\mathbf{P} + \mathbf{P}^*)$. It follows by resolvability that there is a single voter profile $\mathbf{Q}$ such that $F(\mathbf{P} + \mathbf{P}^* + \mathbf{Q}) = \{x\}$. Since for any $a, b \in X(\mathbf{P})$ with $Margin_{\mathbf{P}}(a, b) \geq 0$, we have $Margin_{\mathbf{P}+\mathbf{P}^*+\mathbf{Q}}(a, b) \geq Margin_{\mathbf{P}}(a, b)$, the profile $\mathbf{P} + \mathbf{P}^* + \mathbf{Q}$ is the desired profile $\mathbf{P}^+$ for rejectability. □

Another use of the term 'resolvability' discussed in the literature (see Schulze 2011, § 4.2.1) concerns the proportion of profiles with multiple winners as the number of voters goes to infinity.

Definition 6.19. For $k \in \mathbb{N}$, a voting method F satisfies *asymptotic resolvability for k candidates* if the proportion of profiles $\mathbf{P}$ with $|X(\mathbf{P})| = k$ and $|V(\mathbf{P})| = n$ for which $|F(\mathbf{P})| > 1$ approaches 0 as n approaches infinity.

In Section 3.3, we discussed the tradeoff between a voting method being more resolute and satisfying desirable variable-candidate properties, such as stability for winners. The next result illustrates this tradeoff in the case of resolvability. We impose an assumption that is satisfied by all voting methods based on majority margins that we know of—say that a voting method F satisfies the *triangle property* if for any uniquely weighted profile $\mathbf{P}$ (recall Section 4.4) whose margin graph is of the form shown below, where i is the smallest margin and k is the largest margin (i.e., $i < j < k$), we have $x \in F(\mathbf{P})$:

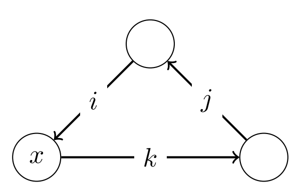

The proof of Theorem 6.20 makes essential use of a theorem of Harrison-Trainor (2020) that answers one of our conjectures.

Theorem 6.20. Suppose F is a voting method satisfying stability for winners and the triangle property. Then:

1. F does not satisfy single-voter resolvability;
2. F does not satisfy asymptotic resolvability for any $k > 3$.

Proof. We will use the fact that stability for winners implies the following: for any profile $\mathbf{P}$, defining $\mathbf{P}^G = \mathbf{P}_{|GETCHA(\mathbf{P})}$ (recall Section 6.1), we have $F(\mathbf{P}^G) \subseteq F(\mathbf{P})$. To see this, let $X(\mathbf{P}) \setminus GETCHA(\mathbf{P}) = \{b_1, \ldots, b_n\}$. Suppose $a \in F(\mathbf{P}^G)$, so $a \in GETCHA(\mathbf{P})$. Then $a \to b_1$, so stability for winners implies $a \in F(\mathbf{P}_{|GETCHA(\mathbf{P}) \cup \{b_1\}})$. Then since $a \to b_2$, stability for winners implies $a \in F(\mathbf{P}_{|GETCHA(\mathbf{P}) \cup \{b_1, b_2\}})$, and so on, until we obtain $a \in F(\mathbf{P})$.

We will also use the notion of a *qualitative margin graph*, which is a pair $\mathbb{M} = (M, \prec)$ where M is an asymmetric directed graph and $\prec$ is a strict weak order on the set of edges of M. We say that $\mathbb{M}$ is *uniquely weighted* if $\prec$ is a strict linear order. Given a profile $\mathbf{P}$, let the *qualitative margin graph* $\mathbb{M}(\mathbf{P})$ *of* $\mathbf{P}$ be the pair $(M(\mathbf{P}), \prec_{\mathbf{P}})$ where $M(\mathbf{P})$ is the majority graph of $\mathbf{P}$, and $\prec_{\mathbf{P}}$ is the relation on the set of edges of $M(\mathbf{P})$ defined by $(a, b) \prec_{\mathbf{P}} (c, d)$ if $Margin_{\mathbf{P}}(a, b) < Margin_{\mathbf{P}}(c, d)$. It follows from Debord's Theorem that every qualitative margin graph is realized by some profile. Harrison-Trainor (2020) proves that for any $k \geq 1$ and uniquely weighted qualitative margin graph $\mathbb{M}$ with k candidates, the proportion of profiles with k candidates and n voters realizing $\mathbb{M}$ does not go to 0 as n goes to infinity. Thus, asymptotic resolvability for k candidates implies the following condition ($\star$): there is no uniquely weighted qualitative margin graph $\mathbb{M}$ with k candidates such that for every profile $\mathbf{P}$ realizing $\mathbb{M}$, $|F(\mathbf{P})| > 1$. This also follows from single-voter resolvability: for if there exists a uniquely weighted $\mathbb{M}$ such that every $\mathbf{P}$ realizing $\mathbb{M}$ has $|F(\mathbf{P})| > 1$, then we can pick a profile $\mathbf{P}$ realizing $\mathbb{M}$ with sufficiently many voters (note that if $\mathbf{P}$ realizes $\mathbb{M}$, so does $\mathbf{P} + \mathbf{P}^*$ where $\mathbf{P}^*$ is a copy of $\mathbf{P}$ with a disjoint set of voters) such that for any single-voter profile $\mathbf{P}'$, $\mathbf{P} + \mathbf{P}'$ still realizes $\mathbb{M}$ (since the differences between distinct margins are too large in $\mathbf{P}$ for one voter to change the qualitative margin graph), so that $|F(\mathbf{P} + \mathbf{P}')| > 1$, in violation of single-voter resolvability.

Now consider any profile $\mathbf{P}$ with $X(\mathbf{P}) > 3$ realizing a qualitative margin graph $\mathbb{M}$ that *when restricted to* $GETCHA(\mathbf{P})$ has the following form, where $\alpha \prec \gamma \prec \beta$ and $\gamma \prec \varphi \prec \psi$:

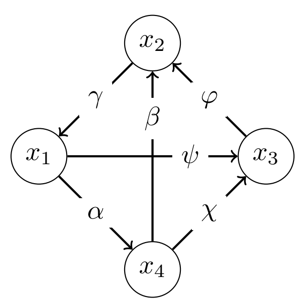

Since $\alpha \prec \gamma \prec \beta$, by the triangle property we have $x_4 \in F((\mathbf{P}^G)_{-x_3})$. Then given $x_4 \to x_3$, from stability for winners we have $x_4 \in F(\mathbf{P}^G)$ and hence $x_4 \in F(\mathbf{P})$ by the first paragraph of the proof. Since $\gamma \prec \varphi \prec \psi$, by the triangle property we have $x_1 \in F((\mathbf{P}^G)_{-x_4})$. Then given $x_1 \to x_4$, from stability for winners we have

$x_1 \in F(\mathbf{P}^G)$ and hence $x_1 \in F(\mathbf{P})$ by the first paragraph of the proof. Thus, $|F(\mathbf{P})| > 1$. Since this holds for every $\mathbf{P}$ realizing $\mathbb{M}$, condition ($\star$) above does not hold, so neither version of resolvability holds either. $\square$

It is easy to see that Split Cycle satisfies asymptotic resolvability for $k = 2$ and $k = 3$ (for $k = 3$, this follows from Corollary 4.14). For $k > 3$, since Split Cycle satisfies the triangle property and stability for winners, Theorem 6.20 yields the following.

Corollary 6.21. For $k > 3$, Split Cycle does not satisfy asymptotic resolvability for k candidates.

While it is certainly of theoretical interest to know whether the proportion of profiles with multiple winners goes to 0 as the number of voters goes to infinity, for real world applications, what matters is the proportion of profiles with multiple winners for realistic numbers of voters. In Appendix D, we provide a quantitative analysis. For instance, our results show that when there are 7 candidates and up to a few thousand voters, Split Cycle produces multiple winners on only about 1% more of such profiles than Beat Path, which satisfies resolvability in both forms above. Our results also show that this difference in frequency of multiple winners decreases as the number of candidates decreases. In addition, our analysis shows that Split Cycle is substantially more resolute than GETCHA. In Figure 3 in Section 4.4 and Figure 9 in Appendix D, we also compare the resoluteness of Split Cycle with that of Copeland, another Condorcet consistent voting method that fails resolvability in both forms above (see Appendix C.4).

6.5 Positive and Negative Involvement

Like rejectability and resolvability, the next two criteria we consider—positive and negative involvement—also concern adding voters to an election. In this case, the concern is about perverse changes to the set of winners in light of who the new voters ranked as their favorite (resp. least favorite) candidate. Recall our discussion in Section 1.2 of violations of positive or negative involvement as "strong no show paradoxes."

First of all, we note that if an election has a single Split Cycle winner, then adding a single new voter to the election can never remove that candidate from the Split Cycle winning set.

Definition 6.22. A voting method F satisfies *winner continuity* if for any profiles $\mathbf{P}$ and $\mathbf{P}'$ with $X(\mathbf{P}) = X(\mathbf{P}')$, $V(\mathbf{P}) \cap V(\mathbf{P}') = \varnothing$, and $|V(\mathbf{P}')| = 1$, if $F(\mathbf{P}) = \{x\}$, then $x \in F(\mathbf{P} + \mathbf{P}')$.

Note, for example, that Plurality satisfies winner continuity, while the Borda voting method does not (for the definition of Borda, see Pacuit 2019, § 2.1). We prove the following in Appendix B.1.

Proposition 6.23. Split Cycle satisfies winner continuity.

Now we consider the case where x is not initially the unique winner but is rather *among* the winners. We also consider the case where more than one new voter joins the election (Lemma 6.25). The criterion of positive (resp. negative) involvement ensures that if x is among the winners (resp. losers) and we add a group of voters who all rank x as their favorite (resp. least favorite), then x will still be a winner (resp. loser).

Definition 6.24. F satisfies *positive involvement* if for any profiles $\mathbf{P}$ and $\mathbf{P}'$ with $X(\mathbf{P}) = X(\mathbf{P}')$, $V(\mathbf{P}) \cap V(\mathbf{P}') = \varnothing$, and $|V(\mathbf{P}')| = 1$, if $x \in F(\mathbf{P})$ and for $i \in V(\mathbf{P}')$, $x\mathbf{P}'_i y$ for all $y \in X(\mathbf{P}') \setminus \{x\}$, then $x \in F(\mathbf{P} + \mathbf{P}')$.

F satisfies *negative involvement* if for any profiles $\mathbf{P}$ and $\mathbf{P}'$ with $X(\mathbf{P}) = X(\mathbf{P}')$, $V(\mathbf{P}) \cap V(\mathbf{P}') = \varnothing$, and $|V(\mathbf{P}')| = 1$, if $x \notin F(\mathbf{P})$ and for $i \in V(\mathbf{P}')$, $y\mathbf{P}'_i x$ for all $y \in X(\mathbf{P}') \setminus \{x\}$, then $x \notin F(\mathbf{P} + \mathbf{P}')$.

Lemma 6.25. If F satisfies positive involvement (resp. negative involvement), then it satisfies the analogous coalitional properties that drop the restriction that $|V(\mathbf{P}')| = 1$.

Proof. To prove the properties for a coalition of more than one voter, add each voter in the coalition one at a time, applying positive (resp. negative involvement) at each step. This can be iterated because the property of x belonging to (resp. not belonging to) the winning set is preserved at each step. □

Remark 6.26. It is important to distinguish positive and negative involvement from the *participation* criterion (recall Section 1.2), which we discuss further in Appendix B.2. It is also important that since we are assuming linear ballots, positive (resp. negative) involvement applies only when adding a voter for whom x is their *unique* favorite (resp. least favorite) candidate. For voting methods that accept non-linear ballots, one may consider a related criterion concerning voters for whom x is *among* their favorite (resp. least favorite) candidates (see Duddy 2014). But we see no problem with the addition of voters who rank x and y as tied changing the winner of an election with majority cycles from x to y, given how the new voters change x's and y's pairwise performance against other candidates.

None of Beat Path, Ranked Pairs, Copeland, GETCHA/GOCHA, or Uncovered Set satisfies positive or negative involvement. The failure of positive and negative involvement has been called "a common flaw in Condorcet voting correspondences" (Pérez, 2001). However, Split Cycle does not have this flaw.

Proposition 6.27. Split Cycle satisfies positive and negative involvement.

Proof. First, consider positive involvement. We prove the contrapositive. Suppose $x \notin SC(\mathbf{P} + \mathbf{P}')$. Hence there is a $z \in X(\mathbf{P})$ that defeats x in $\mathbf{P} + \mathbf{P}'$, i.e., such that

$$Margin_{\mathbf{P}+\mathbf{P}'}(z, x) > Cycle\#_{\mathbf{P}+\mathbf{P}'}(z, x). \tag{5}$$

Since $|V(\mathbf{P}')| = 1$, we have

$$Cycle\#_{\mathbf{P}+\mathbf{P}'}(z, x) \geq Cycle\#_{\mathbf{P}}(z, x) - 1, \tag{6}$$

and since $x\mathbf{P}'_i z$, we have

$$Margin_{\mathbf{P}+\mathbf{P}'}(z, x) = Margin_{\mathbf{P}}(z, x) - 1. \tag{7}$$

It follows from (5)–(7) that

$$Margin_{\mathbf{P}}(z, x) > Cycle\#_{\mathbf{P}}(z, x),$$

so $x \notin SC(\mathbf{P})$.

Next, consider negative involvement. Suppose $x \notin F(\mathbf{P})$. Hence there is a $z \in X(\mathbf{P})$ that defeats x in $\mathbf{P}$, i.e., such that

$$Margin_{\mathbf{P}}(z, x) > Cycle\#_{\mathbf{P}}(z, x). \tag{8}$$

Since $|V(\mathbf{P}')| = 1$, we have

$$Cycle\#_{\mathbf{P}+\mathbf{P}'}(z, x) \leq Cycle\#_{\mathbf{P}}(z, x) + 1, \tag{9}$$

and since $z\mathbf{P}'_i x$, we have

$$Margin_{\mathbf{P}+\mathbf{P}'}(z, x) = Margin_{\mathbf{P}}(z, x) + 1. \tag{10}$$

It follows from (8)–(10) that

$$Margin_{\mathbf{P}+\mathbf{P}'}(z,x) > Cycle\#_{\mathbf{P}+\mathbf{P}'}(z,x),$$

so $x \notin SC(\mathbf{P}+\mathbf{P}')$. □

Thus, with Split Cycle the strong no show paradox discussed in Section 1.2 is impossible.

7 Conclusion

In this paper, we have proposed the Split Cycle voting method, which can be distinguished from all methods we know of in any of the following three ways, using positive involvement, negative involvement, or immunity to spoilers:

- Only Split Cycle satisfies independence of clones, positive involvement, narrowing, and at least one of Condorcet consistency, monotonicity, and immunity to spoilers.
- Only Split Cycle satisfies independence of clones, negative involvement, and narrowing.
- Only Split Cycle satisfies independence of clones, immunity to spoilers, and rejectability.

Moreover, Split Cycle can be motivated by the three key ideas of Section 3:

1. Group incoherence raises the threshold for defeat, but not infinitely.
2. Incoherence can be localized.
3. Defeat is direct.

We think the third idea is especially important for justifying election outcomes to supporters of a candidate who was not among the winners of the election. To try to explain to supporters of a candidate x that the reason x is not among the winners is that another candidate y "defeated" x *even though a majority of voters prefer* x *to* y (as is possible with the Beat Path voting method, for example) seems a recipe for complaints of illegitimacy and resulting social instability.

