

Measuring Violations of Positive Involvement in Voting*

Wesley H. Holliday

University of California, Berkeley

wesholliday@berkeley.edu

Eric Pacuit

University of Maryland

epacuit@umd.edu

In the context of computational social choice, we study voting methods that assign a set of winners to each profile of voter preferences. A voting method satisfies the property of positive involvement (PI) if for any election in which a candidate x would be among the winners, adding another voter to the election who ranks x first does not cause x to lose. Surprisingly, a number of standard voting methods violate this natural property. In this paper, we investigate different ways of measuring the extent to which a voting method violates PI, using computer simulations. We consider the probability (under different probability models for preferences) of PI violations in randomly drawn profiles vs. profile-coalition pairs (involving coalitions of different sizes). We argue that in order to choose between a voting method that satisfies PI and one that does not, we should consider the probability of PI violation conditional on the voting methods choosing different winners. We should also relativize the probability of PI violation to what we call voter potency, the probability that a voter causes a candidate to lose. Although absolute frequencies of PI violations may be low, after this conditioning and relativization, we see that under certain voting methods that violate PI, much of a voter’s potency is turned against them—in particular, against their desire to see their favorite candidate elected.

1 Introduction

Voting provides a mechanism for resolving conflicts between the preferences of multiple agents in order to arrive at a group choice. Although traditional voting theory is largely motivated by the example of democratic political elections, the field of computational social choice [3, 4, 16] views voting theory as applicable to many other preference aggregation problems for multiagent systems.

One of the most basic ideas in voting is that an unequivocal increase in support for a candidate should not result in that candidate going from being a winner to being a loser. There are at least two ways to formalize this idea. First, there is the fixed-electorate axiom of *monotonicity*: if a candidate x is a winner given a preference profile \mathbf{P} , and \mathbf{P}' is obtained from \mathbf{P} by one voter moving x up in their ranking, then x should still be a winner given \mathbf{P}' . Second, there is the variable-electorate axiom of *positive involvement* [43, 41]: if a candidate x is a winner given \mathbf{P} , and \mathbf{P}^* is obtained from \mathbf{P} by adding a new voter who ranks x in first place, then x should still be a winner given \mathbf{P}^* . These axioms are logically independent. For example, Instant Runoff Voting (see Section 2.2) satisfies positive involvement but not monotonicity; and a number of well-known voting methods satisfy monotonicity but not positive involvement.

There are at least two reactions to a voting method’s violating monotonicity or positive involvement. One is that the method’s violating one of these axioms is a serious problem if and only if violations are sufficiently frequent (according to some probability model). Another reaction is that the method’s violating one of these axioms is a sign that the principle by which the voting method selects winners is fundamentally misconceived, which casts suspicion on its selection of winners in general, not just in those cases where the axiom is violated; but of course it is even worse for the method if in addition to having a misconceived principle for selecting winners, it frequently witnesses violations of the axiom.

*We thank the three anonymous reviewers for TARK for their helpful comments.

In this paper, we discuss different ways of measuring the extent to which a voting method violates positive involvement (in particular, the eight different ways shown in Figure 1), using computer simulations. Our main aim is to determine appropriate measures of positive involvement violation and identify the main parameters affecting these measures—e.g., whether they are affected by variants of voting methods, numbers of candidates and voters, probability models, etc.—which is a necessary precursor to using positive involvement as part of an argument in favor of some voting methods over others in future work.

First, we make some stage-setting conceptual points in Sections 1.1-1.2, discuss related work in Section 1.3, and recall technical preliminaries in Section 2. Section 3 contains our main discussion, as well as results of our simulations. Our methodological points in Section 3 can be applied to measuring the extent of violation of other axioms in addition to positive involvement. We conclude in Section 4.

1.1 Strategic considerations

It is important to distinguish the perspective on positive involvement adopted above—we view it as an axiom, like monotonicity, that rules out perverse responses to unequivocal increases in support for a candidate—from a *strategic perspective* on positive involvement, which is also adopted in the literature. A voting method's violating positive involvement may give a voter an incentive for *strategic abstention*: if a voter i knows that by casting a sincere ballot with her favorite candidate a ranked first, this would kick a out of the set of winners, then i may prefer not to vote rather than to cast a sincere ballot. This is so whenever a would be the unique winner were i not to vote but a would not be were i to vote sincerely.

However, there are two important qualifications. First, if we are considering strategic abstention, we may also consider *strategic voting*. For certain voting methods, whenever a candidate a would be the unique winner were i not to vote, there is always *some* linear order with a ranked first that i can cast as her ballot to keep a the unique winner,¹ so i has no incentive to abstain if she can vote strategically.² Second, we must consider the case where a would not be the *unique* winner were i not to vote. E.g., suppose that if i were not to vote, the result would be a tie between a and c , where c is i 's least favorite candidate, whereas if i were to vote sincerely, the unique winner would be b , who is i 's second favorite candidate. Depending on i 's utility function, i may well prefer b to a tiebreaking process applied to $\{a, c\}$, in which case i would prefer to vote sincerely rather than to abstain. So not every situation witnessing a failure of positive involvement is one in which the voter would prefer to abstain rather than vote sincerely.

Thus, the problem with violating positive involvement is not just that it may incentivize strategic abstention. Even in a society in which all voters always vote and always vote sincerely, so there is no risk of strategic abstention instead of sincere voting, we still find violations of positive involvement perverse responses to unequivocal increases in support for a candidate.

1.2 Positive involvement vs. participation

It is also important to distinguish positive involvement from the axiom of *participation* [37]. Participation is usually stated for *resolute* voting methods that map each profile of voter preferences to a *unique* winning candidate.³ A resolute voting method satisfies participation (and an arbitrary voting method

¹When a voter casts a ballot with her true favorite a ranked first but with deviations from her sincere preference lower down on the ballot in order to get a elected, Dowding and Van Hees [14] call this a *sincere manipulation*.

²This is obvious if i can submit as her ballot a strict weak order instead of a linear order, since then she may submit a fully indifferent ballot, which for many voting methods is equivalent to abstention, or a ballot with a on top followed by an indifference class containing all other candidates. However, here we assume ballots are linear orders.

³Resolute voting methods (defined on any profile of voter preferences, as in Definition 3) either fail to treat voters equally—by violating the axiom of anonymity—or fail to treat candidates equally—by violating the axiom of neutrality (see [54, § 2.3]).

satisfies what could be called *resolute participation*) if adding a new voter who ranks x above y cannot result in a change from x being the unique winner to y being the unique winner.⁴ Crucially, it is not required that the new voter ranks x in first place. Thus, the new voter may rank other candidates above x , thereby hurting x 's prospects. Participation says that ranking those other candidates above x should not hurt x more than it hurts y , so it should not result in y becoming the unique winner. But this is not so clear in profiles containing majority cycles. In the presence of cycles, the main threat to x may be a candidate z who does not threaten y . Then a new voter with a ranking of the form $z \dots xy \dots$ may do more harm to x than to y . This is only a sketch of an argument raising doubts about the participation axiom, but the key point is this: unlike violations of positive involvement, some violations of participation are not perverse responses to unequivocal increases in support for a candidate x , as some violations involve adding a voter who ranks other candidates above x . Of course, violations of participation give a voter an incentive to abstain rather than to vote sincerely, but that is a separate strategic issue that we set aside.

The term “No Show Paradox” was introduced by Fishburn and Brams [23] for violations of what is now called *negative involvement*.⁵ Later Moulin [37] changed the meaning of “No Show Paradox” to refer to violations of participation. Violations of positive and negative involvement are called instances of the “Strong No Show Paradox” by Pérez [41].⁶ Pérez concludes that the Strong No Show Paradox is a common flaw of many *Condorcet consistent* voting methods, which are methods that always pick a Condorcet winner—a candidate who is majority preferred to every other candidate—if one exists.

1.3 Related Work

This study fits into a line of recent work using computer simulations to estimate the frequency of violations of various voting criteria for different voting methods [42, 25, 6, 5]. The most closely related previous study is that of Brandt et al. [6] on the frequency of the “No Show Paradox” in Moulin’s sense [37], i.e., the frequency of violations of the participation axiom. As the authors observe, “While it is known that certain voting rules suffer from this paradox in principle, the extent to which it is of practical concern is not well understood” (p. 520), a comment that also applies to violations of positive involvement. To fill the gap in the case of the No Show Paradox, Brandt et al. use computer simulations—as well as analytic results obtained using Ehrhart theory—for six Condorcet consistent voting methods, three of which overlap with our list: Baldwin, Copeland, and Nanson. They find that “for few alternatives, the probability of the NSP is rather small (less than 4% for four alternatives and all considered preference models, except for Copeland’s rule). As the number of alternatives increases, the NSP becomes much more likely and which rule is most susceptible to abstention strongly depends on the underlying distribution of preferences” (p. 520). This observation about alternatives matches what we find for positive involvement, where the frequency of violations significantly increases with the number of alternatives.

⁴For other definitions of participation for irresolute voting methods, see [41, 31, 44].

⁵The axiom of negative involvement [43, 41] states that adding a new voter who ranks a candidate last should not result in the candidate going from being a loser to a winner. The analysis of this paper can also be applied to negative involvement, but for the sake of space we focus on positive involvement.

⁶Plassman and Tideman [42] use “Strong No Show Paradox” to refer to violations of participation, contradicting Pérez’s use. Pérez [41] 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 [20] and Felsenthal and Nurmi [18, 19] call them the “P-TOP” and “P-BOT” paradoxes, respectively.