There are several natural next steps for future research. For theoretical purposes, it would be desirable to have a set of axioms that single out Split Cycle not only from known voting methods but from all possible voting methods, providing a complete axiomatic characterization of Split Cycle (for steps in this direction, see Holliday and Pacuit 2020). For both theoretical and applied purposes, it would be desirable to have a more detailed quantitative analysis of how Split Cycle performs on profiles with various numbers of candidates and voters. The code we are making available at https://github.com/epacuit/splitcycle allows any researcher to perform such analyses. Ultimately, of course, the best test of Split Cycle would come from the use of the method in practice.

Acknowledgements

For helpful comments, we wish to thank Felix Brandt, Yifeng Ding, Mikayla Kelley, Chase Norman, Markus Schulze, Nicolaus Tideman, John Weymark, students in the Fall 2019 seminar on Voting and Democracy at UC Berkeley, and students in the Spring 2020 seminar on Preference and Judgment Aggregation at the

University of Maryland. We are also grateful for feedback received from the audiences at presentations of this work in the Logic Seminar at Stanford University on November 20, 2019, the 2nd Games, Agents, and Incentives Workshop (GAIW@AAMAS 2020) on May 10, 2020, and the Online Social Choice and Welfare Seminar Series on February 2, 2021.

A Independence of Clones

In this Appendix, we prove that Split Cycle satisfies independence of clones. In the following, fix a profile $\mathbf{P}$ with a set C of clones and $c \in C$. Then obviously we have the following.

Lemma A.1.

1. For any $a, b \in X(\mathbf{P}) \setminus \{c\}$, $Margin_{\mathbf{P}}(a, b) = Margin_{\mathbf{P}_{-c}}(a, b)$.
2. For any $b \in X(\mathbf{P}) \setminus C$ and $e \in C \setminus \{c\}$, $Margin_{\mathbf{P}}(c, b) = Margin_{\mathbf{P}_{-c}}(e, b)$.

Next we show that certain cycle numbers do not change from $\mathbf{P}$ to $\mathbf{P}_{-c}$. For this we use the following key lemma.

Lemma A.2. For any $c_1, c_2 \in C$ with $c_1 \neq c_2$ and simple cycle ρ in $\mathcal{M}(\mathbf{P})$ that contains c_1 and some non-clone, the sequence ρ' obtained from ρ by replacing all clones in ρ by c_2 and then replacing any subsequence $c_2, \dots, c_2$ by c_2 is a simple cycle in $\mathcal{M}(\mathbf{P})$ such that $Split\#(\rho') \geq Split\#(\rho)$.

Proof. For any $a \in X(\mathbf{P}) \setminus C$ and $d \in C$, if $a \to d$ (resp. $d \to a$) occurs in ρ with margin α in $\mathcal{M}(\mathbf{P})$, then by the definition of a set of clones (Definition 4.8), we have $a \to c_2$ (resp. $c_2 \to a$) with margin α in $\mathcal{M}(\mathbf{P})$. It follows that ρ' is a simple cycle in $\mathcal{M}(\mathbf{P})$ and that the margins between successive candidates in ρ' already occurred as margins between successive candidates in ρ, which implies $Split\#(\rho') \geq Split\#(\rho)$. ☐

Lemma A.3. For any $a \in X(\mathbf{P}) \setminus \{c\}$ and $b \in X(\mathbf{P}) \setminus C$, we have $Cycle\#_{\mathbf{P}}(a, b) = Cycle\#_{\mathbf{P}_{-c}}(a, b)$.

Proof. First, observe that any simple cycle in $\mathcal{M}(\mathbf{P}_{-c})$ is also a simple cycle with the same margins in $\mathcal{M}(\mathbf{P})$. Hence $Cycle\#_{\mathbf{P}}(a, b) \geq Cycle\#_{\mathbf{P}_{-c}}(a, b)$.

Second, to show $Cycle\#_{\mathbf{P}}(a, b) \leq Cycle\#_{\mathbf{P}_{-c}}(a, b)$, it suffices to show that for every simple cycle ρ in $\mathcal{M}(\mathbf{P})$ extending $a \to b$, there is a simple cycle ρ' in $\mathcal{M}(\mathbf{P}_{-c})$ extending $a \to b$ such $Split\#(\rho') \geq Split\#(\rho)$. If ρ does not contain c, then take $\rho' = \rho$. Suppose ρ does contain c. Case 1: $a \in C$. Then apply Lemma A.2 with $c_1 := c$ and $c_2 := a$ to obtain a simple cycle ρ' in $\mathcal{M}(\mathbf{P})$ extending $a \to b$, but not containing c, such that $Split\#(\rho') \geq Split\#(\rho)$; since ρ' does not contain c, it is also a simple cycle in $\mathcal{M}(\mathbf{P}_{-c})$ with the desired properties. Case 2: $a \notin C$. Then apply Lemma A.2 with $c_1 := c$, $c_2 \in C \setminus \{c\}$ and reason as in Case 1. ☐

Lemma A.4. Let $d \in C$ and $e \in C \setminus \{c\}$.

1. For any $b \in X(\mathbf{P}) \setminus C$, $Cycle\#_{\mathbf{P}}(d, b) = Cycle\#_{\mathbf{P}_{-c}}(e, b)$;
2. For any $a \in X(\mathbf{P}) \setminus C$, $Cycle\#_{\mathbf{P}}(a, d) = Cycle\#_{\mathbf{P}_{-c}}(a, e)$.

Proof. For part 1, for any simple cycle ρ in $\mathcal{M}(\mathbf{P})$ extending $d \to b$, by Lemma A.2 with $c_1 := d$ and $c_2 := e$, there is a simple cycle ρ' in $\mathcal{M}(\mathbf{P})$ extending $e \to b$, but not containing c (since $e \in C \setminus \{c\}$), such that $Split\#(\rho') \geq Split\#(\rho)$. Since ρ' does not contain c, it is also a simple cycle in $\mathcal{M}(\mathbf{P}_{-c})$ extending $e \to b$ with the same margins. Hence $Cycle\#_{\mathbf{P}}(d, b) \leq Cycle\#_{\mathbf{P}_{-c}}(e, b)$. Next, suppose ρ is a simple cycle in

$\mathcal{M}(\mathbf{P}_{-c})$ extending $e \to b$. Then ρ is also a simple cycle in $\mathcal{M}(\mathbf{P})$ extending $e \to b$ with the same margins. Thus, by Lemma A.2 with $c_1 := e$ and $c_2 := d$, there is a simple cycle ρ' in $\mathcal{M}(\mathbf{P})$ extending $d \to b$ such that $Split\#(\rho') \geq Split\#(\rho)$. Hence $Cycle\#_{\mathbf{P}}(d,b) \geq Cycle\#_{\mathbf{P}_{-c}}(e,b)$.

The proof of part 2 is analogous. □

Proposition A.5. For any $b \in X(\mathbf{P}) \setminus C$, we have $b \in SC(\mathbf{P})$ if and only if $b \in SC(\mathbf{P}_{-c})$. Hence Split Cycle is such that non-clone choice is independent of clones.

Proof. Suppose $b \notin SC(\mathbf{P}_{-c})$, so there is an $a \in X(\mathbf{P}) \setminus \{c\}$ such that a defeats b in $\mathbf{P}_{-c}$. Then by Lemmas A.1.1 and A.3, a defeats b in $\mathbf{P}$. Now suppose $b \notin SC(\mathbf{P})$, so there is an $a \in X(\mathbf{P})$ such that a defeats b in $\mathbf{P}$. Case 1: $a \neq c$. Then by Lemmas A.1.1 and A.3 again, a defeats b in $\mathbf{P}_{-c}$. Case 2: $a = c$. Then by Lemmas A.1.2 and A.4.1 with $d := c$, each $e \in C \setminus \{c\}$ defeats b in $\mathbf{P}_{-c}$. □

Proposition A.6. $C \cap SC(\mathbf{P}) \neq \varnothing$ if and only if $C \setminus \{c\} \cap SC(\mathbf{P}_{-c}) \neq \varnothing$. Hence Split Cycle is such that clone choice is independent of clones.

Proof. Suppose $C \cap SC(\mathbf{P}) = \varnothing$. Hence every clone in C is defeated in $\mathbf{P}$. Since the defeat graph for $\mathbf{P}$ contains no cycles (Lemma 3.9), it follows that there is some $a \in X(\mathbf{P}) \setminus C$ that defeats some $d \in C$ in $\mathbf{P}$. It then follows by the definition of a set of clones (Definition 4.8), Lemma A.1.1, and Lemma A.4.2 that a defeats every $e \in C \setminus \{c\}$ in $\mathbf{P}_{-c}$. Hence $C \setminus \{c\} \cap SC(\mathbf{P}_{-c}) = \varnothing$. Similarly, if $C \setminus \{c\} \cap SC(\mathbf{P}_{-c}) = \varnothing$, then there is some $a \in X(\mathbf{P}_{-c}) \setminus C$ that defeats some $e \in C \setminus \{c\}$ in $\mathbf{P}_{-c}$. It then follows by Definition 4.8, Lemma A.1.1, and Lemma A.4.2 that a defeats every $d \in C$ in $\mathbf{P}$. Hence $C \cap SC(\mathbf{P}) = \varnothing$. □

Theorem A.7. Split Cycle satisfies independence of clones.

Proof. By Propositions A.5 and A.6. □

B Winner Continuity and Participation

B.1 Winner Continuity

Recall that a voting method F satisfies *winner continuity* if for any profiles $\mathbf{P}$ and $\mathbf{P}'$ with $X(\mathbf{P}) = X(\mathbf{P}')$, $V(\mathbf{P}) \cap V(\mathbf{P}') = \varnothing$, and $|V(\mathbf{P}')| = 1$, if $F(\mathbf{P}) = \{x\}$, then $x \in F(\mathbf{P}+\mathbf{P}')$.

Proposition B.1. Split Cycle satisfies winner continuity.

Proof. Suppose $SC(\mathbf{P}) = \{x\}$. Further suppose $x \notin SC(\mathbf{P}+\mathbf{P}')$, so there is some $z \in X(\mathbf{P})$ such that

$$Margin_{\mathbf{P}+\mathbf{P}'}(z,x) > Cycle\#_{\mathbf{P}+\mathbf{P}'}(z,x). \tag{11}$$

Since $z \notin SC(\mathbf{P})$, by Lemma 3.11 there are distinct $y_1, \dots, y_n$ with $y_1 = x$ and $y_n = z$ such that $y_1 D y_2 D \dots D y_{n-1} D y_n$ in the defeat graph of $\mathbf{P}$.

Since $|V(\mathbf{P}')| = 1$, it follows from (11) that $Margin_{\mathbf{P}}(z,x) = 0$ or $Margin_{\mathbf{P}}(z,x) > 0$.

Case 1: $Margin_{\mathbf{P}}(z,x) = 0$. Then for each $i \in \{1, \dots, n-1\}$, $Margin_{\mathbf{P}}(y_i, y_{i+1})$ is even, and since $y_i D y_{i+1}$, it is greater than 0, so $Margin_{\mathbf{P}}(y_i, y_{i+1}) \geq 2$. Since $|V(\mathbf{P}')| = 1$, together $Margin_{\mathbf{P}}(z,x) = 0$

and (11) imply $Margin_{\mathbf{P}+\mathbf{P}'}(z,x) = 1$. In addition, since $|V(\mathbf{P}')| = 1$, from $Margin_{\mathbf{P}}(y_i, y_{i+1}) \geq 2$ we have $Margin_{\mathbf{P}+\mathbf{P}'}(y_i, y_{i+1}) \geq 1$. Thus, we have a simple cycle

$$\rho = y_1 \xrightarrow{\gamma_1} y_2 \xrightarrow{\gamma_2} \dots \xrightarrow{\gamma_{n-1}} y_n \xrightarrow{\delta} y_1$$

in the margin graph of $\mathbf{P} + \mathbf{P}'$ in which δ, i.e., $Margin_{\mathbf{P}+\mathbf{P}'}(z,x)$, is not greater than any γ_i. But this contradicts (11).

Case 2: $Margin_{\mathbf{P}}(z,x) > 0$. Together with $y_1 D y_2 D \dots D y_{n-1} D y_n$, this means there is a simple cycle

$$\rho = y_1 \xrightarrow{\alpha_1} y_2 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_{n-1}} y_n \xrightarrow{\beta} y_1$$

in the margin graph of $\mathbf{P}$. Moreover, from $y_1 D y_2 D \dots D y_{n-1} D y_n$, it follows that for each $i \in \{1, \dots, n-1\}$, α_i is greater than the splitting number of ρ; hence β, i.e., $Margin_{\mathbf{P}}(z,x)$, is the splitting number of ρ. Thus, for each $i \in \{1, \dots, n-1\}$, we have $Margin_{\mathbf{P}}(y_i, y_{i+1}) \geq Margin_{\mathbf{P}}(z,x) + 2$ since the parity of all margins must be the same. Since $|V(\mathbf{P}')| = 1$, it follows that there is a simple cycle

$$\rho^\star = y_1 \xrightarrow{\alpha_1^\star} y_2 \xrightarrow{\alpha_2^\star} \dots \xrightarrow{\alpha_{n-1}^\star} y_n \xrightarrow{\beta^\star} y_1$$

in the margin graph of $\mathbf{P} + \mathbf{P}'$ in which $\beta^\star$, i.e., $Margin_{\mathbf{P}+\mathbf{P}'}(z,x)$, is not greater than any $\alpha_i^\star$. But this contradicts (11). □

B.2 Participation

In Section 6.5 on positive and negative involvement, we mentioned the related *participation* criterion. Participation is usually stated for resolute voting methods: if x is the winner in a profile, and we add to the profile a new voter who strictly prefers x to y, then y is not the winner in the resulting profile. (Note that there is no requirement that x be at the top of the new voter's ballot or that y be at the bottom, a point to which we return below.) When applied to irresolute voting methods, we call this "resolute" participation.[34]

Definition B.2. A voting method F satisfies *resolute participation* if for any profiles $\mathbf{P}$ and $\mathbf{P}'$ with $X(\mathbf{P}) = X(\mathbf{P}')$, $V(\mathbf{P}) \cap V(\mathbf{P}') = \varnothing$, and $|V(\mathbf{P}')| = 1$, and any $x, y \in X$, if $F(\mathbf{P}) = \{x\}$ and $x\mathbf{P}'_i y$ for $i \in V(\mathbf{P}')$, then $F(\mathbf{P} + \mathbf{P}') \neq \{y\}$.

Corollary B.3. Split Cycle satisfies resolute participation.

Proof. Immediate from Proposition B.1. □

While positive and negative involvement entail the analogous coalitional properties (recall Lemma 6.25), resolute participation does not entail the analogous coalitional property.

[34] Several authors have investigated what could be called "irresolute" participation-like criteria (recall Footnote 10), where one changes the initial assumption from $F(\mathbf{P}) = \{x\}$ to $x \in F(\mathbf{P})$. For example, Perez (2001) considers the following axiom, called *VC-participation*: if $x \in F(\mathbf{P})$ and $\mathbf{P}'$ is a one-voter profile with a new voter i having $x\mathbf{P}'_i y$, then $y \in F(\mathbf{P} + \mathbf{P}')$ implies $x \in F(\mathbf{P} + \mathbf{P}')$. He then observes that no Condorcet consistent voting method satisfies VC-participation. However, it is not clear that this criterion is a plausible normative requirement on a voting method. Suppose, for example, that i's ranking is $zxyw$, and i's joining the election results in a change from $F(\mathbf{P}) = \{x, w\}$ to $F(\mathbf{P} + \mathbf{P}') = \{z, y\}$. It is not clear that we should impose a criterion that prohibits such a change, which seems to be a strict improvement from i's point of view.

Definition B.4. A voting method F satisfies *resolute coalitional participation* if for any profiles $\mathbf{P}$ and $\mathbf{P}'$ with $X(\mathbf{P}) = X(\mathbf{P}')$ and $V(\mathbf{P}) \cap V(\mathbf{P}') = \varnothing$ and any $x, y \in X$, if $F(\mathbf{P}) = \{x\}$ and $x\mathbf{P}'_i y$ for all $i \in V(\mathbf{P}')$, then $F(\mathbf{P} + \mathbf{P}') \neq \{y\}$.

Proposition B.5. Split Cycle does not satisfy resolute coalitional participation.

Proof. By Debord's Theorem, there is a profile $\mathbf{P}$ whose margin graph is shown on the left below:

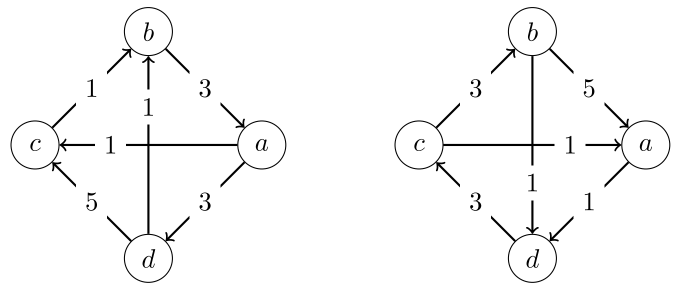

On the right, we show the margin graph of the profile $\mathbf{P}+\mathbf{P}'$ where $\mathbf{P}'$ is a two-voter profile whose two voters both have $c\mathbf{P}'_i b\mathbf{P}'_i d\mathbf{P}'_i a$. Although both voters have $b\mathbf{P}'_i d$, we go from $SC(\mathbf{P}) = \{b\}$ to $SC(\mathbf{P}+\mathbf{P}') = \{d\}$. □

In our view, the example in the proof of Proposition B.5 shows that resolute coalitional participation is too strong to require. Its violation can be rationalized as follows. In $\mathbf{P}$, d is defeated by a; yet with the two new voters having $d\mathbf{P}'_i a$, d is no longer defeated by a (or anyone else) in $\mathbf{P} + \mathbf{P}'$. In $\mathbf{P}$, b is not defeated by c (or anyone else); yet with the two new voters having $c\mathbf{P}'_i b$, b becomes defeated by c in $\mathbf{P} + \mathbf{P}'$. In short, the two new voters help d against its main threat, a, and hurt b against its main threat, c, resulting in the change of the winning set from $\{b\}$ to $\{d\}$. It does not matter, in this case, that the two new voters help b against d, because b and d do not threaten to defeat each other in the presence of the cycles.

More decisively, if b is the unique winner for the margin graph on the left above, then d must be the unique winner for the margin graph on the right above, assuming a neutrality property for margin graphs—that the names assigned to nodes do not matter—satisfied by Split Cycle (and the methods in Appendices C.1-C.7).

Definition B.6. A voting method F satisfies *margin graph neutrality* if for any profiles $\mathbf{P}$ and $\mathbf{P}'$, if there is a weighted directed graph isomorphism $h : \mathcal{M}(\mathbf{P}) \to \mathcal{M}(\mathbf{P}')$, then $F(\mathbf{P}') = h[F(\mathbf{P})]$.

For example, the map $c \mapsto a$, $a \mapsto c$, $b \mapsto d$, $d \mapsto b$ is a weighted directed graph isomorphism from the margin graph on the left above to the margin graph on the right above (imagine turning the margin graph on the left 180°—then it matches the margin graph on the right except for the names of nodes). Despite the fact that the right margin graph is obtained from the left margin graph by adding two voters who rank b over d, candidate b on the left and candidate d on the right are in isomorphic situations. Thus, if b is the winner on the left, d must be the winner on the right by margin graph neutrality.

Finally, note that the phenomenon with b and d above can happen only when b and d are in a cycle. Indeed, we have the following version of participation when the two relevant candidates are cycle free.

Proposition B.7. For any profiles $\mathbf{P}$ and $\mathbf{P}'$ with $X(\mathbf{P}) = X(\mathbf{P}')$ and $V(\mathbf{P}) \cap V(\mathbf{P}') = \varnothing$ and any $x, y \in X$, if $x \in SC(\mathbf{P})$ and $x\mathbf{P}'_i y$ for all $i \in V(\mathbf{P}')$, and there is no cycle in $\mathcal{M}(\mathbf{P})$ or $\mathcal{M}(\mathbf{P} + \mathbf{P}')$ containing x and y, then $y \notin SC(\mathbf{P} + \mathbf{P}')$.

Proof. Since $x \in SC(\mathbf{P})$, y does not defeat x in $\mathbf{P}$. Since there is no cycle in $\mathcal{M}(\mathbf{P})$ involving x and y, it follows that $Margin_{\mathbf{P}}(x, y) \geq 0$. Hence $Margin_{\mathbf{P}+\mathbf{P}'}(x, y) > 0$, and by hypothesis, there is no cycle involving in $\mathcal{M}(\mathbf{P} + \mathbf{P}')$. Hence x defeats y in $\mathbf{P} + \mathbf{P}'$, so $y \notin SC(\mathbf{P} + \mathbf{P}')$. □

C Other Methods

In this appendix, we provide definitions of the other voting methods mentioned in the main text, as well citations or proofs for our claims about their properties.

C.1 Ranked Pairs (Tideman 1987)

Let $\mathbf{P}$ be a profile and $T \in \mathcal{L}(X(\mathbf{P}) \times X(\mathbf{P}))$ a linear order on the set of pairs of candidates (the tie-breaking ordering). We say that a pair (x, y) of candidates has a *higher priority* than a pair (x', y') of candidates using the tie-breaking ordering T when either $Margin_{\mathbf{P}}(x, y) > Margin_{\mathbf{P}}(x', y')$ or $Margin_{\mathbf{P}}(x, y) = Margin_{\mathbf{P}}(x', y')$ and $(x, y)\ T\ (x', y')$. Given a profile $\mathbf{P}$ and a tie-breaking ordering $T \in \mathcal{L}(X(\mathbf{P}) \times X(\mathbf{P}))$, we construct a *Ranked Pairs ranking* $\succ_{\mathbf{P},T} \in \mathcal{L}(X)$ according to the following procedure:

1. Initialize $\succ_{\mathbf{P},T}$ to $\varnothing$.
2. If all pairs (x, y) with $Margin_{\mathbf{P}}(x, y) \geq 0$ have been considered, then return $\succ_{\mathbf{P},T}$. Otherwise let (a, b) be the pair with the highest priority among those with $Margin_{\mathbf{P}}(a, b) \geq 0$ that have not been considered so far.
3. If $\succ_{\mathbf{P},T} \cup \{(a, b)\}$ is acyclic, then add (a, b) to $\succ_{\mathbf{P},T}$; otherwise, add (b, a) to $\succ_{\mathbf{P},T}$. Go to step 2.

When the procedure terminates, $\succ_{\mathbf{P},T}$ is a linear order. A linear order L on $X(\mathbf{P})$ is a *Ranked Pairs ranking for* $\mathbf{P}$ if $L = \succ_{\mathbf{P},T}$ for some tie-breaking ordering $T \in \mathcal{L}(X(\mathbf{P}) \times X(\mathbf{P}))$. Then the set $RP(\mathbf{P})$ of Ranked Pairs winners is the set of all $x \in X(\mathbf{P})$ such that x is the maximum of some Ranked Pairs ranking for $\mathbf{P}$.

Another useful characterization of Ranked Pairs rankings is given by Zavist and Tideman (1989). Given a profile $\mathbf{P}$, a linear order L on $X(\mathbf{P})$ is a *stack for* $\mathbf{P}$ if for any $a, b \in X(\mathbf{P})$, if aLb, then there are distinct $x_1, \ldots, x_n \in X(\mathbf{P})$ with $x_1 = a$ and $x_n = b$ such that $x_i L x_{i+1}$ and $Margin_{\mathbf{P}}(x_i, x_{i+1}) \geq Margin_{\mathbf{P}}(b, a)$ for all $i \in \{1, \ldots, n-1\}$.

Lemma C.1 (Zavist and Tideman 1989)**.** For any profile $\mathbf{P}$ and linear order L on $X(\mathbf{P})$, L is a Ranked Pairs ranking for $\mathbf{P}$ if and only if L is a stack for $\mathbf{P}$.[35]

See Brill and Fischer 2012, Wang et al. 2019, and Lamboray 2008 for discussion of the computational and axiomatic properties of Ranked Pairs.

Core criteria Proofs that Ranked Pairs satisfies the Condorcet winner, Condorcet loser, and monotonicity criteria can be found in Tideman 1987. See Remark 4.11 on the status of independence of clones. Narrowing follows from asymptotic resolvability below.

[35] Using this lemma, we can also prove a strengthened version of Lemma 3.18: for any profile $\mathbf{P}$ and $a, b \in X(\mathbf{P})$, if b defeats a according to Split Cycle (Definition 3.3), then bLa for any Ranked Pairs ranking L for $\mathbf{P}$.

New criteria Proposition 5.2 shows that Ranked Pairs violates immunity to spoilers and hence the stronger criteria of stability for winners and amalgamation. Corollary 5.16 shows that it satisfies rejectability.

Other criteria For the satisfaction of Smith, ISDA, and reversal symmetry, see Schulze 2011, Table 2. For the satisfaction of single-voter resolvability, see Tideman 1987. Asymptotic resolvability follows from the fact that the proportion of profiles that are uniquely weighted goes to 1 as the number of voters goes to infinity, and Ranked Pairs selects a unique winner in any uniquely weighted profile. For the failure of positive involvement and negative involvement, see Pérez 2001, p. 612.

C.2 Beat Path (Schulze 2011)

Let $\mathcal{M}$ be a margin graph. A (*simple*) *path from* x *to* y *in* $\mathcal{M}$ is a sequence $\langle x_1, \dots, x_k \rangle$ of distinct nodes in $\mathcal{M}$ where $x_1 = x$, $x_k = y$, and for $i \in \{1, \dots, k-1\}$, $x_i \overset{\alpha_i}{\to} x_{i+1}$. The strength of a path $\langle x_1, \dots, x_k \rangle$ in $\mathcal{M}$ is

$$S_{\mathcal{M}}(\langle x_1, \dots, x_k \rangle) = \min\{\alpha_i \mid x_i \overset{\alpha_i}{\to} x_{i+1}, 1 \leq i \leq k-1\}.$$

Given a profile $\mathbf{P}$, let $Path_{\mathbf{P}}(x, y)$ be the set of all paths from x to y in $\mathcal{M}(\mathbf{P})$. The strength of x over y in $\mathbf{P}$ is

$$Strength_{\mathbf{P}}(x, y) = \begin{cases} \max\{S_{\mathcal{M}(\mathbf{P})}(p) \mid p \in Path_{\mathbf{P}}(x, y)\} & Path_{\mathbf{P}}(x, y) \neq \varnothing \\ 0 & \text{otherwise.} \end{cases}$$

Then the set $BP(\mathbf{P})$ of Beat Path winners is the set of all $x \in X(\mathbf{P})$ such that there is no $y \in X(\mathbf{P})$ such that $Strength_{\mathbf{P}}(y, x) > Strength_{\mathbf{P}}(x, y)$.

Core criteria For proofs that Beat Path satisfies the Condorcet winner, Condorcet loser, monotonicity, and independence of clones criteria, see Schulze 2018. Narrowing follows from asymptotic resolvability below.

New criteria Proposition 5.2 shows that Beat Path violates immunity to spoilers and hence also the stronger criteria of stability for winners and amalgamation. The failure of amalgamation is also shown in Example 5.7 (d is a Beat Path winner in both $\mathbf{P}$ and $\mathbf{P}'$ but not in $\mathbf{Q}$). Corollary 5.16 shows that Beat Path satisfies rejectability.

Other criteria For the satisfaction of Smith, ISDA, single-voter and asymptotic resolvability, and reversal symmetry, see Schulze 2018. For an example of the simultaneous failure of positive and negative involvement, see Example 7 of Schulze 2018, pp. 68-95.

C.3 Minimax (Simpson 1969; Kramer 1977)

The set of winners for Minimax, also known as the Simpson-Kramer method, are the candidates whose largest majority loss is the smallest, i.e., for a profile $\mathbf{P}$,

$$Minimax(\mathbf{P}) = \text{argmin}_{x \in X} \max(\{Margin_{\mathbf{P}}(y, x) \mid y \in X\}).$$

Core criteria Examples of violations of the Condorcet loser criterion and proofs of the Condorcet winner and monotonicity criteria can be found in Felsenthal 2012. Violation of independence of clones is discussed in Tideman 1987. Narrowing follows from asymptotic resolvability below.

New criteria For immunity to spoilers, if $a \in Minimax(\mathbf{P}_{-b})$, $Margin_{\mathbf{P}}(a,b) > 0$, and $b \notin Minimax(\mathbf{P})$, then a must still be among the candidates in $\mathbf{P}$ whose largest majority loss is smallest, so $a \in Minimax(\mathbf{P})$. Proposition 5.4 shows that Minimax violates stability for winners and hence the stronger property of amalgamation. For rejectability, given $x \in Minimax(\mathbf{P})$, modify the margin graph $\mathcal{M}(\mathbf{P})$ to $\mathcal{M}'$ such that for all $y \in X(\mathbf{P}) \setminus \{x\}$, (i) if there is no edge from x to y in $\mathcal{M}(\mathbf{P})$, add an edge from x to y in $\mathcal{M}'$, and (ii) increase the weights of all incoming edges to y to be larger than the largest majority loss of x in $\mathbf{P}$, such that all weights in $\mathcal{M}'$ are even or all weights in $\mathcal{M}'$ are odd. Then x is clearly the unique Minimax winner in $\mathcal{M}'$, and $\mathcal{M}'$ is the margin graph of a profile $\mathbf{P}'$ by Debord's Theorem. Finally, since Minimax clearly satisfies the overwhelming majority criterion (recall Lemma 5.14), $\mathbf{P}'$ may be used to obtain the $\mathbf{P}^{+}$ required for rejectability as in the proof of Proposition 5.15.

Other criteria For the failure of Smith and hence ISDA, see Darlington 2016, p. 10. For the satisfaction of single-voter resolvability, see Tideman 1987, and for asymptotic resolvability, the argument is the same as given for Ranked Pairs in Appendix C.1. For the failure of reversal symmetry, see Felsenthal 2012 (under 'preference inversion'). For the satisfaction of positive and negative involvement, see Pérez 2001, p. 613.

C.4 Copeland (Copeland 1951)

The Copeland score of a candidate x is the number of candidates to whom x is majority preferred minus the number majority preferred to x. The Copeland winners are the candidates with maximal Copeland score:

$$Copeland(\mathbf{P}) = \text{argmax}_{x \in X} \, |\{y \in X(\mathbf{P}) \mid Margin_{\mathbf{P}}(x,y) > 0\}| - |\{y \in X(\mathbf{P}) \mid Margin_{\mathbf{P}}(y,x) > 0\}|.$$

Core criteria It is easy to see that Copeland satisfies the Condorcet winner, Condorcet loser, and monotonicity criteria. For the failure of independence of clones, see Tideman 1987. For the satisfaction of the narrowing criterion, see Fey 2008, § 4.