2 Preliminaries

2.1 Profiles

Fix infinite sets \mathcal{V} and \mathcal{X} of *voters* and *candidates*, respectively. A given election will use only finite subsets $V \subseteq \mathcal{V}$ and $X \subseteq \mathcal{X}$. We consider elections in which each voter submits a ranking of all the candidates, which we assume is a strict linear order. For a set X , let $\mathcal{L}(X)$ be the set of all strict linear orders on X . For $P \in \mathcal{L}(X)$, we write ‘ xPy ’ for $(x, y) \in P$; for $a \in X$, let $\text{Rank}(a, P) = |\{b \in X \mid bPa\}| + 1$.

Definition 1. A *profile* is a function $\mathbf{P} : V \rightarrow \mathcal{L}(X)$ for some nonempty finite $V \subseteq \mathcal{V}$, which we denote by $V(\mathbf{P})$ (called the set of *voters in P*), and nonempty finite $X \subseteq \mathcal{X}$, which we denote by $X(\mathbf{P})$ (called the set of *candidates in P*). We call $\mathbf{P}(i)$ voter i ’s *ballot* and write ‘ \mathbf{P}_i ’ for $\mathbf{P}(i)$.

Definition 2. For a profile \mathbf{P} and set $C \subseteq V(\mathbf{P})$ of voters, let \mathbf{P}_{-C} be the restriction of \mathbf{P} to $V(\mathbf{P}) \setminus C$. For a single voter $i \in V(\mathbf{P})$, let $\mathbf{P}_{-i} = \mathbf{P}_{-\{i\}}$. Given profiles \mathbf{P} and \mathbf{P}' such that $X(\mathbf{P}) = X(\mathbf{P}')$ and $V(\mathbf{P}) \cap V(\mathbf{P}') = \emptyset$, we define the profile $\mathbf{P} + \mathbf{P}' : V(\mathbf{P}) \cup V(\mathbf{P}') \rightarrow \mathcal{L}(X(\mathbf{P}))$ as follows: if $i \in V(\mathbf{P})$, then $(\mathbf{P} + \mathbf{P}')(i) = \mathbf{P}(i)$, and if $i \in V(\mathbf{P}')$, then $(\mathbf{P} + \mathbf{P}')(i) = \mathbf{P}'(i)$.

2.2 Voting Methods

Definition 3. A *voting method* is a function F on the domain of all profiles such that for any profile \mathbf{P} , $\emptyset \neq F(\mathbf{P}) \subseteq X(\mathbf{P})$.

In this paper, we consider the following voting methods. We chose these methods because they give a broad representation of different classes of methods (e.g., scoring rules, iterative methods, and Condorcet methods, including C1 and C2 methods) and can be efficiently computed (except for Ranked Pairs).

Positional scoring rules. A *scoring vector* is a vector $\langle s_1, \dots, s_n \rangle$ of numbers such that for each $m \in \{1, \dots, n-1\}$, $s_m \geq s_{m+1}$. Given a profile \mathbf{P} with $|X(\mathbf{P})| = n$, $x \in X(\mathbf{P})$, a scoring vector \vec{s} of length n , and $i \in V(\mathbf{P})$, define $\text{score}_{\vec{s}}(x, \mathbf{P}_i) = s_r$ where $r = \text{Rank}(x, \mathbf{P}_i)$. Let $\text{score}_{\vec{s}}(x, \mathbf{P}) = \sum_{i \in V(\mathbf{P})} \text{score}_{\vec{s}}(x, \mathbf{P}_i)$. A voting method F is a *positional scoring rule* if there is a map \mathcal{S} assigning to each natural number n a scoring vector of length n such that for any profile \mathbf{P} with $|X(\mathbf{P})| = n$, $F(\mathbf{P}) = \text{argmax}_{x \in X(\mathbf{P})} \text{score}_{\mathcal{S}(n)}(x, \mathbf{P})$. Two well-known examples of scoring rules are:

$$\mathbf{Plurality}: \mathcal{S}(n) = \langle 1, 0, \dots, 0 \rangle \quad \text{and} \quad \mathbf{Borda}: \mathcal{S}(n) = \langle n-1, n-2, \dots, 1, 0 \rangle.$$

Instant Runoff (also known as Alternative Vote, Ranked Choice, and Single Transferable Vote): Iteratively remove all candidates with the fewest number of voters who rank them first, until there is a candidate who is a majority winner (i.e., ranked first by a strict majority of voters).⁷ 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.

Coombs [11, 27]: Iteratively remove all candidates with the most number of voters who rank them last, until there is a candidate who is a majority winner. If, at some stage of the removal process, all remaining candidates have the same number voters who rank them last (so all candidates would be removed), then all remaining candidates are selected as winners.

Baldwin [1]: Iteratively remove all candidates with the smallest Borda score, until there is a single candidate remaining. If, at some stage of the removal process, all remaining candidates have the same Borda score (so all candidates would be removed), then all remaining candidates are selected as winners.

⁷When there is more than one candidate with the fewest number of voters who rank them first, this definition of Instant Runoff, taken from [49, p. 7], eliminates all of them. For another way of handling such ties, see [51] and Footnote 14.

Nanson [38]: There are two versions of this voting method [40]. **Strict Nanson** (resp. **Weak Nanson**) iteratively removes all candidates whose Borda score is strictly less than (resp. less than or equal to) the average Borda score of the candidates remaining at that stage, until one candidate remains. If, at some stage of the removal process, all remaining candidates have the same Borda score (so all candidates would be removed), then all remaining candidates are selected as winners. Although Nanson seems to have intended **Weak Nanson** (see [40]), the literature on computational social choice usually interprets ‘Nanson’ as **Strict Nanson** (see, e.g., [39, 54, 21, 10]), so we focus on **Strict Nanson** in this paper.

Bucklin [28]: Given a candidate a in a profile \mathbf{P} and a positive integer n , say that a is an n -th level majority winner in \mathbf{P} if a strict majority of voters rank a in n -th place or higher (thus, a majority winner in the usual sense is a 1st level majority winner). Where k is the smallest positive integer for which there is at least one k -th level majority winner, Bucklin selects as winners the k -th level majority winners for whom the most voters rank them in k -th place or higher.⁸

The other methods that we study in this paper are from a broad class of so-called majority or margin-based methods. We need the following notation for defining these methods.

Definition 4. For a profile \mathbf{P} and $a, b \in X(\mathbf{P})$, let $\text{Margin}_{\mathbf{P}}(a, b) = |\{i \in V(\mathbf{P}) \mid a\mathbf{P}_i b\}| - |\{i \in V(\mathbf{P}) \mid b\mathbf{P}_i a\}|$ be the margin of a over b in \mathbf{P} . The *margin graph* of \mathbf{P} is the weighted directed graph whose set of vertices is $X(\mathbf{P})$ with an edge from a to b when $\text{Margin}_{\mathbf{P}}(a, b) > 0$, weighted by $\text{Margin}_{\mathbf{P}}(a, b)$.

Copeland [12]: The Copeland score of $a \in X(\mathbf{P})$ is the number of $b \in X(\mathbf{P})$ such that $\text{Margin}_{\mathbf{P}}(a, b) > 0$ minus the number of $b \in X(\mathbf{P})$ such that $\text{Margin}_{\mathbf{P}}(b, a) > 0$. Then $\text{Copeland}(\mathbf{P})$ is the set of $x \in X(\mathbf{P})$ with maximal Copeland score.⁹

Top Cycle [48, 46] (also known as Smith and GETCHA): For a profile \mathbf{P} , let $a \succ_{\mathbf{P}} b$ if $\text{Margin}_{\mathbf{P}}(a, b) > 0$. Let $\succ_{\mathbf{P}}^*$ be the transitive closure of $\succ_{\mathbf{P}}$. Then $\text{TopCycle}(\mathbf{P}) = \{x \in X(\mathbf{P}) \mid \text{for all } y \in X(\mathbf{P}), x \succ_{\mathbf{P}}^* y\}$.¹⁰

Uncovered Set (Gillies version) [26]: Given $a, b \in X(\mathbf{P})$, a *Gillies covers* b in \mathbf{P} if $\text{Margin}_{\mathbf{P}}(a, b) > 0$ and for all $c \in X(\mathbf{P})$, if $\text{Margin}_{\mathbf{P}}(c, a) > 0$, then $\text{Margin}_{\mathbf{P}}(c, b) > 0$. Then $\text{UC}(\mathbf{P})$ is the set of candidates who are not Gillies covered in \mathbf{P} . This is called the Gillies Uncovered Set in [15].

Ranked Pairs [50]: For a profile \mathbf{P} and $T \in \mathcal{L}(\{(x, y) \mid x \neq y \text{ and } \text{Margin}_{\mathbf{P}}(x, y) \geq 0\})$, called the *tie-breaking ordering*, a pair (x, y) of candidates has a *higher priority* than a pair (x', y') of candidates according to T when either $\text{Margin}_{\mathbf{P}}(x, y) > \text{Margin}_{\mathbf{P}}(x', y')$ or $\text{Margin}_{\mathbf{P}}(x, y) = \text{Margin}_{\mathbf{P}}(x', y')$ and $(x, y) T (x', y')$. We construct a *Ranked Pairs ranking* $\succ_{\mathbf{P}, T} \in \mathcal{L}(X)$ as follows:

1. Initialize $\succ_{\mathbf{P}, T}$ to \emptyset .
2. If all pairs (x, y) with $x \neq y$ and $\text{Margin}_{\mathbf{P}}(x, y) \geq 0$ have been considered, then return $\succ_{\mathbf{P}, T}$. Otherwise let (a, b) be the pair with the highest priority among those with $a \neq b$ and $\text{Margin}_{\mathbf{P}}(a, b) \geq 0$ 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. The set $\text{RP}(\mathbf{P})$ of Ranked Pairs winners is the set of all $x \in X(\mathbf{P})$ such that x is the maximum of $\succ_{\mathbf{P}, T}$ for some tie-breaking ordering T . This is the

⁸The **Simplified Bucklin** method selects as winners all k -th level majority winners. We ran our simulations with Simplified Bucklin as well as Bucklin; its frequency of violating positive involvement is similar to Bucklin but slightly worse.