New criteria For immunity to spoilers, if $a \in Copeland(\mathbf{P}_{-b})$, so a's Copeland score is maximal in $\mathbf{P}_{-b}$, and $Margin_{\mathbf{P}}(a,b) > 0$, then a's Copeland score in $\mathbf{P}$ is maximal *among the original candidates in* $X(\mathbf{P}_{-b})$; if in addition $b \notin Copeland(\mathbf{P})$, then a's Copeland score in $\mathbf{P}$ is maximal among all candidates in $X(\mathbf{P})$, so $a \in Copeland(\mathbf{P})$. However, if we do not assume $b \notin Copeland(\mathbf{P})$, then it is easy to construct profiles in which b has a higher Copeland score in $\mathbf{P}$ than a does (this requires $|X(\mathbf{P})| \geq 4$ with an even number of voters or $|X(\mathbf{P})| \geq 6$ with an odd number of voters), so that $a \notin Copeland(\mathbf{P})$. Thus, Copeland violates stability for winners and hence amalgamation. For the failure of rejectability, see Proposition 5.13.

Other criteria Smith, ISDA, and reversal symmetry are easy to check. That Copeland fails single-voter resolvability follows from the same reasoning as in the proof of Proposition 5.13. For the failure of asymptotic resolvability for $k \geq 3$, consider a majority graph consisting of three candidates in a top cycle followed by a linear order of the remaining candidates, so the top three candidates are Copeland winners. The proportion of profiles realizing such a majority graph does not go to 0 as the number of voters goes to infinity (Harrison-Trainor 2020). Finally, for the failure of positive and negative involvement, see Pérez 2001, § 4.1.

C.5 GETCHA (Smith 1973)

For the definition of GETCHA, see Definition 6.1 in Section 6.1.

First of all, to see that GETCHA fails Pareto, consider the following example:

1	1	1
a	b	c
x	c	a
b	a	x
c	x	b

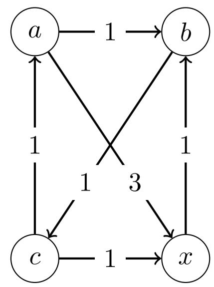

All voters prefer a to x, but x is among the GETCHA winners.

Core criteria It is easy to see that GETCHA satisfies the Condorcet winner, Condorcet loser, and monotonicity criteria. It is also not difficult to see that GETCHA satisfies independence of clones, using an alternative characterization of GETCHA. Given a profile $\mathbf{P}$, let $a \succsim_{\mathbf{P}} b$ mean that $Margin_{\mathbf{P}}(a,b) \geq 0$. Let $\succsim^*_{\mathbf{P}}$ be the transitive closure of $\succsim_{\mathbf{P}}$.

Lemma C.2 (Schwartz 1986, Corollary 6.2.2)**.** For any profile $\mathbf{P}$,

$$GETCHA(\mathbf{P}) = \{x \in X(\mathbf{P}) \mid \text{for all } y \in X(\mathbf{P}) : x \succsim^*_{\mathbf{P}} y\}.$$

Finally, for the failure of narrowing, see Section 4.4.

New criteria It is easy to see that GETCHA satisfies immunity to spoilers, stability for winners, and amalgamation. For the failure of rejectability, see Proposition 5.13.

Other criteria By definition, GETCHA satisfies the Smith criterion, and we proved ISDA in Footnote 33. For reversal symmetry, note that if $\mathbf{P}$ is a profile with $GETCHA(\mathbf{P}) = \{x\}$, then x is a Condorcet winner. Thus, x is a Condorcet loser in $\mathbf{P}^r$, so $x \notin GETCHA(\mathbf{P}^r)$. That GETCHA fails single-voter resolvability follows from the same reasoning as in the proof of Proposition 5.13. The failure of asymptotic resolvability for $k \geq 3$ follows from the fact that GETCHA selects a unique winner only if there is a Condorcet winner, and for 3 or more candidates, the proportion of profiles with a Condorcet winner does not go to 1 as the number of voters goes to infinity (DeMeyer and Plott 1970). To see that GETCHA fails positive involvement, observe that in a three-candidate, two-voter profile $\mathbf{P}$ such that $a\mathbf{P}_i b\mathbf{P}_i c$ and $c\mathbf{P}_j a\mathbf{P}_j b$, we have $GETCHA(\mathbf{P}) = \{a,b,c\}$, yet adding one voter k such that $b\mathbf{P}_k a\mathbf{P}_k c$ results in a profile $\mathbf{P}'$ in which a is the Condorcet winner, so $GETCHA(\mathbf{P}') = \{a\}$. Finally, to see that GETCHA fails negative involvement, consider any three-candidate profile $\mathbf{P}$ with a Condorcet winner a, so $GETCHA(\mathbf{P}) = \{a\}$, in which $Margin_{\mathbf{P}}(a,c) = 1$;

then adding a voter with the ranking $c\mathbf{P}_i a\mathbf{P}_i b$ results in a profile $\mathbf{P}'$ in which $Margin_{\mathbf{P}'}(a,b) > 0$ but $Margin_{\mathbf{P}'}(a,c) = 0$, which implies $GETCHA(\mathbf{P}') = \{a,b,c\}$.

C.6 GOCHA (Schwartz 1986)

For the definition of GOCHA, see Definition 6.4 in Section 6.1. Note that GOCHA violates Pareto, using the same example given for GETCHA in Section C.5.

Core criteria Clearly GOCHA satisfies the Condorcet winner, Condorcet loser, and monotonicity criteria (see Felsenthal 2012). For the satisfaction of independence of clones, see Tideman 1987. Finally, for narrowing, as the number of voters goes to infinity, the proportion of profiles for which there are no zero margins between distinct candidates goes to 1, and in such profiles GOCHA is equivalent to GETCHA. Thus, GETCHA's violating the narrowing criterion implies that GOCHA does as well.

New criteria It is easy to see that GOCHA satisfies immunity to spoilers and stability for winners using Lemma 6.5. Example 5.7 shows that GOCHA does not satisfy amalgamation. For rejectability, see Proposition 5.13.

Other criteria It is well known that GOCHA satisfies the Smith criterion, and ISDA can be proved using Lemma 6.5.[36] For reversal symmetry, see Felsenthal 2012. GOCHA fails single-voter and asymptotic resolvability for $k \geq 3$ by the same reasoning as for GETCHA, using the fact above that in the limit as the number of voters goes infinity, GOCHA is equivalent to GETCHA. For the failure of positive involvement and negative involvement, see Pérez 2001, § 4.1, where GOCHA is called "Top Cycle", or Felsenthal and Nurmi 2016, where GOCHA is called "Schwartz". As a simple example of the failure of positive involvement, consider a three cycle with $y \overset{5}{\to} x \overset{3}{\to} z \overset{1}{\to} y$, so $GOCHA(\mathbf{P}) = \{x,y,z\}$, and add a new voter with $x\mathbf{P}'_i y\mathbf{P}'_i z$ to obtain $y \overset{4}{\to} x \overset{4}{\to} z$ with y and z tied, so $GOCHA(\mathbf{P}') = \{y\}$.

C.7 Uncovered Set (Fishburn 1977; Miller 1980)

The Uncovered Set in voting is usually attributed to Fishburn (1977) and Miller (1980), though the covering relation appears in earlier game-theoretic work of Gillies (1959). Fishburn defined his version of the Uncovered Set for arbitrary margin graphs, whereas Miller defined his only for *tournaments*, i.e., directed graphs in which the edge relation $\to$ is not only asymmetric but also *weakly complete*: for all distinct nodes x, y, either $x \to y$ or $y \to x$. Fishburn and Miller's definitions are equivalent for tournaments but not for margin graphs that are not weakly complete, which may arise from profiles with an even number of voters. Several non-equivalent definitions of the Uncovered Set for arbitrary margin graphs appear in the literature

[36] Suppose $a \notin GOCHA(\mathbf{P}_{-x})$, so there is a $b \in X(\mathbf{P}_{-x})$ such that $b \to^*_{\mathbf{P}_{-x}} a$ but $a \not\to^*_{\mathbf{P}_{-x}} b$. Then $b \to^*_{\mathbf{P}} a$. If $a \not\to^*_{\mathbf{P}} b$, then we are done: $a \notin GOCHA(\mathbf{P})$. If $a \to^*_{\mathbf{P}} b$, then given $a \not\to^*_{\mathbf{P}_{-x}} b$, it follows that x is on the path witnessing $a \to^*_{\mathbf{P}} b$, which together with $b \to^*_{\mathbf{P}} a$ implies that there is a path from x to a in $\mathcal{M}(\mathbf{P})$. Then since $x \notin GETCHA(\mathbf{P})$ and there can be no path from a candidate outside $GETCHA(\mathbf{P})$ to one inside $GETCHA(\mathbf{P})$, it follows that $a \notin GETCHA(\mathbf{P})$. Hence $a \notin GOCHA(\mathbf{P})$, since GOCHA satisfies the Smith criterion. Conversely, suppose $a \notin GOCHA(\mathbf{P})$, so there is a $b \in X(\mathbf{P})$ such that $b \to^*_{\mathbf{P}} a$ but $a \not\to^*_{\mathbf{P}} b$. If x is not on the path witnessing $b \to^*_{\mathbf{P}} a$, then $b \to^*_{\mathbf{P}_{-x}} a$ but $a \not\to^*_{\mathbf{P}_{-x}} b$, so we are done: $a \notin GOCHA(\mathbf{P}_{-x})$. If x is on the path witnessing $b \to^*_{\mathbf{P}} a$, then since $x \notin GETCHA(\mathbf{P})$ and there can be no path from a candidate outside $GETCHA(\mathbf{P})$ to one inside $GETCHA(\mathbf{P})$, it follows that $a \notin GETCHA(\mathbf{P})$. Hence by ISDA for GETCHA, $a \notin GETCHA(\mathbf{P}_{-x})$ and hence $a \notin GOCHA(\mathbf{P}_{-x})$, since GOCHA satisfies the Smith criterion.

(see Bordes 1983; Peris and Subiza 1999; Penn 2006; Duggan 2013), and some of these versions differ in their axiomatic properties. As examples, we will consider the versions due to Fishburn and Gillies.

Given a margin graph $\mathcal{M}$ and nodes x, y in $\mathcal{M}$, say that y *left-covers* x *in* $\mathcal{M}$ if for all nodes z in $\mathcal{M}$, if $z \to y$, then $z \to x$.[37] Then the Fishburn and Gillies versions of the Uncovered Set are defined by:

$$\begin{aligned} UC_{Fish}(\mathbf{P}) &= \{x \in X(\mathbf{P}) \mid \text{there is no } y \in X(\mathbf{P})\text{: } y \text{ left-covers } x \text{ but } x \text{ does not left-cover } y \text{ in } \mathcal{M}(\mathbf{P})\}; \\ UC_{Gill}(\mathbf{P}) &= \{x \in X(\mathbf{P}) \mid \text{there is no } y \in X(\mathbf{P})\text{: } y \to x \text{ and } y \text{ left-covers } x \text{ in } \mathcal{M}(\mathbf{P})\}. \end{aligned}$$

Note that $UC_{Fish}(\mathbf{P}) \subseteq UC_{Gill}(\mathbf{P})$. A useful alternative characterization of UC_{Gill} is given by the following "two-step" principle (see, e.g., Duggan 2013, Proposition 12(ii)): $x \in UC_{Gill}(\mathbf{P})$ if and only if for all $y \in X(\mathbf{P}) \setminus \{x\}$, $Margin_{\mathbf{P}}(x, y) \geq 0$ or there is a $z \in X(\mathbf{P})$ such that $Margin_{\mathbf{P}}(x, z) \geq 0$ and $Margin_{\mathbf{P}}(z, y) > 0$.

Before turning to the core criteria, for a proof that the Uncovered Set satisfies Pareto, see Duggan 2013, Proposition 52.

Core criteria For the satisfaction of the Condorcet criterion under various definitions of the Uncovered Set, see Duggan 2013, Propositions 4, 5, and 13. The Condorcet loser and monotonicity criteria are also straightforward to check. That UC_{Gill} satisfies independence of clones can be seen using the two-step characterization above. However, UC_{Fish} does not, as shown by the following profile $\mathbf{P}$:

1	1	1	1	1	1
e	d	f	b	b	b
d	f	e	e	d	f
f	e	d	d	f	e
b	b	b	f	e	d

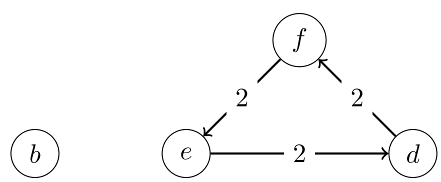

Observe that $\{d, e, f\}$ is a set of clones in $\mathbf{P}$, $UC_{Fish}(\mathbf{P}) = \{b\}$, but $UC_{Fish}(\mathbf{P}_{-f}) = \{b, e\}$, so UC_{Fish} does not satisfy the condition that clone choice is independent of clones (recall Definition 4.9). Finally, for narrowing, see Section 4.4.

New criteria The method UC_{Fish} satisfies stability for winners and hence immunity to spoilers, since if $Margin_{\mathbf{P}}(a, b) > 0$, then b does not left-cover a. However, it violates strong stability for winners, as shown by the profile $\mathbf{P}$ above where $UC_{Fish}(\mathbf{P}_{-b}) = \{d, e, f\}$, $Margin_{\mathbf{P}}(e, b) \geq 0$, but $UC_{Fish}(\mathbf{P}) = \{b\}$, and hence it violates amalgamation. By contrast, UC_{Gill} satisfies amalgamation, as one can easily see using the two-step characterization above. For rejectability, see Proposition 5.13.

Other criteria For the satisfaction of the Smith criterion under all standard definitions of the Uncovered Set, see Duggan 2013, Propositions 4, 5, and 14. For the satisfaction of ISDA, suppose $x \in X(\mathbf{P}) \setminus GETCHA(\mathbf{P})$ and $UC \in \{UC_{Fish}, UC_{Gill}\}$. To see that $UC(\mathbf{P}_{-x}) \subseteq UC(\mathbf{P})$, suppose $a \in UC(\mathbf{P}_{-x})$. By the Smith criterion for UC and ISDA for GETCHA, $UC(\mathbf{P}_{-x}) \subseteq GETCHA(\mathbf{P}_{-x}) = GETCHA(\mathbf{P})$, so together $a \in UC(\mathbf{P}_{-x})$ and $x \in X(\mathbf{P}) \setminus GETCHA(\mathbf{P})$ imply $Margin_{\mathbf{P}}(a, x) > 0$, so $a \in UC(\mathbf{P})$ by stability

[37] Miller's (1980) definition uses the right-sided version: y *right-covers* x in $\mathcal{M}$ if for all z, if $x \to z$, then $y \to z$. If $\to$ is weakly complete, then left-covering and right-covering are equivalent.

for winners. For $UC(\mathbf{P}) \subseteq UC(\mathbf{P}_{-x})$, if $\to$ in $\mathcal{M}(\mathbf{P})$ is not weakly complete, then $GETCHA(\mathbf{P}) = X(\mathbf{P})$, so there is no $x \in X(\mathbf{P}) \setminus GETCHA(\mathbf{P})$. Hence we may suppose that $\to$ in $\mathcal{M}(\mathbf{P})$ is weakly complete, in which case the various definitions of the Uncovered Set are equivalent to each other and to the following: $a \in UC(\mathbf{P})$ if and only if for all $b \in X(\mathbf{P}) \setminus \{a\}$, $a \to b$ or there is a $c \in X(\mathbf{P})$ such that $a \to c \to b$. Now suppose $a \in UC(\mathbf{P})$. To show $a \in UC(\mathbf{P}_{-x})$, we must show that for any $b \in X(\mathbf{P}_{-x}) \setminus \{a\}$, we have $a \to b$ or there is a $c \in X(\mathbf{P}_{-x})$ such that $a \to c \to b$. Since $a \in UC(\mathbf{P})$, either $a \to b$ or there is a $c \in X(\mathbf{P})$ such that $a \to c \to b$. If $a \to b$, we are done, so suppose it is not the case that $a \to b$, but instead there is a $c \in X(\mathbf{P})$ such that $a \to c \to b$. Since it is not the case that $a \to b$, we have $b \to a$ by the weak completeness of $\to$ in $\mathcal{M}(\mathbf{P})$. Then since $a \in UC(\mathbf{P}) \subseteq GETCHA(\mathbf{P})$, it follows from $c \to b \to a$ that $c \in GETCHA(\mathbf{P})$, so $c \neq x$. Hence $c \in X(\mathbf{P}_{-x})$ and $a \to c \to b$, so we are done. For reversal symmetry, the argument is the same as we gave for GETCHA in Appendix C.5. Uncovered Set violates single-voter and asymptotic resolvability for $k \geq 3$ by the same reasoning as for GETCHA (i.e., Uncovered Set selects a unique winner only if there is a Condorcet winner). For the failure of positive and negative involvement under all standard definitions of the Uncovered Set, see Pérez 2001, § 4.1.