⁹An equivalent definition of Copeland defines the score of $a \in X(\mathbf{P})$ as the number of $b \in X(\mathbf{P})$ with $\text{Margin}_{\mathbf{P}}(a, b) > 0$ plus $1/2$ times the number of $b \in X(\mathbf{P})$ with $\text{Margin}_{\mathbf{P}}(a, b) = 0$. We also ran all of our simulations using the variant of Copeland known as Llull, where the score of $a \in X(\mathbf{P})$ is the number of $b \in X(\mathbf{P})$ with $\text{Margin}_{\mathbf{P}}(a, b) > 0$ plus the number of $b \in X(\mathbf{P})$ with $\text{Margin}_{\mathbf{P}}(a, b) = 0$ (see [17]). Llull’s performance was very similar to that of Copeland.

¹⁰We also ran our simulations on the GOCHA method [46] that differs from Top Cycle only in profiles with margins of 0 between some candidates. Its performance was similar or slightly better than that of Top Cycle, depending on the parameters of the simulation.

“parallel universe” version of Ranked Pairs called RP in [8, 51], distinguished from other non-neutral, non-anonymous, or probabilistic versions of Ranked Pairs.

Since calculating $RP(\mathbf{P})$ is an NP-complete problem [8], we also consider the non-anonymous version of Ranked Pairs proposed by Zavist and Tideman [53], which we call **Ranked Pairs ZT**. Zavist and Tideman propose to use a distinguish voter’s ranking to derive the tie-breaking ordering T . In particular, given $i \in V(\mathbf{P})$, let $T(\mathbf{P}_i)$ be the lexicographic order on $\{(x, y) \mid x \neq y \text{ and } Margin_{\mathbf{P}}(x, y) \geq 0\}$ derived from \mathbf{P}_i . Since different profiles have different sets of voters, we cannot use the same distinguished voter for all profiles. Given a linear order L of \mathcal{V} , for any profile \mathbf{P} , we define $RPZT_L(\mathbf{P})$ to be the set of all $x \in X(\mathbf{P})$ such that x is the maximum of $\succ_{\mathbf{P}, T(\mathbf{P}_i)}$ where i is the minimal element of $V(\mathbf{P})$ according to L .

Beat Path [45]: For $a, b \in X(\mathbf{P})$, a *path from a to b in \mathbf{P}* is a sequence $\rho = x_1, \dots, x_n$ of distinct candidates in $X(\mathbf{P})$ with $x_1 = a$ and $x_n = b$ such that for $1 \leq k \leq n - 1$, $Margin_{\mathbf{P}}(x_k, x_{k+1}) > 0$. The *strength of ρ* is $\min\{Margin_{\mathbf{P}}(x_k, x_{k+1}) \mid 1 \leq k \leq n - 1\}$. Then a defeats b in \mathbf{P} according to Beat Path if the strength of the strongest path from a to b is greater than the strength of the strongest path from b to a . $BP(\mathbf{P})$ is the set of undefeated candidates.

Split Cycle [29]: A *majority cycle in \mathbf{P}* is a sequence $\rho = x_1, \dots, x_n$ of distinct candidates in $X(\mathbf{P})$ except $x_1 = x_n$ such that for $1 \leq k \leq n - 1$, $Margin_{\mathbf{P}}(x_k, x_{k+1}) > 0$. The *strength of ρ* is defined as above for Beat Path. Then a defeats b in \mathbf{P} according to Split Cycle if $Margin_{\mathbf{P}}(a, b)$ is positive and greater than the strength of the strongest majority cycle containing a and b . $SC(\mathbf{P})$ is the set of undefeated candidates.

2.3 Positive Involvement

As explained in Section 1, our interest in this paper is the axiom of positive involvement.

Definition 5. A voting method F satisfies *positive involvement* (PI) if for any profile \mathbf{P} and $x \in F(\mathbf{P})$, if \mathbf{P}' is obtained from \mathbf{P} by adding a new voter who ranks x in first place, then $x \in F(\mathbf{P}')$.

PI is sometimes discussed in terms of the addition of several voters who rank x first. An obvious inductive argument shows that the coalitional version of PI follows from the single-voter version in Definition 5.

Lemma 6. If a voting method F satisfies PI, then it satisfies *coalitional* PI: for any profile \mathbf{P} and $x \in F(\mathbf{P})$, if \mathbf{P}' is a profile with $X(\mathbf{P}) = X(\mathbf{P}')$, $V(\mathbf{P}) \cap V(\mathbf{P}') = \emptyset$, and every voter in \mathbf{P}' ranks x in first place, then $x \in F(\mathbf{P} + \mathbf{P}')$.

The voting methods in Section 2.2 can be classified by their satisfying or violating PI as follows.

Proposition 7. The voting methods Baldwin, Beat Path, Bucklin, Coombs, Copeland, Ranked Pairs, Strict Nanson, Top Cycle, and Uncovered Set all violate PI—see the Appendix for examples. All positional scoring rules, Instant Runoff, and Split Cycle satisfy PI (see [41, 29]).¹¹

Remark 8. For results on violation or satisfaction of PI by other voting methods, see [41, 19]. Note that few known voting methods satisfy both PI and Condorcet consistency. Indeed, Pérez [41] observed that with the exception of the Minimax method [47, 33], which satisfies PI, “all the Condorcet correspondences that (to the best of our knowledge) are proposed in the literature” violate PI (p. 601). Other examples of Condorcet methods violating PI, besides those listed in Proposition 7 (Baldwin, Beat Path,

¹¹Another well-known method satisfying PI is Plurality with Runoff [19, § 4.7.1]. This result depends on a specific way of breaking ties when more than two candidates qualify for the runoff (a candidate a qualifies for the runoff if either a has maximal plurality score or there is a unique candidate with maximal plurality score and a is among the candidates with the second highest plurality score). The version of Plurality with Runoff satisfying PI is the “parallel universe” version that considers all possible duels between candidates who qualify for the runoff; a candidate is a winner just in case they win in one of these duels. If instead all candidates who qualify for the runoff are promoted to a second round decided by Plurality, then the resulting version of Plurality with Runoff does not satisfy PI. Thanks to an anonymous referee for discussion of these points.

	profiles	profile-coalition pairs
absolute frequency	column 1 of Figure 2	column 1 of Figure 4
conditional on disagreement with F_2	columns 2-4 of Figure 2	columns 2-4 of Figure 4
relativized to voter potency	column 1 of Figure 3	column 1 of Figure 5
conditional on disagreement & relativized	columns 2-4 of Figure 3	columns 2-4 of Figure 5

Figure 1: ways of measuring PI violation

Copeland, Ranked Pairs, Strict Nanson, Top Cycle, and Uncovered Set), are Dodgson [13], Kemeny [32], and Young [52] (see [41, 19]).¹² However, Split Cycle is Condorcet consistent and satisfies PI.

3 Quantitative Analysis

We now turn to our discussion of different ways to quantify the extent to which a voting method F violates PI. They are summarized in Figure 1 and explained in the subsections to follow. The most obvious idea is to simply consider the probability that a randomly drawn profile \mathbf{P} (according to some probability model) *witnesses a violation of PI for F* , meaning that there is some voter i in \mathbf{P} such that i 's favorite candidate wins in \mathbf{P}_{-i} according to F but not in \mathbf{P} . However, for medium to large sized electorates, we should expect this probability to be low for the mundane reason that the addition of any *single* voter will rarely cause someone to lose. We will explore several ways to deal with this issue.

One natural idea is to look at violations of *coalitional* PI: say that a profile \mathbf{P} *witnesses a violation of coalitional PI for F* if there is some coalition C of voters in \mathbf{P} with the same ranking (or at least the same top-ranked candidate), and their favorite candidate wins in \mathbf{P}_{-C} but not in \mathbf{P} according to F . Ideally, we would like to estimate the probability that a random profile witnesses a violation of coalitional PI (for various coalition sizes). Unfortunately, we find it too computationally expensive to check for every coalition of voters whether removing that coalition shows a violation of coalitional PI, especially for the thousands of profiles needed for reliable simulation results. To deal with this problem, in Section 3.2 we will assess the probability that a random *profile-coalition pair* witnesses a violation of PI.

However, we will begin in Section 3.1 with the probability that a random profile witnesses a violation of single-voter PI, since this provides a baseline with which to compare everything else that follows.

The probability model for profiles that will serve as our baseline is the Impartial Culture (IC) model: when sampling profiles with n candidates and m voters, the IC model gives every such profile equal probability. In Section 3.3, we consider several other probability models for sampling profiles. For each data point in our graphs, we sampled 25,000 profiles with the indicated even number m of voters and 25,000 profiles with the indicated odd number $m + 1$ of voters, in order to have a mix of even and odd-sized electorates. We used the preflib [36] implementation for each of the probability models. Our code is in the online supplementary material at <https://github.com/epacuit/posinvolvement>.

¹²We did not include Dodgson, Kemeny, and Young in our simulations due to the computational complexity of determining winners for these methods (see [9] and [21, § 4.2]).