C.8 Ranked Choice

Ranked Choice iteratively removes all candidates with the fewest number of voters who rank them first, until there is a candidate with a majority of voters ranking them first. If there are two or more candidates with the fewest number of voters ranking them first, then remove all such candidates. If, at some stage of the removal process, all remaining candidates have the same number of voters who rank them first (so all candidates would be removed), then all remaining candidates are selected as winners. For a profile $\mathbf{P}$, let $RC(\mathbf{P})$ be the set of Ranked Choice winners. See Wang et al. 2019 and Freeman et al. 2014 for computational and axiomatic properties of Ranked Choice.

Core criteria It is well known that Ranked Choice violates the Condorcet winner and monotonicity criteria but satisfies the Condorcet loser and independence of clones criteria (see Felsenthal 2012). Narrowing follows from asymptotic resolvability below.

New criteria Example 1.2 shows that Ranked Choice violates immunity to spoilers and hence the stronger criteria of stability for winners and amalgamation. That Ranked Choice satisfies rejectability follows from Lemma 6.18 given that it satisfies resolvability and clearly homogeneity.

Other criteria The failure of the Smith and ISDA criteria follows from the failure of the Condorcet winner criterion. For the failure of reversal symmetry, see Felsenthal 2012 (under "preference inversion" for "Alternative vote"). For the satisfaction of single-voter resolvability, see Tideman 1987 (where Ranked Choice is called "Alternative vote"). That asymptotic resolvability holds follows from the fact that for any number of candidates, the proportion of profiles in which there is a tie in the number of first place votes for two candidates goes to 0. For the failure of negative involvement, see Fishburn and Brams 1983 (under the "no show paradox"). For positive involvement, suppose $x \in RC(\mathbf{P})$, so x is not eliminated at any stage of the iteration procedure starting from $\mathbf{P}$. Then where $\mathbf{P}'$ is a one-voter profile whose voter ranks x in first place, clearly x is not eliminated at any stage of the iteration procedure starting from $\mathbf{P}+\mathbf{P}'$, so $x \in RC(\mathbf{P}+\mathbf{P}')$.

C.9 Plurality

The Plurality score of a candidate is the number of voters who rank that candidate in first place. The Plurality voting method selects as winners all candidates whose Plurality scores are maximal. The problems with Plurality Rule are well known (see Laslier 2012).

Core criteria It is well known that Plurality violates the Condorcet winner and Condorcet loser criteria (see Felsenthal 2012). It is clear that Plurality satisfies monotonicity. Example 1.1 shows that Plurality does not satisfy independence of clones. Narrowing follows from asymptotic resolvability below.

New criteria Example 1.1 shows that Plurality violates immunity to spoilers and hence the stronger criteria of stability for winners and amalgamation. That Plurality satisfies rejectability follows from Lemma 6.18 given that it satisfies resolvability and clearly homogeneity.

Other criteria The failure of the Smith and ISDA criteria follows from the failure of the Condorcet winner criterion. For the failure of reversal symmetry, see Felsenthal 2012. The satisfaction of single-voter and asymptotic resolvability and positive and negative involvement is obvious.

D Frequency of Irresoluteness

The graphs in this section show the frequency with which Split Cycle and several other voting methods select more than one winner, as well as the sizes of the winning sets, before tiebreaking. We evaluated these voting methods on profiles with different numbers of candidates and voters.

First, comparing Split Cycle, Beat Path, and GETCHA, we ran three simulations using different probabilistic models to generate the profiles, with the results shown in Figures 5, 6, and 7. The probabilistic models are explained in the captions of the figures. The impartial culture model (Figure 5) can be viewed as a "worst case scenario" for irresoluteness (cf. Tsetlin et al. 2003), so we expect that in practice the frequency of multiple winners will be substantially lower. The graphs on the left show the frequency of multiple winners for Split Cycle (the blue line), Beat Path (the green line), and GETCHA (the red line) as the number of candidates ranges over the set $\{5, 7, 10, 20, 30\}$ and the number of voters ranges over the set $\{5, 9, 25, 55, 101, 501, 1001, 5001\}$. On the right we use the standard box plot representation of the quartiles of the sizes of the winning sets, when there are multiple winners, where the blue boxes on the left are for Split Cycle, the green boxes in the middle are for Beat Path, and the red boxes on the right are for GETCHA. The black dots outside of the boxes are the "outliers." For each choice of a number of candidates and a number of voters, we generated 10,000 profiles and counted the number of times Split Cycle, Beat Path, and GETCHA selected more than one winner. For instance, the graphs on the left in Figure 5 show that with 10 candidates and 5,001 voters, GETCHA selects more than one winner in about 50% of the profiles, Split Cycle selects more than one winner in about 20% of the profiles, and Beat Path selects more than one winner in about 2% of the profiles. The box plots show that when there are multiple winners, there are very few for Split Cycle and Beat Path, whereas there are many for GETCHA.

Second, we compared Ranked Pairs to Split Cycle, Beat Path, and GETCHA. Due to the computational difficulty of computing Ranked Pairs winning sets (see Brill and Fischer 2012; Wang et al. 2019), we only tested profiles for 5, 6, or 7 candidates and 55, 101, 501, or 1001 voters. In Figure 8, we only show the data for the worst case scenario of the impartial culture model. The results for Ranked Pairs for the other probabilistic models are similar to those for Beat Path, mentioned above.

Third, in Figure 9 we include a comparison with the Condorcet consistent voting method Copeland, also under the impartial culture model. We chose Copeland because like Split Cycle, GETCHA, and GOCHA, Copeland does not satisfy the resolvability criteria of Definitions 6.15 and 6.19, yet Copeland is one of the most discriminating of all C1 voting methods (recall Section 5.3 for the definition of C1, and see Brandt and Seedig 2014 on the discriminating power of different C1 methods). Unlike in the cases of Split Cycle, Beat Path, and Ranked Pairs, in the case of Copeland (as in the case of GETCHA), the proportion of profiles with multiple winners seems largely insensitive to the number of voters. Unlike in the cases of Split Cycle, Beat Path, Ranked Pairs, and GETCHA, in the case of Copeland, the proportion of profiles with multiple winners also seems largely insensitive to the number of candidates.

Fourth, in Figure 10 we compare Split Cycle, Copeland, and the Uncovered Set (since all sampled profiles have an odd number of voters, we do not distinguish between different versions of the Uncovered Set here). It follows from a result of Moulin (1986, Theorem 1) that for any C1 voting method F satisfying neutrality and amalgamation and profile $\mathbf{P}$ with an odd number of voters, $UC(\mathbf{P}) \subseteq F(\mathbf{P})$ (for an analogous result for an even number of voters, using a definition of the Uncovered Set that satisfies amalgamation, see Peris and Subiza 1999, Theorem 1). Thus, by comparing Split Cycle to the Uncovered Set, we are comparing Split Cycle to the most discriminating of all C1 methods satisfying amalgamation, and by comparing Split Cycle to Copeland, we are comparing Split Cycle to one of the most discriminating of all C1 methods.

We also compared Split Cycle, Beat Path, GETCHA, Copeland, and the Uncovered Set (the Gillies version) on 315 profiles using real election data, with different numbers of candidates and voters, available at www.preflib.org (Mattei and Walsh 2013). The results are shown in Figure 11. The number of candidates range from 3 to 242, and the number of voters range from 4 to 14,081.

Interestingly, on all 315 profiles from the preflib database, Split Cycle and Beat Path produced the same set of winners. To investigate further the probability that Split Cycle and Beat Path will agree, we used the same three probability models used for Figures 5, 6, and 7. The results are shown in Figure 13. For two of the three probability models, as the number of candidates increases, the probability that Split Cycle and Beat Path disagree also increases.

Finally, Figures 14 and 15 estimate how much the greater resoluteness of Beat Path relative to Split Cycle can be attributed to Beat Path violating the stability for winners criterion (Section 5.1). Using two probability models, we sampled profiles and checked for each profile, first, if $BP(\mathbf{P}) \subsetneq SC(\mathbf{P})$, and if so, whether there is some candidate $a \in SC(\mathbf{P}) \setminus BP(\mathbf{P})$ and some other candidate $b \in X(\mathbf{P})$ such that $\mathbf{P}$, a, and b witness a violation of stability for winners as stated in Definition 5.3.

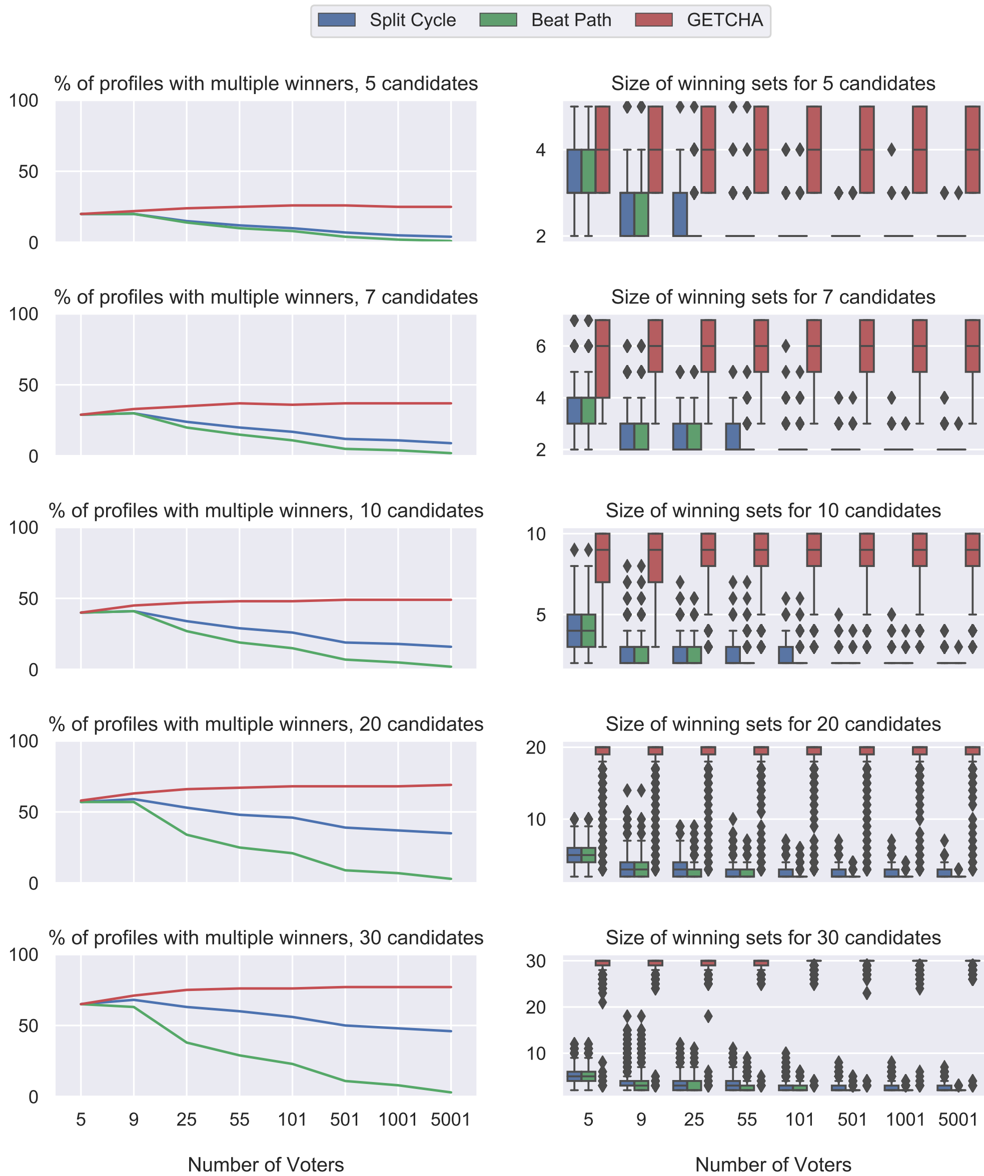

Figure 5: The profiles were generated using the impartial culture probability model. This assumes that each linear order has the same probability of being submitted by a voter (so the probability of a linear order is $\frac{1}{n!}$ where n is the number of candidates), and the voters' choices of rankings are independent.

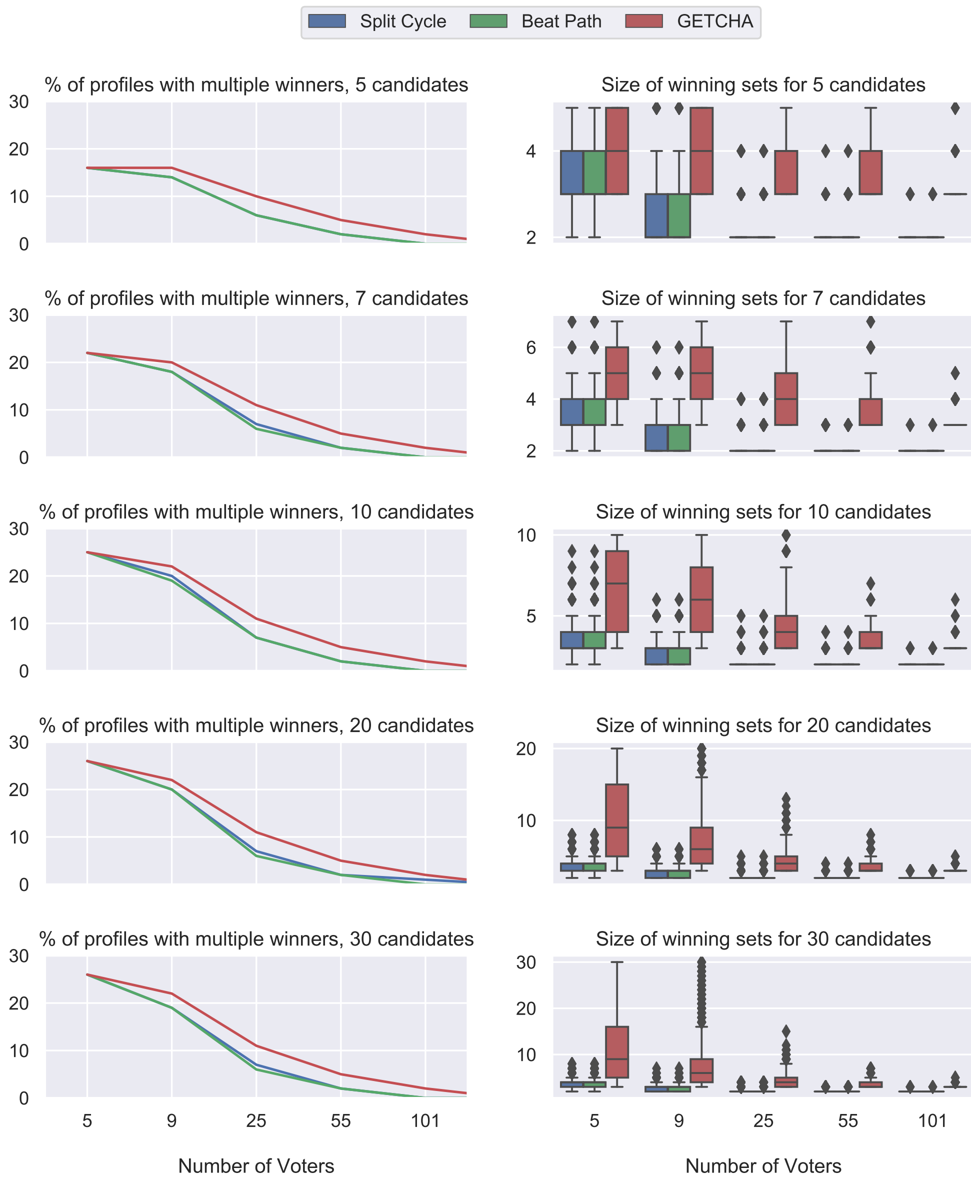

Figure 6: The profiles were generated using the Mallows-ϕ model (see Marden 1995). The main idea is to fix a reference linear ordering of the candidates and to assign to each voter a ranking that is "close" to this reference ranking. Closeness to the reference ranking is defined using the Kendall-tau distance between rankings, depending on a *dispersion* parameter. Setting the dispersion parameter to 0 means that every voter is assigned the reference ranking, and setting it to 1.0 is equivalent to the impartial culture model. We used a dispersion parameter of 0.8 (the same used in Brandt et al. 2019, 2020). The reason that only the voting sizes 5, 9, 25, 55 and 101 are shown is that there were no profiles with 501, 1001, and 5001 voters that had more than a single Split Cycle, Beat Path, or GETCHA winner.