3.1 Random profiles

3.1.1 Probability of PI violation

The leftmost column in Figure 2 shows, for several voting methods F that violate PI, the estimated probability that a random profile (according to IC) witnesses a violation of PI for F , as defined above. The important qualitative observations are that (i) the probability of PI violation increases as the number of candidates increases and (ii) the probability of PI violation decreases as the number of voters increases (above 20), for the mundane reason mentioned above—as the number of voters increases, any single voter has less influence in the election (a point to which we return in Section 3.1.3).

3.1.2 Probability of PI violation conditional on voting method disagreement

Although the absolute frequency of PI violation in the leftmost column of Figure 2 is relevant, it is a mistake to think that if a voting method F_1 has a low frequency of violating PI, then this undermines the use of PI as a reason to favor another voting method F_2 , which satisfies PI, over F_1 . First of all, when using PI to help choose between a method F_1 that violates the axiom and a method F_2 that satisfies it, we should ask: *in the profiles in which F_1 and F_2 disagree*, with what frequency does F_1 violate PI?¹³ That is, we should consider the following conditional probability, for a random profile \mathbf{P} (according to the given probability model, for a given number of candidates and voters):

$$Pr(\mathbf{P} \text{ witnesses a violation of PI for } F_1 \mid F_1(\mathbf{P}) \neq F_2(\mathbf{P})). \quad (1)$$

Columns 2-4 of Figure 2 show estimates of this conditional probability for several voting methods F_1 that violate PI compared to three choices for F_2 : Borda (2nd column), Instant Runoff (3rd column), and Split Cycle (4th column). We see a striking increase in the probability that F_1 will violate PI when we condition on F_1 disagreeing with the PI-satisfying method F_2 . Note that if the probability of F_1 violating PI is higher conditional on F_1 disagreeing with F_2 than it is conditional on F_1 disagreeing with F'_2 , then PI provides a stronger argument in the context of deciding between F_1 and F_2 than it does in the context of deciding between F_1 and F'_2 . E.g., let F_1 be Baldwin, F_2 be Split Cycle, and F'_2 be Borda.

Remark 9. Since violation of (single voter) PI requires a close election in order for one voter to affect the outcome, we checked (i) to what extent the frequency of PI violation by Baldwin and Coombs depends on the manner of handling ties at intermediate stages of their elimination procedures and (ii) to what extent the frequency of PI violation by Uncovered Set depends on which of several variants of the Uncovered Set—which can disagree only when some candidates have a margin of 0—one chooses to analyze.

The tie-handling issue arises for Baldwin (resp. Coombs) in case there are multiple candidates with the smallest Borda score (resp. most last place votes) at a given stage of iteration. For each iterative elimination method F , there is a “parallel-universe tie-handling” (PUT) variant (cf. [24]): a candidate a wins under the PUT variant of F just in case there is some linear order L on the set of candidates such that a wins according to F_L , where F_L is defined in the same way as F except that if multiple candidates meet the criterion for elimination at some stage (e.g., have the smallest Borda score, or the most last place votes), only the L -minimal candidate among those candidates is eliminated. We found that the PUT versions of Baldwin and Coombs performed similarly or slightly better than the versions defined in

¹³We think that this methodology (which, as far as we know, is new) should be applied to the study of voting axioms in general, not only to PI.

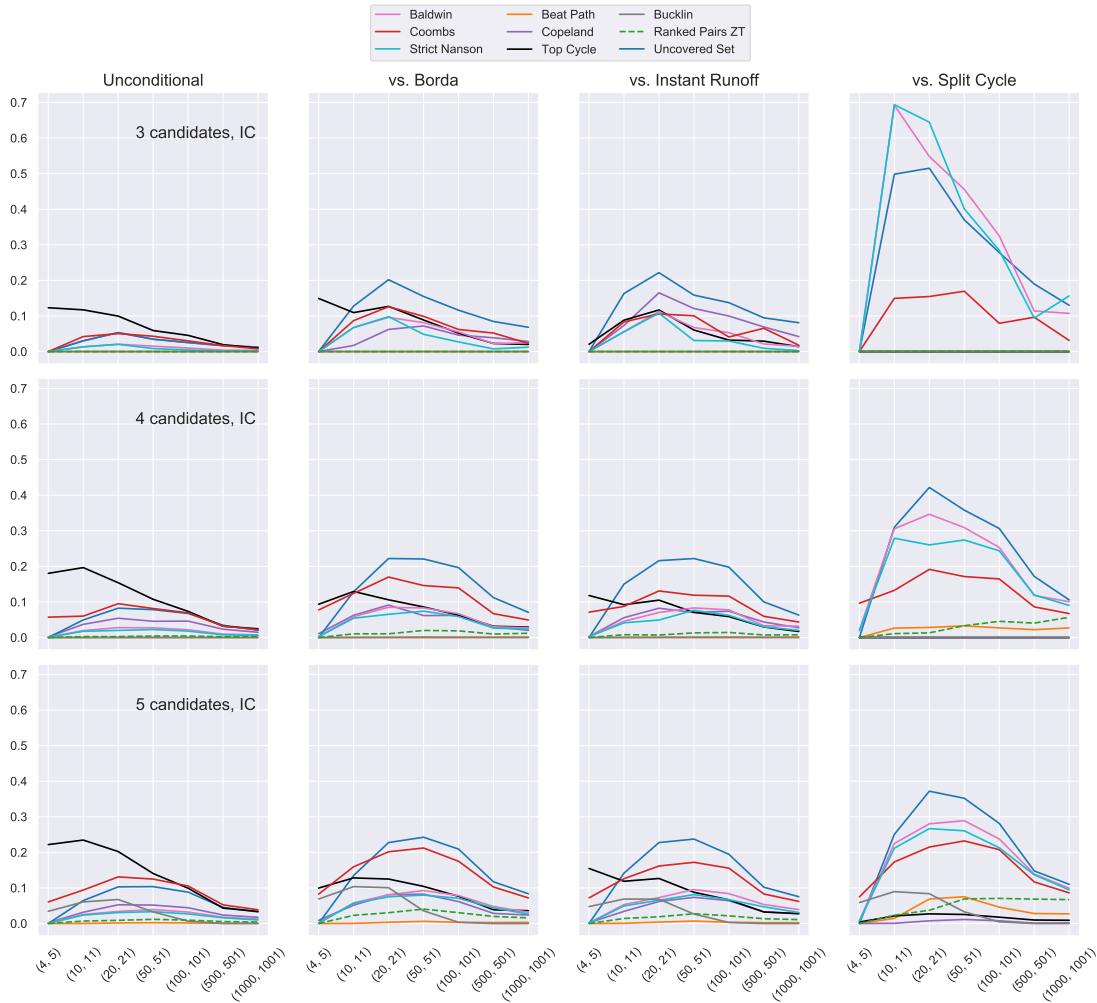


Figure 2: probability of a profile \mathbf{P} witnessing a violation of PI for F either unconditionally (far left), conditional on F disagreeing with Borda in \mathbf{P} (second from left), conditional on F disagreeing with Instant Runoff in \mathbf{P} (second from right), or conditional on F disagreeing with Split Cycle in \mathbf{P} (far right). Ranked Pairs was not included due to computational limitations (cf. [8]). Ranked Pairs ZT is shown with a dashed line to mark it out as our only non-anonymous voting method.

Section 2.2.¹⁴ But given the computational difficulty of determining winners for the PUT versions, we use the versions defined in Section 2.2 for the rest of our analysis.

As for Uncovered Set, we considered three variants in addition to the Gillies Uncovered Set in Section 2.2: the Bordes Uncovered Set, the McKelvey Uncovered Set, and the Fishburn Uncovered Set.¹⁵ Although the absolute frequencies of PI violation for these variants are similar (with Gillies and Fish-

¹⁴We also checked (for 3 and 4 candidates) the difference between conditioning on disagreement with Instant Runoff, as in Figure 2, and conditioning on disagreement with the PUT version of Instant Runoff. We found the results to be similar.

¹⁵Following Duggan [15], say that a Bordes covers b in \mathbf{P} if $\text{Margin}_{\mathbf{P}}(a, b) > 0$ and for all $c \in X(\mathbf{P})$, if $\text{Margin}_{\mathbf{P}}(c, a) \geq 0$, then $\text{Margin}_{\mathbf{P}}(c, b) \geq 0$; and say that a McKelvey covers b if a Gillies covers and Bordes covers b . For Fishburn’s [22] variant, say that a Fishburn covers b in \mathbf{P} if for all $c \in X(\mathbf{P})$, if $\text{Margin}_{\mathbf{P}}(c, a) > 0$, then $\text{Margin}_{\mathbf{P}}(c, b) > 0$, and there is a $c \in X(\mathbf{P})$ such that $\text{Margin}_{\mathbf{P}}(c, b) > 0$ and $\text{Margin}_{\mathbf{P}}(c, a) \leq 0$.

burn almost indistinguishable, and Bordes and McKelvey almost indistinguishable), the frequencies of PI violation conditional on disagreement with Borda, Instant Runoff, and Split Cycle are surprisingly different, with Gillies far worse than the others (see our online supplementary material). In this paper, we focus on the Gillies variant for a worst-case analysis of the family of Uncovered Set variants.

3.1.3 Probability of PI violation relativized to voter potency

Although the results of conditioning on voting method disagreement are striking, even those conditional probabilities are bound to decline as we increase the number of voters, given a single voter’s declining influence. What we should be asking, we think, is how likely a PI violation is compared to how likely it is that any single voter causes a candidate to lose. For this, we introduce the notion of voter *potency*.

Definition 10. For any profile \mathbf{P} and $i \in V(\mathbf{P})$, we say that i is *potent* in \mathbf{P} if $F(\mathbf{P}_{-i}) \not\subseteq F(\mathbf{P})$.¹⁶

Then our proposal is to measure how badly a voting method violates single-voter PI by the ratio

$$\frac{\Pr(\text{there is a voter } i \text{ triggering a PI violation, i.e., } \max(\mathbf{P}_i) \in F_1(\mathbf{P}_{-i}) \setminus F_1(\mathbf{P}))}{\Pr(\text{there is a voter } i \text{ who is potent in } \mathbf{P})}, \quad (2)$$

possibly with these probabilities conditioned on F_1 disagreeing in \mathbf{P} with an F_2 that satisfies PI. For voting methods violating PI, the numerator may be small (as opposed to 0 for methods that satisfy PI), but the denominator is also small; thus, the ratio may be surprisingly large. The results are shown in Figure 3.

Remark 11. One curious feature of the last column of Figure 3 is that for 3 or 4 candidates, the probability of Copeland and Top Cycle violating PI conditional on their disagreeing with Split Cycle in \mathbf{P} goes *down to zero*. This is initially puzzling: since Split Cycle satisfies PI, mustn’t Copeland and Top Cycle disagree with Split Cycle whenever they violate PI? The answer is ‘yes’, but they may disagree with Split Cycle in the *smaller profile* \mathbf{P}_{-i} rather than in \mathbf{P} ; indeed, this always happens when the cited methods violate single-voter PI for 3 or 4 candidates. Thus, there is an ambiguity in the idea of “conditioning on disagreement with F_2 ”—we could condition on disagreement in \mathbf{P}_{-i} , in \mathbf{P} , in both, or in one or the other. So far we have only conditioned on disagreement in \mathbf{P} , since so far in our random sampling we only draw a profile \mathbf{P} ; but in Section 3.2.2, when we randomly sample profile-voter pairs, we will be able to conveniently condition on disagreement in \mathbf{P}_{-i} or \mathbf{P} .

3.2 Random profile-coalition pairs

Given the computational difficulty of checking whether a profile witnesses a violation of coalitional PI, we now turn to randomly sampling *profile-coalition pairs*, i.e., pairs $(\mathbf{P}, \mathbf{P}_C)$ of a random profile \mathbf{P} (for a given number of candidates and voters) and a “coalitional” profile \mathbf{P}_C obtained by randomly selecting a single ranking and then assigning it to all voters in C (with $C \cap V(\mathbf{P}) = \emptyset$). Such a profile-coalition pair *witnesses a violation of PI* just in case the unanimously top-ranked candidate in \mathbf{P}_C is a winner in \mathbf{P} but not in $\mathbf{P} + \mathbf{P}_C$. Note that even if the pair $(\mathbf{P}, \mathbf{P}_C)$ does not witness a violation of PI, the profile $\mathbf{P} + \mathbf{P}_C$ may witness a violation of coalitional PI by the removal of a different coalition C' . Thus, this approach misses violations of coalitional PI involving other coalitions. The probability that a random profile-coalition pair witnesses a violation of PI may be considerably lower than the probability that a random profile witnesses a violation of coalitional PI. Nonetheless, there are benefits of this approach, such as (i) our being able to feasibly study coalitions of more than one voter, (ii) our being able to search profiles up to 5,000 voters or up to 10 candidates, and (iii) our being able to conveniently condition on voting method disagreement before or after the new coalition of voters joins the election (see Remark 11).

¹⁶Note that i being potent in \mathbf{P} is stronger than i being *pivotal* in \mathbf{P} in the sense that $F(\mathbf{P}_{-i}) \neq F(\mathbf{P})$.




Figure 3: probability of a profile \mathbf{P} witnessing a violation of PI for F , divided by the probability of \mathbf{P} having a potent voter, either unconditionally (far left), conditional on F disagreeing with Borda in \mathbf{P} (second from left), conditional on F disagreeing with Instant Runoff in \mathbf{P} (second from right), or conditional on F disagreeing with Split Cycle in \mathbf{P} (far right).

3.2.1 Probability of PI violation

The leftmost column of Figure 4 shows estimates for the probability that a randomly selected profile-coalition pair $(\mathbf{P}, \mathbf{P}_C)$ witnesses a violation of PI, depending on whether C is a coalition of a single voter (first row), a coalition of 0.25% of the total voter size (second row), a coalition of 0.5% (third row), or a coalition of 0.75% (fourth row). We show only the results for 500, 1,000, and 5,000 voter profiles, since with 100 voters or fewer, all the coalition sizes round to 1 voter. As expected, in the single voter case, the probability that a random profile-voter pair witnesses a violation of PI is much lower than the probability that a random profile witnesses a violation of single-voter PI (compare the leftmost column and third row of Figure 2). Indeed, the difference is roughly an order of magnitude. The probability does increase for

larger coalition sizes, but the probability of a random profile-0.5%-coalition pair witnessing a violation of PI is still less than half the probability of a random profile \mathbf{P} witnessing a violation of single-voter PI.

3.2.2 Probability of PI violation conditional on disagreement

As suggested by earlier results (Section 3.1.2), conditioning on the probability that in our random profile-coalition pair $(\mathbf{P}, \mathbf{P}_C)$ the voting method F_1 disagrees with the method F_2 that satisfies PI—where disagreement now means that $F_1(\mathbf{P}) \neq F_2(\mathbf{P})$ or $F_1(\mathbf{P} + \mathbf{P}_C) \neq F_2(\mathbf{P} + \mathbf{P}_C)$ (recall Remark 11)—dramatically increases the probability that F_1 violates PI. The results are shown in columns 2-4 of Figure 4.

3.2.3 Probability of PI violation relativized to voter potency

We can also apply the idea of relativizing to voter potency from Section 3.1.3 in the paradigm of sampling profile-coalition pairs. In particular, we want to estimate the value of the following ratio for a random profile-coalition pair $(\mathbf{P}, \mathbf{P}_C)$, a randomly chosen $a \in F(\mathbf{P})$, and a random new voter:

$$\frac{\Pr(\text{the new voters in } C \text{ cause } a \text{ to lose} \mid a \text{ is the favorite of the new voters in } C)}{\Pr(\text{the new voters in } C \text{ cause } a \text{ to lose})}. \quad (3)$$

Intuitively, the numerator should be 0, but as we know, for voting methods that violate PI it is not zero. It is small, but on the other hand, the denominator is also small, since small coalitions of voters have limited influence in any election. Thus, as in Section 3.1.3, the ratio itself can be surprisingly large. Ratios of the form $\Pr(A \mid B)/\Pr(A)$, as in (3), are used in confirmation theory to measure the degree of support that evidence provides for a hypothesis (see, e.g., [30, p. 54]). Note that $\Pr(A \mid B)/\Pr(A) = \Pr(B \mid A)/\Pr(B)$. Thus, we can phrase our question in one of two ways: If you learn that the new voters in C rank a first, to what extent does this support the hypothesis that the voters in C will cause a to lose? Or, alternatively, if you learn that the new voters in C caused a to lose, to what extent does this support the hypothesis that the voters in C ranked a first? Of course for methods satisfying PI the answer is “not at all.”

Results are shown in Figure 5, where columns 2-4 also condition the probabilities in (3) on F_1 disagreeing with F_2 in \mathbf{P} or $\mathbf{P} + \mathbf{P}_C$. Note the strikingly high ratios for some methods—especially Coombs and Top Cycle, but also Uncovered Set, Baldwin, Nanson, and Copeland—as we increase the candidates.

3.3 Other probability models

In addition to sampling profiles with the IC model, we tried several other probability models. In the Pólya-Eggenberger urn model [2], each voter in turn randomly draws a linear order from an urn. Initially the urn is $\mathcal{L}(X)$. If a voter randomly chooses L from the urn, we return L to the urn plus $\alpha \in \mathbb{N}$ copies of L . IC is the special case where $\alpha = 0$. The Impartial Anonymous Culture (IAC) is the special case where $\alpha = 1$. We also considered $\alpha = 10$ (as in [7, 6, 5]) for the model we call URN. In the Mallows’s model (see [34, 35]), given a reference ranking $L_0 \in \mathcal{L}(X)$ and $\phi \in (0, 1]$, the probability that a voter’s ballot is $L \in \mathcal{L}(X)$ is $\Pr_{L_0}(L) = \phi^{\tau(L, L_0)}/C$ where $\tau(L, L_0) = \binom{|X|}{2} - |L \cap L_0|$, the Kendall-tau distance of L to L_0 , and C is a normalization constant. We considered two reference rankings, L_0 and its converse L_0^{-1} (e.g., L_0 ranks candidates from more liberal to more conservative, and L_0^{-1} vice versa), in which case the probability that a voter’s ballot is L is $\frac{1}{2}\Pr_{L_0}(L) + \frac{1}{2}\Pr_{L_0^{-1}}(L)$. We set $\phi = 0.8$ (as in [6, 5]).

The most important finding concerning the different probability models is that as we deviate from the IC model, the main phenomena seen in the previous simulations do not disappear. Figure 6 shows the 5-candidate case of Figure 5 under these three probability models, plus IC again (for easy comparison). We see that the surprisingly high values for the ratio measure in (3) are not artifacts of the IC model.