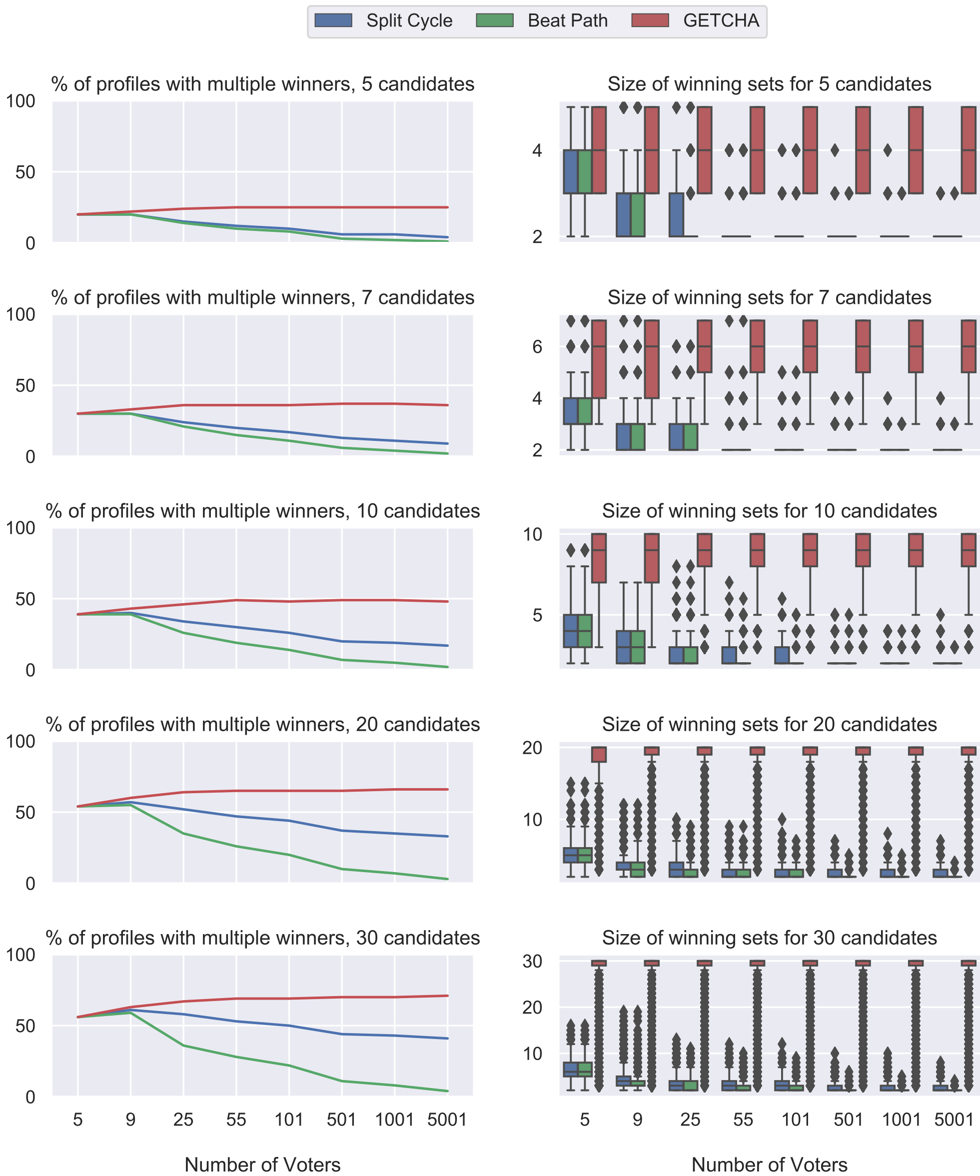

Figure 7: The profiles were generated using the Mallows-ϕ model with *two* reference rankings, which are the reverse of each other. The set of voters is divided into two groups, each associated with one of the reference rankings. Each voter is equally likely to be assigned to either of the two groups. Within a group, the dispersion parameter for selecting a ranking close to the reference ranking of that group is 0.8.

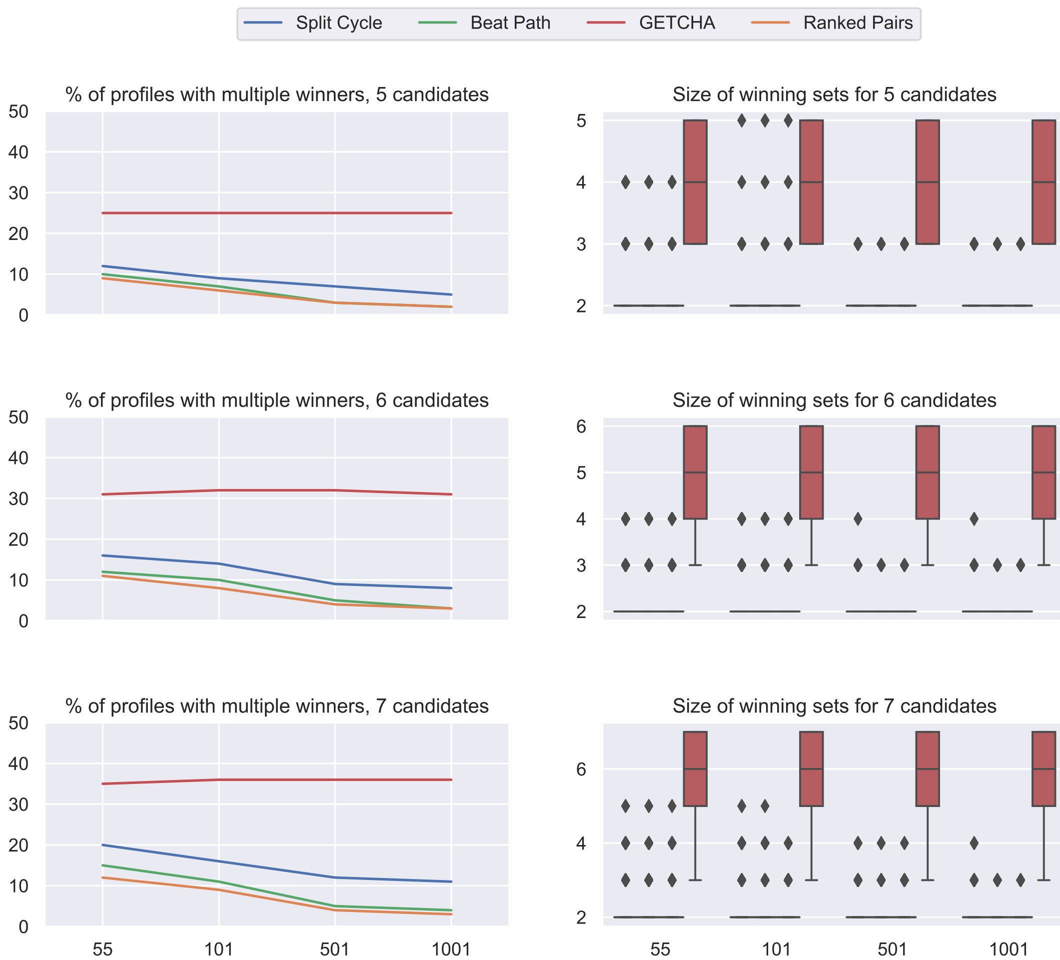

Figure 8: The profiles were generated using the impartial culture probability model. Note that for Split Cycle, Beat Path, and Ranked Pairs, almost all winning sets with multiple winners were of size two.

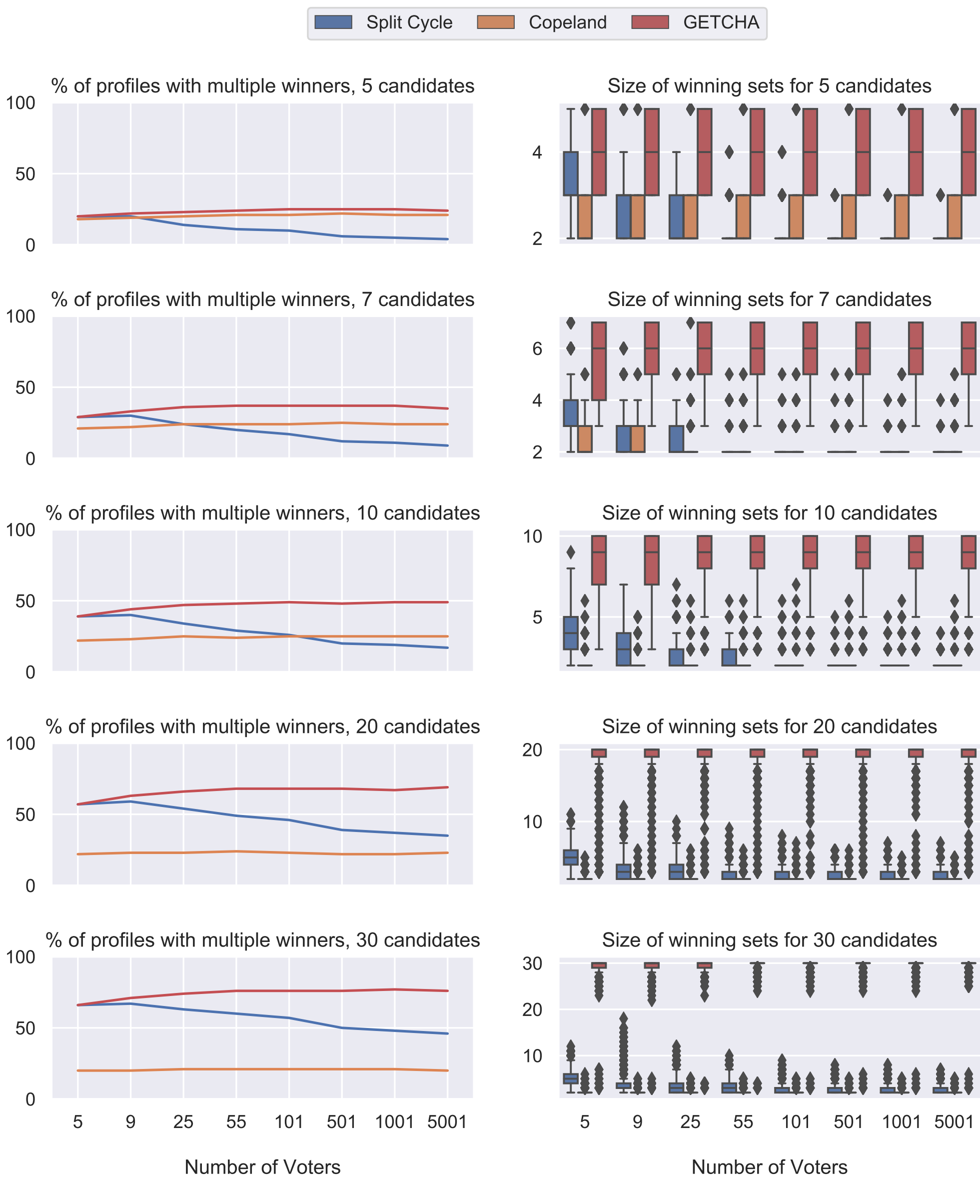

Figure 9: The profiles were generated using the impartial culture probability model.

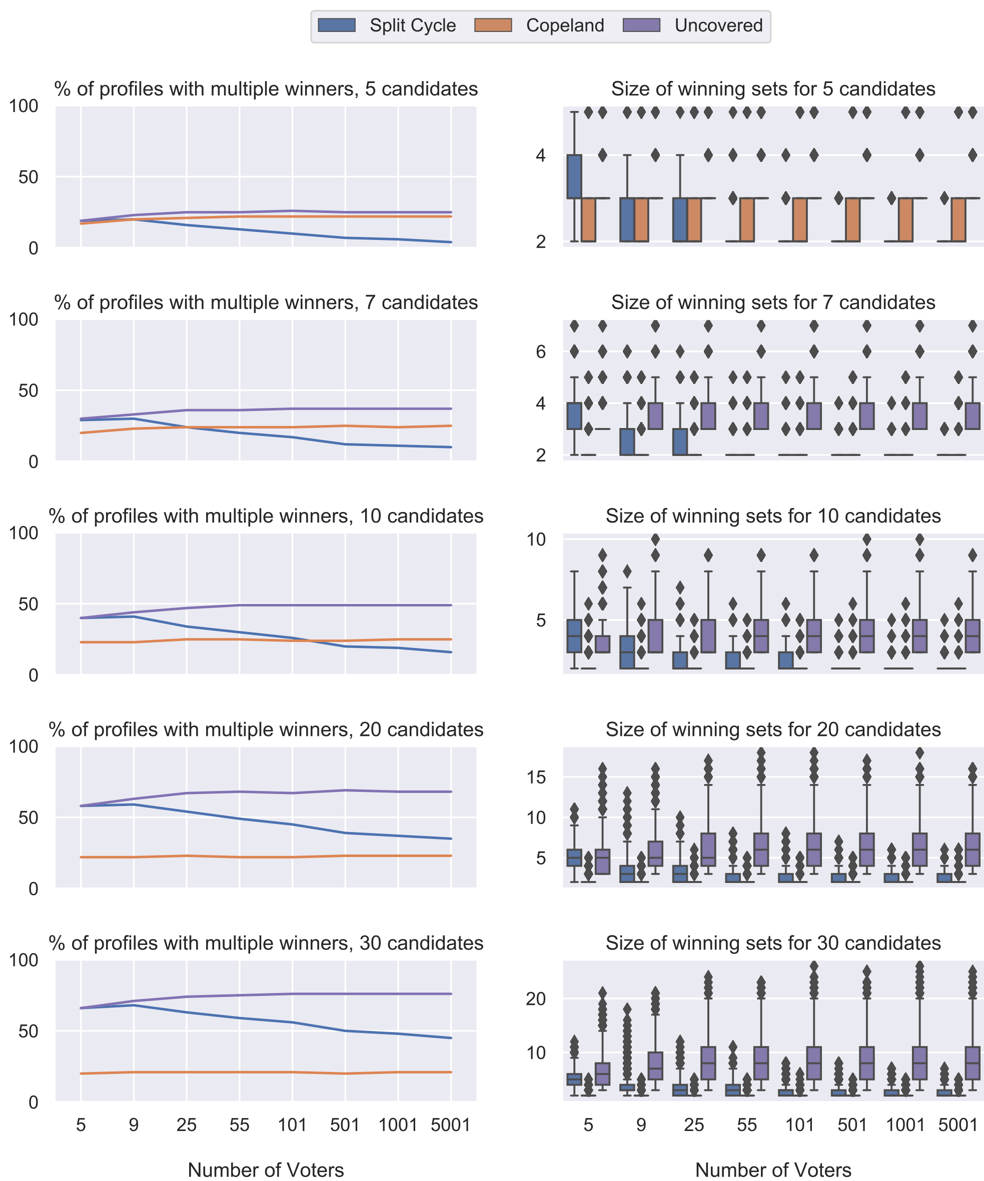

Figure 10: The profiles were generated using the impartial culture probability model.