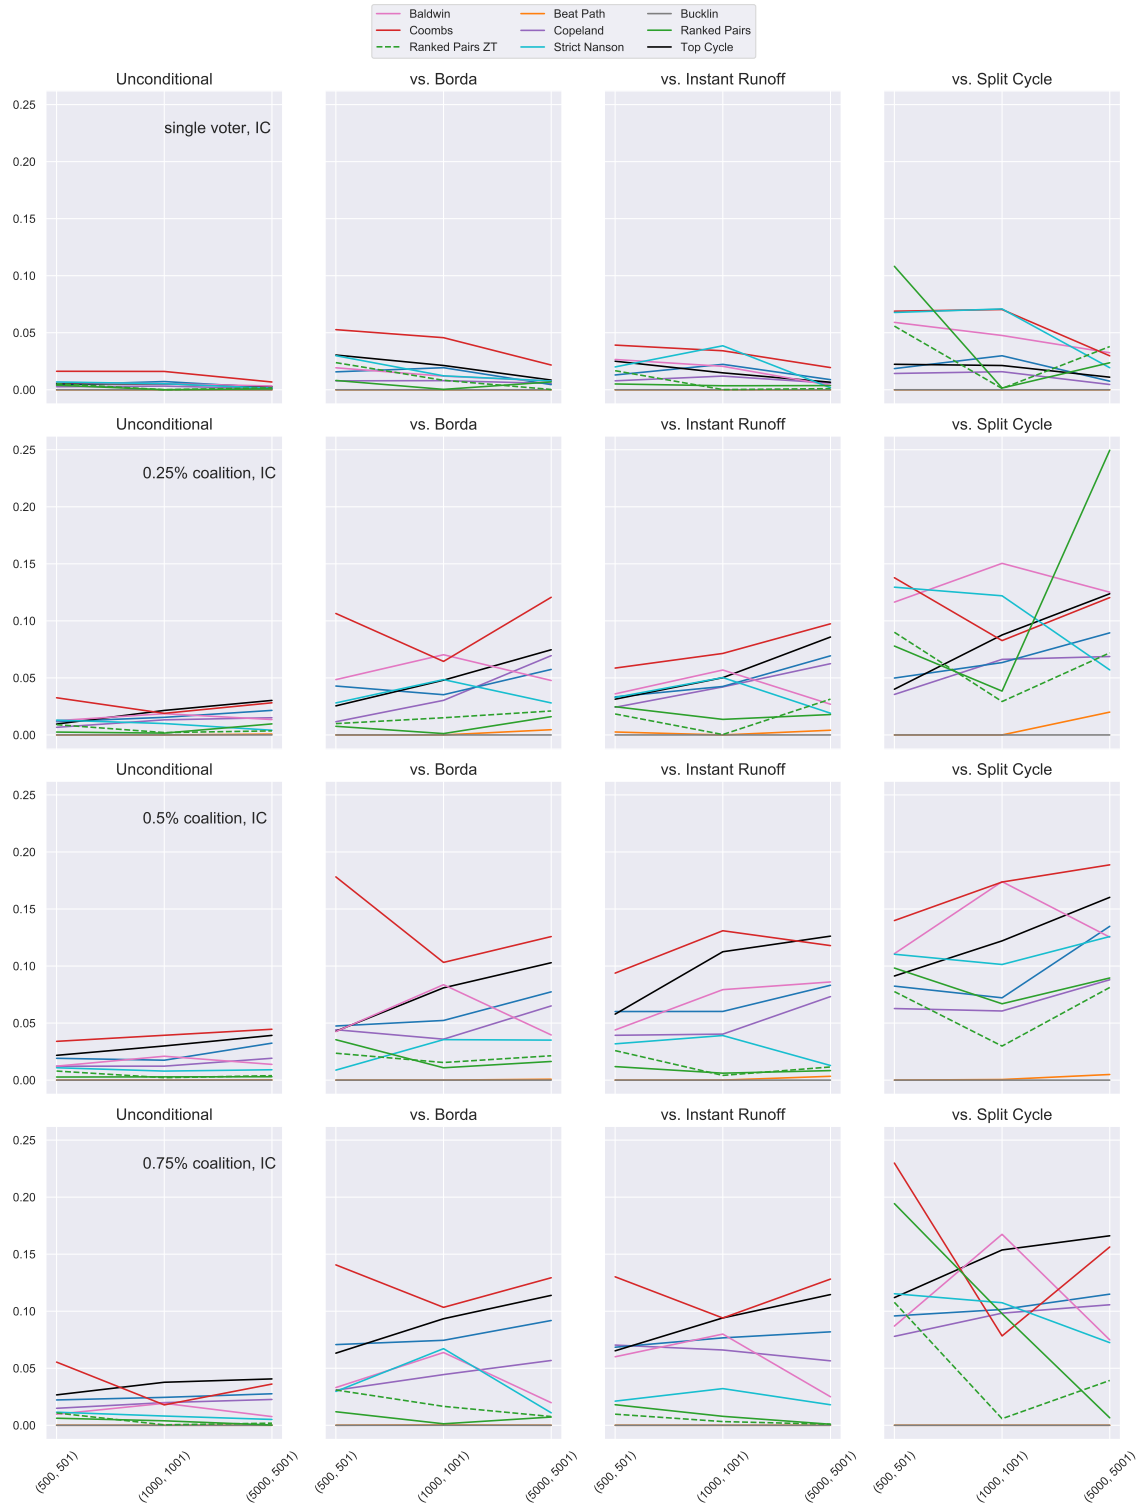


Figure 4: probability of a profile-coalition pair $(\mathbf{P}, \mathbf{P}_C)$ witnessing a violation of PI either unconditionally (far left), conditional on disagreement in \mathbf{P} or $\mathbf{P} + \mathbf{P}_C$ with Borda (second from left), with Instant Runoff (second from right), or with Split Cycle (far right). The first row is for a single voter coalition, the second for coalitions of 0.25% of initial total voter size, the third for 0.5%, and the fourth for 0.75%.



Figure 5: the probability ratio in (3) from Section 3.2.3 for profile-coalition pairs $(\mathbf{P}, \mathbf{P}_C)$ with $|C| = 1$, either unconditionally (far left), conditional on disagreement with Borda in \mathbf{P} or $\mathbf{P} + \mathbf{P}_C$ (second from left), conditional on disagreement with Instant Runoff in \mathbf{P} or $\mathbf{P} + \mathbf{P}_C$, or conditional on disagreement with Split Cycle in \mathbf{P} or $\mathbf{P} + \mathbf{P}_C$ (far right). Due to computational limitations, results for 10 candidates and 5,000/5,001 voters were not obtained, and results for Ranked Pairs were obtained only for 3-5 candidates and 500-5,001 voters (as increasing the number of voters decreases the likelihood of tied margins).