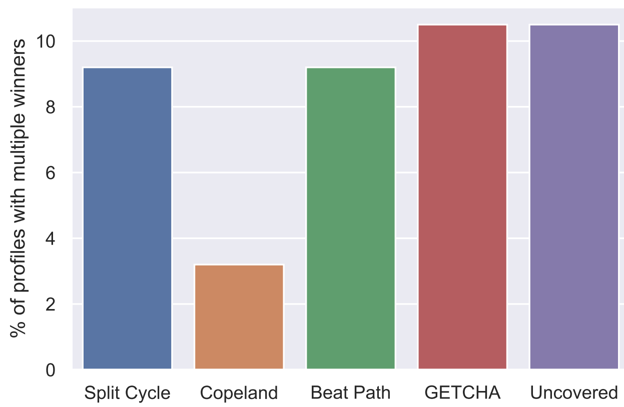

Figure 11: The percentage of profiles with multiple winners (for the Uncovered Set, we use the Gillies version) out of 315 different profiles using real election data available at www.preflib.org.

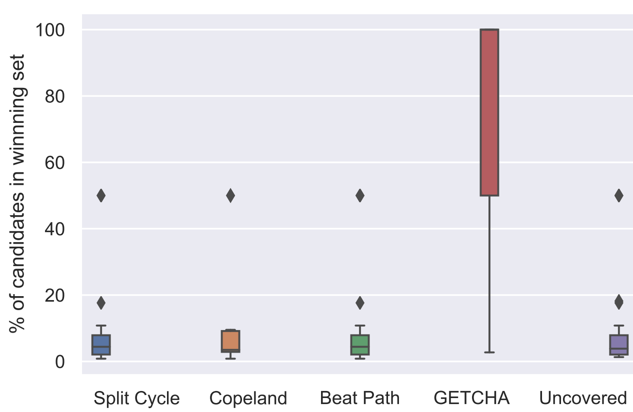

Figure 12: Box plot showing *for the profiles with multiple winners* from the preflib database used for Figure 11, the percentage of candidates in the winning set. Note that the percentages would be lower if we had also included profiles with unique winners.

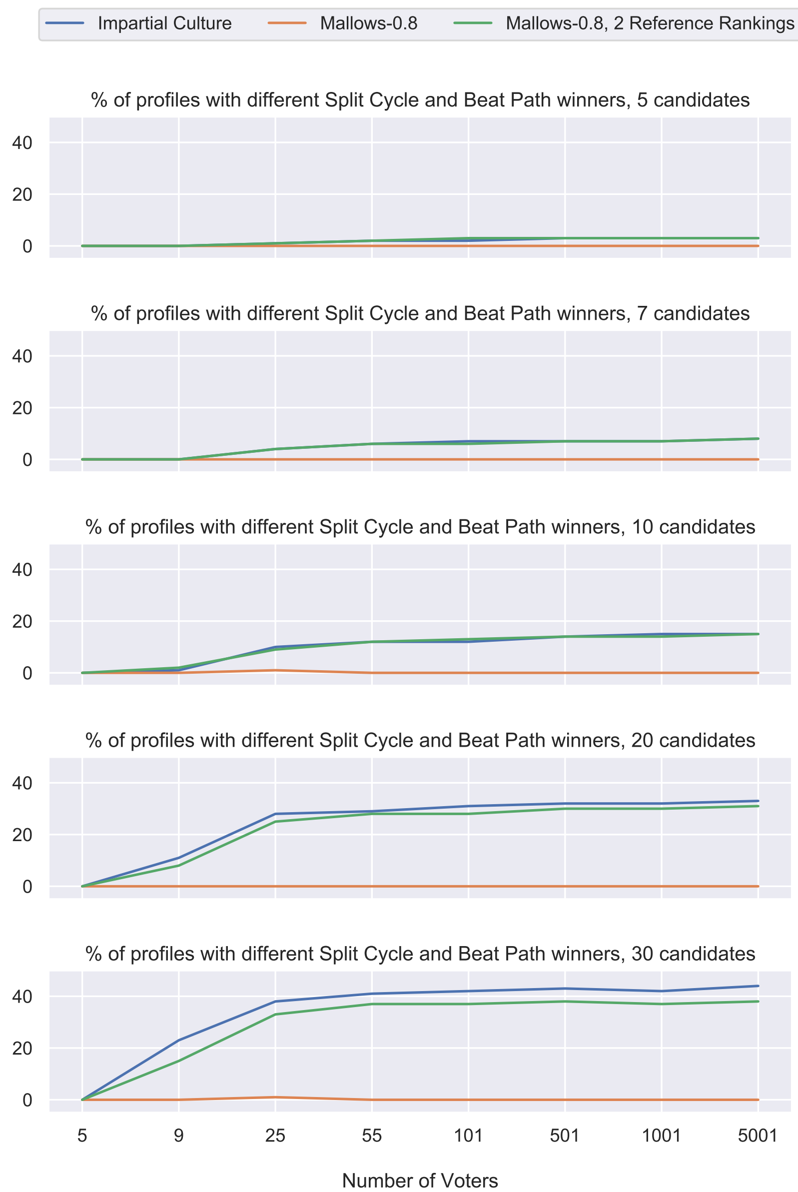

Figure 13: The percentage of profiles with different Split Cycle and Beat Path winners for different numbers of candidates and voters. The profiles were generated using the following three probability models: impartial culture, Mallows-0.8, and Mallows-0.8 with 2 reference rankings.

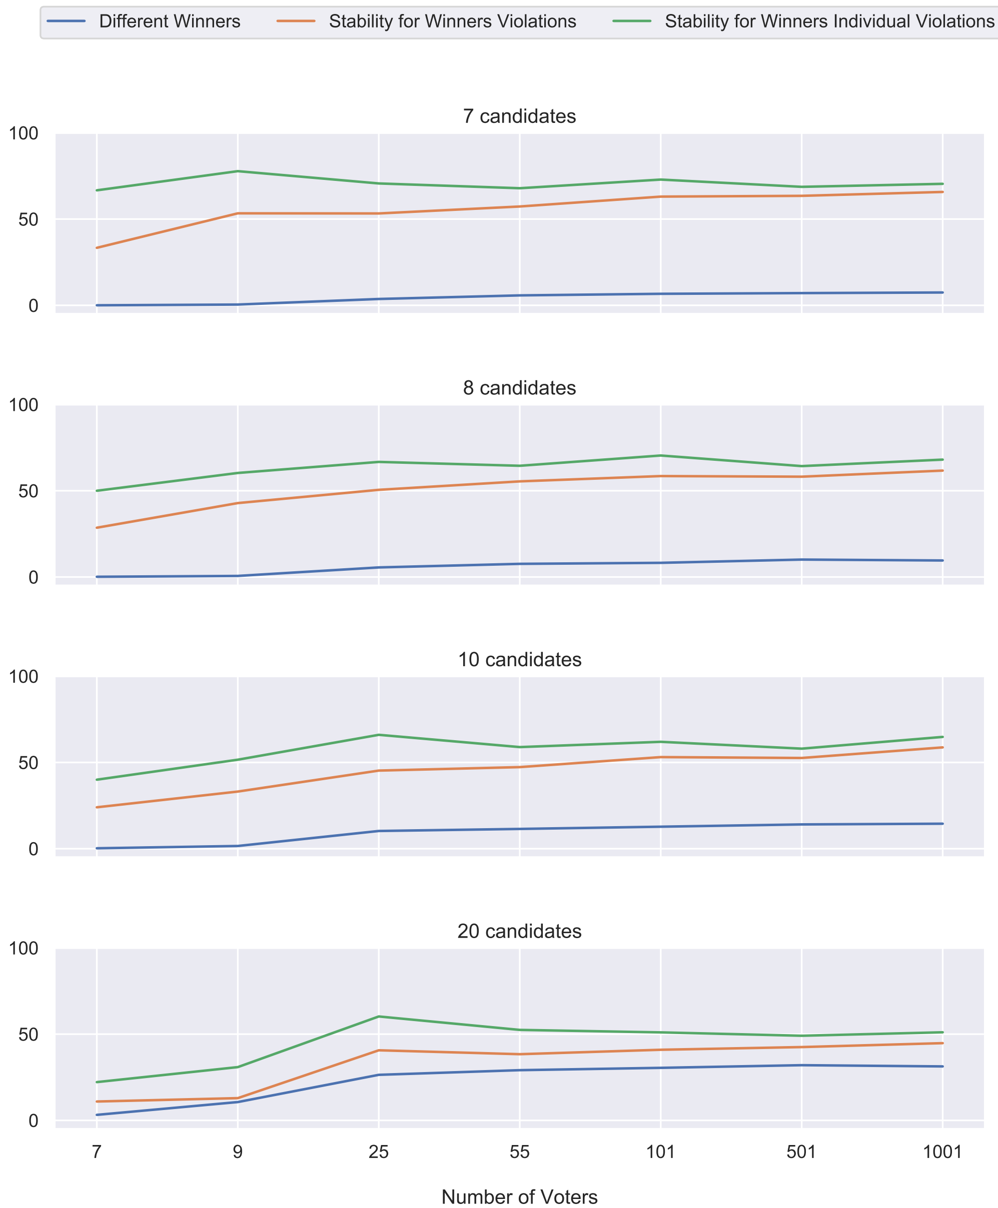

Figure 14: The blue line is the proportion of 10,000 profiles generated using the impartial culture model that have different Split Cycle and Beat Path winners. The orange line is the proportion of profiles, among those with different Split Cycle and Beat Path winners, that witness a violation of the stability for winners criterion for Beat Path, i.e., there exist $a, b \in X(\mathbf{P})$ such that $a \in BP(\mathbf{P}_{-b})$, $Margin_{\mathbf{P}}(a, b) > 0$, and $a \notin BP(\mathbf{P})$. An *individual violation* of stability for winners is a triple $(\mathbf{P}, a, b)$ of a profile and two candidates as just described. Note that there may be multiple individual violations of stability for winners based on the same profile. The green line is the number of individual violations of stability for winners for Beat Path divided by the number of profiles that have different Split Cycle and Beat Path winners.

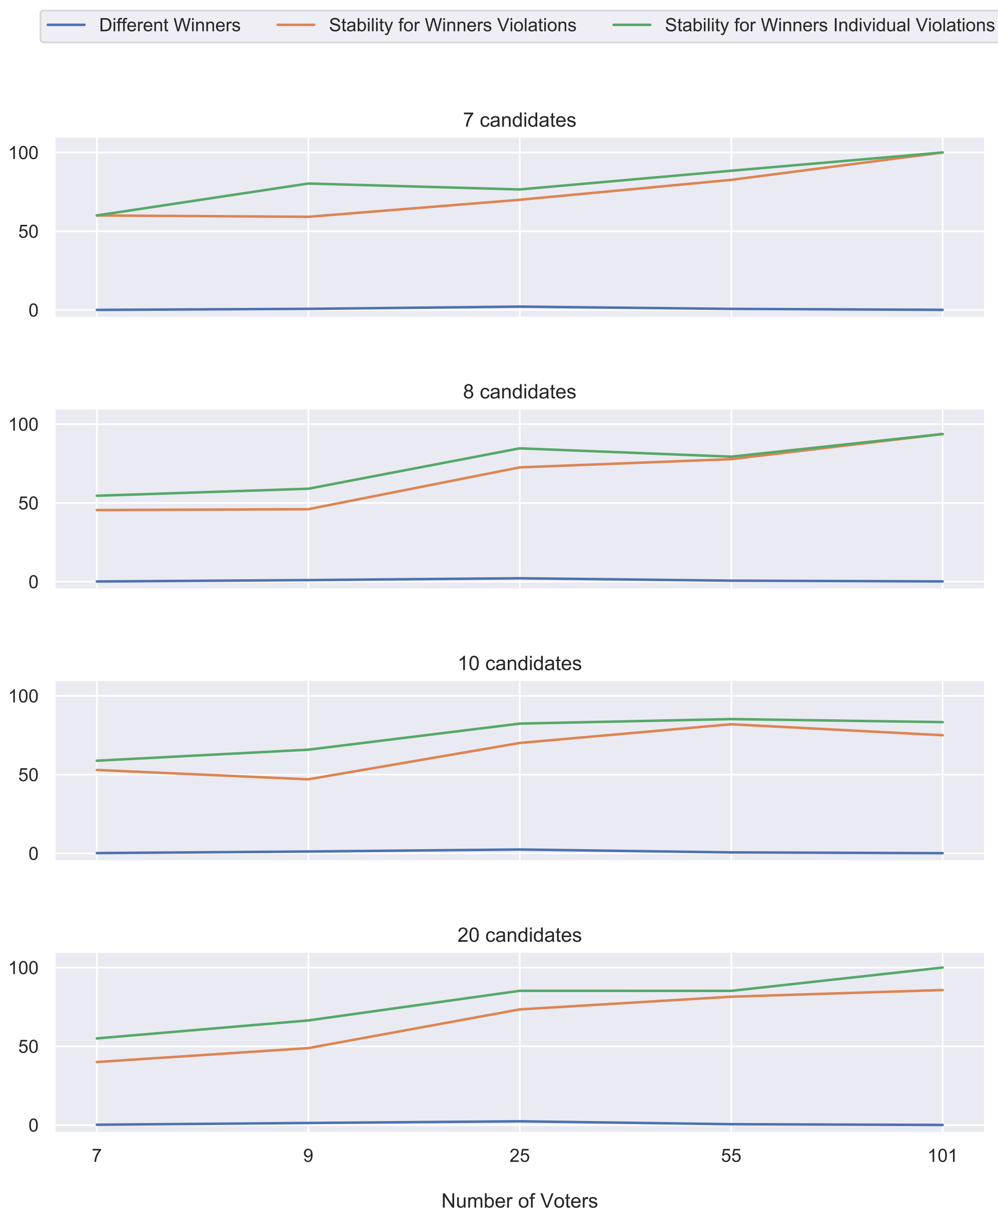

Figure 15: Same graphs as in Figure 14 except that the profiles were generated using a Mallows-0.8 probability model (see the caption to Figure 6). Note that the blue line is in fact slightly above 0 in each of the graphs.

References

Kenneth J. Arrow. *Social Choice and Individual Values*. John Wiley & Sons, Inc., New York, 2nd edition, 1963.

Nick Baigent. Twitching weak dictators. *Journal of Economics*, 47(4):407–411, 1987. doi:10.1007/bf01229471.