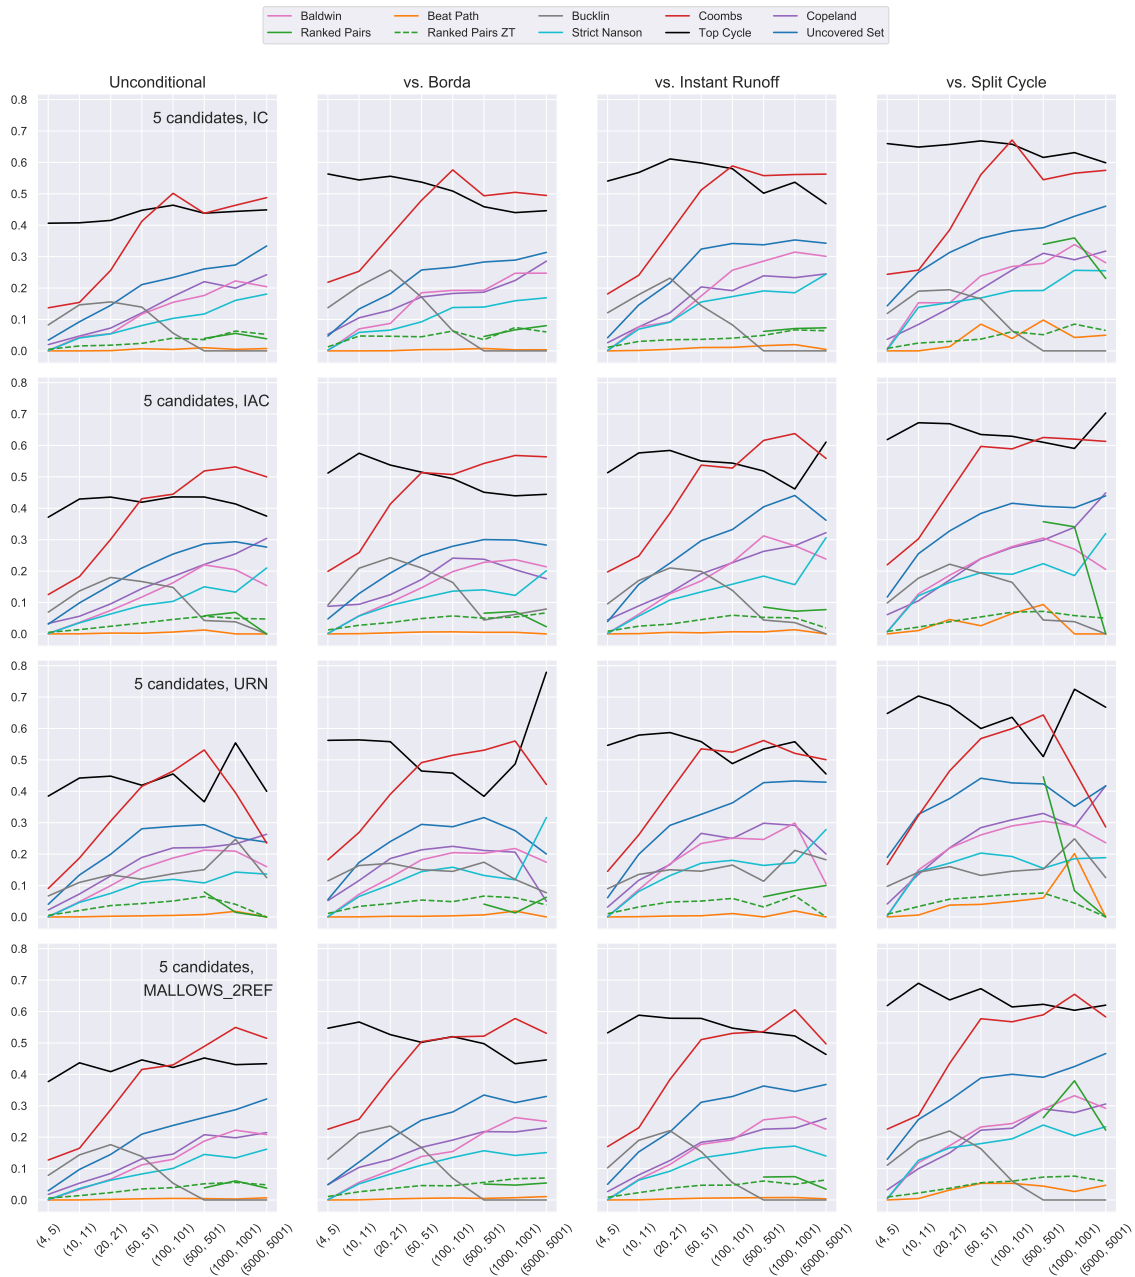


Figure 6: the same measures as in row 3 of Figure 5 but under different probability models. As before, due to computational limitations, results for Ranked Pairs were obtained only for 500-5,001 voters.

4 Conclusion

In their book on failures of monotonicity-type axioms for voting methods, Nurmi and Felsenthal [19] argue that violations of positive (and negative) involvement are the “more dramatic types of no show paradox,” wherein “the issues of legitimacy of outcomes and questionable voter incentives are far more obvious” (p. 86). In this paper, we proposed ways of measuring the extent to which these dramatic violations occur. The next step in future work is to marshal these measures when arguing in favor of some voting methods over others. Under certain voting methods that violate PI, much of a voter’s potency is turned against them—in particular, against their desire to see their favorite candidate elected. It is like adding insult to injury to be told that not only does your vote have little chance of influencing the election, but also your vote may cause your *favorite* candidate to lose, and *the probability of your vote doing so is not at all insignificant compared to the probability of it causing any candidate to lose*. The probabilities may be large enough to raise real concerns in small elections in committees and clubs, unless one somehow knows in advance that none of the probability models used here is relevant for the committee (e.g., one knows in advance that the committee has single-peaked preferences). But for large elections, defenders of voting methods that violate PI may simply respond: don’t worry—your vote probably won’t cause your favorite to lose, because it probably won’t influence the election at all. The same response could be used to try to dismiss concerns about violations of other single-voter axioms, such as monotonicity or single-voter strategyproofness. Such responses raise what is perhaps the ultimate paradox of voting: why do voters vote in large elections? The significance of single-voter axioms for large elections may turn on the answer to that question.

A Appendix

The following are examples of profiles witnessing PI violation for the voting methods from Proposition 7. In the profiles, the new voter’s ranking is highlighted in grey.


Example 12 (Baldwin violates PI). Candidate c is the winner on the left (the order of elimination is a, d, b) but not on the right (the order of elimination is d, b, c):

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

Baldwin winners: $\{c\}$ Baldwin winners: $\{a\}$

Since there are no ties for the lowest Borda score at any stage, the same holds for Baldwin PUT.

Example 13 (Beat Path violates PI). The left graph below can be realized as the margin graph of a profile with 11 voters; the right graph is obtained by adding a voter with the ranking $b a c d$. Then b is a winner on the left but not on the right:



Beat Path: $\{a, b, c, d\}$

Beat Path: $\{a\}$
New ranking: $b a c d$

Example 14 (Bucklin violates PI). Candidate c is the winner on the left (there are no 1st or 2nd level majority winners; a and c are 3rd level majority winners; and more voters rank c in 3rd place or higher than a) but not on the right (because e is the unique 2nd level majority winner):

1	1	1	1
a	a	b	c
b	e	e	d
c	c	c	a
e	b	d	b
d	d	a	e

1	1	1	1	1
a	a	b	c	c
b	e	e	d	e
c	c	c	a	b
e	b	d	b	d
d	d	a	e	a

Bucklin: $\{c\}$

Bucklin: $\{e\}$

The same example works for Simplified Bucklin, as defined in Footnote 8, only the winners on the left are a and c , and the winner on the right is e .

Example 15 (Coombs violates PI). Candidate a is the winner on the left (because b and c are eliminated in the first round, and then a is the majority winner) but not on the right (because b is eliminated first, and then c is the majority winner):

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

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


Coombs: $\{a\}$

Coombs: $\{c\}$

Under Coombs PUT, on the left, since b and c tie for most last place votes, we consider two cases: (i) if we eliminate b in the first round, then c is the majority winner; (ii) if we eliminate c in the first round, then a is the majority winner. Hence both a and c win in the profile on the left. In the profile on the right, c is the unique winner for Coombs PUT for the same reason as for Coombs.

Example 16 (Copeland and Uncovered Set violate PI). Candidate b is a Copeland winner on the left but not on the right:¹⁷

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




Copeland winners: $\{a, b, c\}$

Gillies Uncovered Set winners: $\{a, b, c\}$

Fishburn Uncovered Set winners: $\{a, b, c\}$

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



Copeland winners: $\{a\}$

Gillies Uncovered Set winners: $\{a, c\}$

Fishburn Uncovered Set winners: $\{a\}$

The same example works for the Gillies (resp. Fishburn) version of Uncovered Set (recall Remark 9): on the left, the winners are a , b , and c , while on the right, the winners are a and c (resp. a).

Example 17 (Strict Nanson violates PI). Candidate c is the winner on the left (the order of elimination is b , d , a) but not on the right (both a and b are eliminated in the first round, followed by c):

1	1	3	3	1	1
a	d	c	a	d	d
c	b	b	d	c	c
d	c	a	c	a	b
b	a	d	b	b	a


Strict Nanson winners: $\{c\}$

1	1	3	3	1	1	1
a	d	c	a	d	d	c
c	b	b	d	c	c	d
d	c	a	c	a	b	b
b	a	d	b	b	a	a


Strict Nanson winners: $\{d\}$

¹⁷The same example provides a PI violation for the Llull method defined in Footnote 9. The Llull winners are the same as the Copeland winners in both profiles.


Example 18 (Ranked Pairs violates PI). The left graph below can be realized as the margin graph of a profile with 20 voters; then the right graph below is obtained by adding a voter with the ranking $d a b c$. Then candidate d is a winner on the left (we first lock in the edges (a, c) , (d, c) , and (d, b) , and then there is a choice of whether to prioritize the (b, a) edge, in which case d wins, or the (a, d) edge, in which case a wins) but not on the right (we first lock in the edges (a, c) , (d, c) , and (d, b) , and then there is a choice of whether to prioritize (c, b) or (a, d) , but in either case (a, d) eventually gets locked in):



For Ranked Pairs ZT, in the profile on the left below, where the tiebreaking voter's ranking is indicated with (\star) , b is the winner (we lock in (d, c) , (b, c) , (b, d) , and then (c, a)). But on the right, a is the winner (there are no cycles, so all edges are locked in, and then the ranking (\star) ranks a above b).



Example 19 (Top Cycle violates PI). Candidate b is a winner on the left but not the right:



References

[1] J. M. Baldwin (1926): *A technique of the Nanson preferential majority system of election*. *Transactions and Proceedings of the Royal Society of Victoria* 39, pp. 45–52.

[2] Sven Berg (1985): *Paradox of voting under an urn model: The effect of homogeneity*. *Public Choice* 47, pp. 377–387, doi:10.1007/BF00127533.

[3] Felix Brandt, Vincent Conitzer & Ulle Endriss (2013): *Computational Social Choice*. In Gerhard Weiss, editor: *Multigent Systems*, MIT Press, Cambridge, Mass., pp. 213–283.

- [4] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang & Ariel D. Procaccia, editors (2016): *Handbook of Computational Social Choice*. Cambridge University Press, New York, doi:10.1017/cbo9781107446984.003.
- [5] Felix Brandt, Christian Geist & Martin Strobel (2020): *Analyzing the Practical Relevance of the Condorcet Loser Paradox and the Agenda Contraction Paradox*. In M. Diss & V. Merlin, editors: *Evaluating Voting Systems with Probability Models: Essays by and in Honor of William Gehrlein and Dominique Lepelley*, Springer, Berlin, pp. 97–115, doi:10.1007/978-3-030-48598-6_5.
- [6] Felix Brandt, Johannes Hofbauer & Martin Strobel (2019): *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)*, International Foundation for Autonomous Agents and MultiAgent Systems, pp. 520–528.
- [7] Felix Brandt & Hans Georg Seedig (2014): *On the Discriminative Power of Tournament Solutions*. In M. Lübbecke, A. Koster, P. Letmathe, R. Madlener, B. Peis & G. Walther, editors: *Operations Research Proceedings 2014*, Springer, Cham, pp. 53–58, doi:10.1007/978-3-319-28697-6_8.
- [8] Markus Brill & Felix Fischer (2012): *The Price of Neutrality for the Ranked Pairs Method*. In: *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence (AAAI-12)*, AAAI Press, pp. 1299–1305.
- [9] Ioannis Caragiannis, Edith Hemaspaandra & Lane A. Hemaspaandra (2016): *Dodgson's Rule and Young's Rule*. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang & Ariel D. Procaccia, editors: *Handbook of Computational Social Choice*, Cambridge University Press, New York, pp. 103–126, doi:10.1017/cbo9781107446984.005.
- [10] Vincent Conitzer & Toby Walsh (2016): *Barriers to Manipulation in Voting*. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang & Ariel D. Procaccia, editors: *Handbook of Computational Social Choice*, Cambridge University Press, New York, pp. 127–145, doi:10.1017/cbo9781107446984.006.
- [11] Clyde Hamilton Coombs (1964): *A Theory of Data*. John Wiley and Sons, New York.
- [12] A. H. Copeland (1951): *A 'reasonable' social welfare function*. Notes from a seminar on applications of mathematics to the social sciences, University of Michigan.
- [13] Charles L. Dodgson (1995): *A Method of Taking Votes on More Than Two Issues*. In Iain McLean & Arnold Urken, editors: *Classics of Social Choice*, University of Michigan Press, Ann Arbor, pp. 288–298.
- [14] Keith Dowding & Martin Van Hees (2008): *In Praise of Manipulation*. *British Journal of Political Science* 38(1), pp. 1–15, doi:10.1017/S000712340800001X.
- [15] John Duggan (2013): *Uncovered Sets*. *Social Choice and Welfare* 41(3), pp. 489–535, doi:10.1007/s00355-012-0696-9.
- [16] Ulle Endriss, editor (2017): *Trends in Computational Social Choice*. AI Access.
- [17] Piotr Faliszewski, Edith Hemaspaandra, Lane A. Hemaspaandra & Jörg Rothe (2007): *Llull and Copeland Voting Broadly Resist Bribery and Control*. In: *Proceedings of the Twenty-Second AAAI Conference on Artificial Intelligence (AAAI-07)*, AAAI Press, pp. 724–730.
- [18] Dan S. Felsenthal & Hannu Nurmi (2016): *Two types of participation failure under nine voting methods in variable electorates*. *Public Choice* 168, pp. 115–135, doi:10.1007/s11127-016-0352-5.
- [19] Dan S. Felsenthal & Hannu Nurmi (2017): *Monotonicity failures afflicting procedures for electing a single candidate*. Springer, Cham, doi:10.1007/978-3-319-51061-3.
- [20] Dan S. Felsenthal & Nicolaus Tideman (2013): *Varieties of failure of monotonicity and participation under five voting methods*. *Theory and Decision* 75, pp. 59–77, doi:10.1007/s11238-012-9306-7.
- [21] Felix Fischer, Olivier Hudry & Rolf Niedermeier (2016): *Weighted Tournament Solutions*. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang & Ariel D. Procaccia, editors: *Handbook of Computational Social Choice*, Cambridge University Press, New York, pp. 85–102, doi:10.1017/cbo9781107446984.004.
- [22] Peter C. Fishburn (1977): *Condorcet Social Choice Functions*. *SIAM Journal on Applied Mathematics* 33(3), pp. 469–489, doi:10.1137/0133030.

- [23] Peter C. Fishburn & Steven J. Brams (1983): *Paradoxes of Preferential Voting*. *Mathematics Magazine* 56(4), pp. 207–214, doi:10.2307/2689808.
- [24] Rupert Freeman, Markus Brill & Vincent Conitzer (2015): *General Tiebreaking Schemes for Computational Social Choice*. In: *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015)*, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, pp. 1401–1409.
- [25] William V. Gehrlein & Dominique Lepelley (2017): *Elections, Voting Rules and Paradoxical Outcomes*. Springer, Cham, doi:10.1007/978-3-319-64659-6.
- [26] Donald B. Gillies (1959): *Solutions to general non-zero-sum games*. In A. W. Tucker & R. D. Luce, editors: *Contributions to the Theory of Games*, Princeton University Press, Princeton, New Jersey, doi:10.1515/9781400882168-005.
- [27] Bernard Grofman & Scott L. Feld (2004): *If you like the alternative vote (a.k.a. the instant runoff), then you ought to know about the Coombs rule*. *Electoral Studies* 23(4), pp. 641–659, doi:10.1016/j.electstud.2003.08.001.
- [28] Clarence Hoag & George Hallett (1926): *Proportional Representation*. Macmillan, New York.
- [29] Wesley H. Holliday & Eric Pacuit (2020): *Split Cycle: A New Condorcet Consistent Voting Method Independent of Clones and Immune to Spoilers*. ArXiv:2004.02350.
- [30] Paul Horwich (2016): *Probability and Evidence*. Cambridge University Press, Cambridge, doi:10.1017/CBO9781316494219.
- [31] José L. Jimeno, Joaquín Pérez & Estefanía García (2009): *An extension of the Moulin No Show Paradox for voting correspondences*. *Social Choice and Welfare* 33(3), pp. 343–359, doi:10.1007/s00355-008-0360-6.
- [32] John G. Kemeny (1959): *Mathematics without Numbers*. *Daedalus* 88(4), pp. 577–591.
- [33] Gerald H. Kramer (1977): *A dynamical model of political equilibrium*. *Journal of Economic Theory* 16(2), pp. 310–334, doi:10.1016/0022-0531(77)90011-4.
- [34] C. L. Mallows (1957): *Non-Null Ranking Models. I*. *Biometrika* 44(2), pp. 114–130, doi:10.2307/2333244.
- [35] John Marden (1995): *Analyzing and Modeling Rank Data*. CRC Press, New York, doi:10.1201/b16552.
- [36] Nicholas Mattei & Toby Walsh (2013): *PrefLib: A Library of Preference Data*. In: *Proceedings of Third International Conference on Algorithmic Decision Theory*, Springer, pp. 259–270, doi:10.1007/978-3-642-41575-3_20.
- [37] Hervé Moulin (1988): *Condorcet's Principle Implies the No Show Paradox*. *Journal of Economic Theory* 45(1), pp. 53–64, doi:10.1016/0022-0531(88)90253-0.
- [38] E. J. Nanson (1882): *Methods of election*. *Transactions and Proceedings of the Royal Society of Victoria* 19, pp. 197–240.
- [39] Nina Narodytska, Toby Walsh & Lirong Xia (2011): *Manipulation of Nanson's and Baldwin's Rules*. In: *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence*, AAAI Press, pp. 713–718.
- [40] Emerson M. S. Niou (1987): *A Note on Nanson's Rule*. *Public Choice* 54(2), pp. 191–193, doi:10.1007/BF00123006.
- [41] Joaquín Pérez (2001): *The Strong No Show Paradoxes are a common flaw in Condorcet voting correspondences*. *Social Choice and Welfare* 18(3), pp. 601–616, doi:10.1007/s003550000079.
- [42] Florenz Plassmann & T. Nicolaus Tideman (2014): *How frequently do different voting rules encounter voting paradoxes in three-candidate elections?* *Social Choice and Welfare* 42, pp. 31–75, doi:10.1007/s00355-013-0720-8.
- [43] Donald G. Saari (1995): *Basic Geometry of Voting*. Springer, Berlin, doi:10.1007/978-3-642-57748-2.
- [44] M. Remzi Sanver & William S. Zwicker (2012): *Monotonicity properties and their adaptation to irresolute social choice rules*. *Social Choice and Welfare* 39(2/3), pp. 371–398, doi:10.1007/s00355-012-0654-6.

- [45] Markus Schulze (2011): *A new monotonic, clone-independent, reversal symmetric, and condorcet-consistent single-winner election method*. *Social Choice and Welfare* 36, pp. 267–303, doi:10.1007/s00355-010-0475-4.
- [46] Thomas Schwartz (1986): *The Logic of Collective Choice*. Columbia University Press, New York, doi:10.7312/schw93758.
- [47] Paul B. Simpson (1969): *On Defining Areas of Voter Choice: Professor Tullock on Stable Voting*. *The Quarterly Journal of Economics* 83(3), pp. 478–490, doi:10.2307/1880533.
- [48] John H. Smith (1973): *Aggregation of Preferences with Variable Electorate*. *Econometrica* 41(6), pp. 1027–1041, doi:10.2307/1914033.
- [49] Alan D. Taylor & Allison M. Pacelli (2008): *Mathematics and Politics: Strategy, Voting, Power, and Proof*, 2nd edition. Springer, New York, doi:10.1007/978-0-387-77645-3.
- [50] T. Nicolaus Tideman (1987): *Independence of Clones as a Criterion for Voting Rules*. *Social Choice and Welfare* 4, pp. 185–206, doi:10.1007/bf00433944.
- [51] Jun Wang, Sujoy Sikdar, Tyler Shepherd Zhibing Zhao, Chunheng Jiang & Lirong Xia (2019): *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, doi:10.1609/aaai.v33i01.33012189.
- [52] H. P. Young (1977): *Extending Condorcet's Rule*. *Journal of Economic Theory* 16, pp. 335–353, doi:10.1016/0022-0531(77)90012-6.
- [53] T. M. Zavist & T. Nicolaus Tideman (1989): *Complete Independence of Clones in the Ranked Pairs Rule*. *Social Choice and Welfare* 6, pp. 167–173, doi:10.1007/bf00303170.
- [54] William S. Zwicker (2016): *Introduction to the Theory of Voting*. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang & Ariel D. Procaccia, editors: *Handbook of Computational Social Choice*, Cambridge University Press, New York, pp. 23–56, doi:10.1017/cbo9781107446984.003.