Michel Balinski and Rida Laraki. *Majority Judgement: Measuring, Ranking and Electing*. MIT Press, Boston, 2010. doi:10.7551/mitpress/9780262015134.001.0001.

Colin E. Bell. A random voting graph almost surely has a Hamiltonian cycle when the number of alternatives is large. *Econometrica*, 49(6):597–1603, 1981.

Georges Bordes. On the possibility of reasonable consistent majoritarian choice: Some positive results. *Journal of Economic Theory*, 31(1):122–132, 1983. doi:10.1016/0022-0531(83)90024-8.

Steven J. Brams and Peter C. Fishburn. Voting procedures. In Kenneth J. Arrow, Amartya K. Sen, and Kotaro Suzumura, editors, *Handbook of Social Choice and Welfare*, volume 1, pages 173–236. North-Holland, Amsterdam, 2002. doi:10.1016/s1574-0110(02)80008-x.

Felix Brandt. Rolling the dice: Recent results in probabilistic social choice. In Ulle Endriss, editor, *Trends in Computational Social Choice*, pages 3–26. AI Access, 2017.

Felix Brandt and Hans Georg Seedig. On the discriminative power of tournament solutions. In M. Lübbecke, A. Koster, P. Letmathe, R. Madlener, B. Peis, and G. Walther, editors, *Operations Research Proceedings 2014*, pages 53–58, Cham, 2014. Springer. doi:doi.org/10.1007/978-3-319-28697-6_8.

Felix Brandt, Markus Brill, Hans Georg Seedig, and Warut Suksompong. On the structure of stable tournament solutions. *Economic Theory*, 65(2):483–507, 2018. doi:10.1007/s00199-016-1024-x.

Felix Brandt, Johannes Hofbauer, and Martin Strobel. Exploring the no-show paradox for Condorcet extensions using Ehrhart theory and computer simulations. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS '19)*, pages 520–528. International Foundation for Autonomous Agents and MultiAgent Systems, 2019.

Felix Brandt, Christian Geist, and Martin Strobel. Analyzing the practical relevance of the Condorcet loser paradox and the agenda contraction paradox. In M. Diss and V. Merlin, editors, *Evaluating Voting Systems with Probability Models: Essays by and in Honor of William Gehrlein and Dominique Lepelley*, Berlin, 2020. Springer.

Markus Brill and Felix Fischer. The price of neutrality for the ranked pairs method. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence (AAAI-12)*, pages 1299–1305. AAAI Press, 2012.

Donald E. Campbell and Jerry S. Kelly. Weak independence and veto power. *Economics Letters*, 66(2): 183–189, 2000. doi:10.1016/s0165-1765(99)00209-8.

Vincent Conitzer and Tuomas Sandholm. Common voting rules as maximum likelihood estimators. In *Proceedings of the Twenty-First Conference on Uncertainty in Artificial Intelligence (UAI-05)*, pages 145–152, Arlington, Virginia, 2005. AUAI Press.

A. H. Copeland. A 'reasonable' social welfare function. Notes from a seminar on applications of mathematics to the social sciences, University of Michigan, 1951.

Richard B. Darlington. Minimax is the best electoral system after all. arXiv:1606.04371, 2016.

Bernard Debord. Caractérisation des matrices des préférences nettes et méthodes d'agrégation associées. *Mathématiques et sciences humaines*, 97:5–17, 1987.

Frank DeMeyer and Charles R. Plott. The probability of a cyclical majority. *Econometrica*, 38(2):345–354, 1970.

Conal Duddy. Condorcet's principle and the strong no-show paradoxes. *Theory and Decision*, 77:275–285, 2014. doi:10.1007/s11238-013-9401-4.

John Duggan. Uncovered sets. *Social Choice and Welfare*, 41(3):489–535, 2013. doi:10.1007/s00355-012-0696-9.

Bhaskar Dutta, Matthew O. Jackson, and Michel Le Breton. Strategic candidacy and voting procedures. *Econometrica*, 69(4):1013–1037, 2001. doi:10.1111/1468-0262.00228.

Lars Ehlers and John A. Weymark. Candidate stability and nonbinary social choice. *Economic Theory*, 22 (2):233–243, 2003. doi:10.1007/s00199-002-0279-6.

Hülya Eraslan and Andrew McLellan. Candidate stability and nonbinary social choice. *Journal of Economic Theory*, 117(1):29–54, 2004. doi:10.1016/j.jet.2003.09.005.

Dan S. Felsenthal. Review of paradoxes afflicting procedures for electing a single candidate. In *Electoral Systems: Paradoxes, Assumptions and Procedures*, pages 19–92, Berlin, 2012. Springer. doi:10.1007/978-3-642-20441-8_3.

Dan S. Felsenthal and Hannu Nurmi. Two types of participation failure under nine voting methods in variable electorates. *Public Choice*, 168:115–135, 2016. doi:10.1007/s11127-016-0352-5.

Dan S. Felsenthal and Nicolaus Tideman. Varieties of failure of monotonicity and participation under five voting methods. *Theory and Decision*, 75:59–77, 2013. doi:10.1007/s11238-012-9306-7.

Mark Fey. Choosing from a large tournament. *Social Choice and Welfare*, 31:301–309, 2008. doi:10.1007/s00355-007-0279-3.

Peter C. Fishburn. Condorcet social choice functions. *SIAM Journal on Applied Mathematics*, 33(3):469–489, 1977. doi:10.1137/0133030.

Peter C. Fishburn and Steven J. Brams. Paradoxes of preferential voting. *Mathematics Magazine*, 56(4): 207–214, 1983. doi:10.2307/2689808.

Rupert Freeman, Markus Brill, and Vincent Conitzer. On the axiomatic characterization of runoff voting rules. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence (AAAI-14)*, pages 675–681. AAAI Press, 2014.

Donald B. Gillies. Solutions to general non-zero-sum games. In A. W. Tucker and R. D. Luce, editors, *Contributions to the Theory of Games*. Princeton University Press, 1959.

Matthew Harrison-Trainor. An analysis of random elections with large numbers of voters. arXiv:2009.02979, 2020.

Jobst Heitzig. Social choice under incomplete, cyclic preferences: Majority/minority-based rules, and composition-consistency. arXiv:0201285, 2002.

Jobst Heitzig. Examples with 4 options for immune methods, 2004. URL http://lists.electorama.com/pipermail/election-methods-electorama.com//2004-May/078166.html.

Wesley H. Holliday and Eric Pacuit. Axioms for defeat in democratic elections. arXiv:2008.08451, 2020.

José L. Jimeno, Joaquín Pérez, and Estefanía García. An extension of the Moulin No Show Paradox for voting correspondences. *Social Choice and Welfare*, 33(3):343–359, 2009. doi:10.1007/s00355-008-0360-6.

Anna Purna Kambhampaty. New York City voters just adopted Ranked-Choice Voting in elections. Here's how it works. *Time Magazine*, 2019. November 6, 2019 issue.

Laura Kasper, Hans Peters, and Dries Vermeulen. Condorcet Consistency and the strong no show paradoxes. *Mathematical Social Sciences*, 99:36–42, 2019. doi:10.1016/j.mathsocsci.2019.03.002.

Jerry S. Kelly. *Social Choice Theory: An Introduction*. Springer, Berlin, 1988. doi:10.1007/978-3-662-09925-4.

Gerald H. Kramer. A dynamical model of political equilibrium. *Journal of Economic Theory*, 16(2):310–334, 1977. doi:10.1016/0022-0531(77)90011-4.

Claude Lamboray. A prudent characterization of the ranked pairs rule. *Social Choice and Welfare*, 32:129–155, 2008. doi:10.1007/s00355-008-0319-7.

Jean-François Laslier. *Tournament Solutions and Majority Voting*. Springer, Berlin, 1997.

Jean-François Laslier. And the loser is...plurality voting. In Dan S. Felsenthal and Moshé Machover, editors, *Electoral Systems: Paradoxes, Assumptions and Procedures*, pages 327–351, Berlin, 2012. Springer. doi:10.1007/978-3-642-20441-8_13.

Christopher S. P. Magee. Third-party candidates and the 2000 presidential election. *Social Science Quarterly*, 84(3):29–35, 2003. doi:10.1111/1540-6237.8403006.

John Marden. *Analyzing and Modeling Rank Data*. CRC Press, New York, 1995. doi:10.1201/b16552.

Eric Maskin and Amartya Sen. The rules of the game: A new electoral system. *The New York Review of Books*, 2017a. January 19, 2017 issue.

Eric Maskin and Amartya Sen. A better way to choose presidents. *The New York Review of Books*, 2017b. June 8, 2017 issue.

Nicholas Mattei and Toby Walsh. Preflib: A library of preference data. In *Proceedings of Third International Conference on Algorithmic Decision Theory*, pages 259–270. Springer, 2013. doi:10.1007/978-3-642-41575-3_20.

Nicholas R. Miller. A new solution set for tournaments and majority voting: Further graph-theoretical approaches to the theory of voting. *American Journal of Political Science*, 24(1):68–96, 1980. doi:10.2307/2110925.

Nicholas R. Miller. Reflections on Arrow's theorem and voting rules. *Public Choice*, 179:113–124, 2019. doi:10.1007/s11127-018-0524-6.

J. W. Moon and L. Moser. Almost all tournaments are irreducible. *Canadian Mathematical Bulletin*, 5(1): 61–65, 1962. doi:10.4153/CMB-1962-010-4.

Hervé Moulin. Choosing from a tournament. *Social Choice and Welfare*, 3(4):271–291, 1986. doi:10.1007/BF00292732.

Hervé Moulin. Condorcet's principle implies the no show paradox. *Journal of Economic Theory*, 45(1): 53–64, 1988. doi:10.1016/0022-0531(88)90253-0.

Michael Munie, Pingzhong Tang, and Yoav Shoham. A framework for the quantitative evaluation of voting rules. In *Proceedings of Logic, Game Theory, and Social Choice 6 (LGS6)*, 2009.

Roger B. Myerson. Axiomatic derivation of scoring rules without the ordering assumption. *Social Choice and Welfare*, 12(1):59–74, 1995. doi:10.1007/BF00182193.

Eric Pacuit. Voting methods. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Fall 2019 edition, 2019.

Elizabeth Maggie Penn. Alternate definitions of the uncovered set and their implications. *Social Choice and Welfare*, 27(1):83–87, 2006. doi:10.1007/s00355-006-0114-2.

Joaquín Pérez. The strong no show paradoxes are a common flaw in Condorcet voting correspondences. *Social Choice and Welfare*, 18(3):601–616, 2001. doi:10.1007/s003550000079.

Josep E. Peris and Begoña Subiza. Condorcet choice correspondences for weak tournaments. *Social Choice and Welfare*, 16(2):217–231, 1999. doi:10.1007/s003550050141.

Marcus Pivato. Voting rules as statistical estimators. *Social Choice and Welfare*, 40(2):581–630, 2013a. doi:10.1007/s00355-011-0619-1.

Marcus Pivato. Variable-population voting rules. *Journal of Mathematical Economics*, 49(3):210–221, 2013b. doi:10.1016/j.jmateco.2013.02.001.

Florenz Plassmann and T. Nicolaus Tideman. How frequently do different voting rules encounter voting paradoxes in three-candidate elections? *Social Choice and Welfare*, 42:31–75, 2014. doi:10.1007/s00355-013-0720-8.

Jeffrey T. Richelson. A comparative analysis of social choice functions, III. *Behavioral Science*, 23:169–176, 1978.

Carmelo Rodríguez-Álvarez. Candidate stability and voting correspondences. *Social Choice and Welfare*, 27 (3):545–570, 2006. doi:10.1007/s00355-006-0126-y.

Donald G. Saari. *Geometry of Voting*. Springer, Berlin, 1994. doi:10.1007/978-3-642-48644-9.

Donald G. Saari. *Basic Geometry of Voting*. Springer, Berlin, 1995. doi:10.1007/978-3-642-57748-2.

Donald G. Saari. Explaining all three-alternative voting outcomes. *Journal of Economic Theory*, 87(2): 313–355, 1999. doi:10.1006/jeth.1999.2541.

Christian Saile and Warut Suksompong. Robust bounds on choosing from large tournaments. *Social Choice and Welfare*, 54:87–110, 2020. doi:10.1007/s00355-019-01213-6.

M. Remzi Sanver and William S. Zwicker. Monotonicity properties and their adaptation to irresolute social choice rules. *Social Choice and Welfare*, 39(2/3):371–398, 2012. doi:10.1007/s00355-012-0654-6.

Markus Schulze. A new monotonic, clone-independent, reversal symmetric, and condorcet-consistent single-winner election method. *Social Choice and Welfare*, 36:267–303, 2011. doi:10.1007/s00355-010-0475-4.

Markus Schulze. The Schulze method of voting. arXiv:1804.02973, 2018.

Thomas Schwartz. *The Logic of Collective Choice*. Columbia University Press, New York, 1986. doi:10.7312/schw93758.

Alex Scott and Mark Fey. The minimal covering set in large tournaments. *Social Choice and Welfare*, 38: 1–9, 2012. doi:10.1007/s00355-010-0503-4.

Amartya Sen. Choice functions and revealed preference. *The Review of Economic Studies*, 38(3):307–317, 1971. doi:10.2307/2296384.

Amartya Sen. Internal consistency of choice. *Econometrica*, 61(3):495–521, 1993. doi:10.2307/2951715.

Amartya Sen. *Collective Choice and Social Welfare: An Expanded Edition*. Harvard University Press, Cambridge, Mass., 2017.

Paul B. Simpson. On defining areas of voter choice: Professor Tullock on stable voting. *The Quarterly Journal of Economics*, 83(3):478–490, 1969. doi:10.2307/1880533.

John H. Smith. Aggregation of preferences with variable electorate. *Econometrica*, 41(6):1027–1041, 1973. doi:10.2307/1914033.

Alan D. Taylor. *Social Choice and the Mathematics of Manipulation*. Cambridge University Press, Cambridge, 2005. doi:10.1017/cbo9780511614316.

T. Nicolaus Tideman. Independence of clones as a criterion for voting rules. *Social Choice and Welfare*, 4: 185–206, 1987. doi:10.1007/bf00433944.

Ilia Tsetlin, Michel Regenwetter, and Bernard Grofman. The impartial culture maximizes the probability of majority cycles. *Social Choice and Welfare*, 21:387–398, 2003. doi:10.1007/s00355-003-0269-z.

Jun Wang, Sujoy Sikdar, Tyler Shepherd Zhibing Zhao, Chunheng Jiang, and Lirong Xia. Practical algorithms for multi-stage voting rules with parallel universes tiebreaking. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence (AAAI-19)*. AAAI Press, 2019.

Wikipedia contributors. Spoiler effect — Wikipedia, the free encyclopedia, 2020a. URL `https://en.wikipedia.org/wiki/Spoiler_effect`. [Online; accessed 19-February-2020].

Wikipedia contributors. Condorcet method — Wikipedia, the free encyclopedia, 2020b. URL `https://en.wikipedia.org/wiki/Condorcet_method#Use_of_Condorcet_voting`. [Online; accessed 19-February-2020].

H. P. Young. Social choice scoring functions. *SIAM Journal on Applied Mathematics*, 28(4):313–355, 1975. doi:10.1137/0128067.

T. M. Zavist and T. Nicolaus Tideman. Complete independence of clones in the ranked pairs rule. *Social Choice and Welfare*, 6:167–173, 1989. doi:10.1007/bf00303170.

William S. Zwicker. Introduction to the theory of voting. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, pages 23–56. Cambridge University Press, New York, 2016. doi:10.1017/cbo9781107446984.003.