

The River Voting Method

Michelle Döring¹, Markus Brill², Jobst Heitzig³

¹Hasso Plattner Institute, University of Potsdam, Potsdam, Germany

²Department of Computer Science, University of Warwick, Coventry, United Kingdom

³FutureLab on Game Theory & Networks of Interacting Agents, Complexity Science Department, Potsdam Institute for Climate Impact Research (PIK), Potsdam, Germany
michelle.doering@hpi.de, markus.brill@warwick.ac.uk, jobst.heitzig@pik-potsdam.de

Abstract

We introduce *River*, a novel Condorcet-consistent voting method that is based on pairwise majority margins and can be seen as a simplified variation of Tideman’s *Ranked Pairs* method. *River* is simple to explain, simple to compute even “by hand,” and gives rise to an easy-to-interpret certificate in the form of a directed tree. Like *Ranked Pairs* and Schulze’s *Beat Path* method, *River* is a refinement of the *Split Cycle* method and shares with those many desirable properties, including independence of clones. Unlike the other three methods, *River* satisfies a strong form of resistance to agenda-manipulation that is known as *independence of Pareto-dominated alternatives*.

1 Introduction

The task of making a collective decision on the basis of individual rankings is fundamental to social choice theory (Arrow, Sen, and Suzumura 2002; Brandt et al. 2016) and has a wide range of applications in artificial intelligence (Fürnkranz and Hüllermeier 2003; Asbell et al. 2021; Mishra 2023). For instance, Köpf et al. (2024) have recently employed a rank-aggregation approach to align large language models (LLMs) with human preferences.

We use the framework of *voting theory* (Zwicker 2016) and interpret rankings as preference orders given by a set of *voters* over a set of *alternatives*. A key principle in collective decision making is *majority rule*—the idea that the decision should follow what is seen as “the will of the majority” (May 1952). In the simple case of only two alternatives, majority rule is unambiguous and chooses the alternative that is ranked first by more than half of the voters. This principle is used by many common decision methods which are based on *pairwise comparisons* between the alternatives.

In real-world elections, a *Condorcet winner* often emerges—an alternative that defeats all others in pairwise comparison. This alternative is considered to be the most suitable choice for the election winner. In the absence of a Condorcet winner, each alternative faces at least one majority defeat. To then decide on a winner, one can encode the pairwise majority comparisons as a graph, with each alternative represented by a vertex, and an edge (y, x) indicating

that y defeats x . This forms a *tournament graph*, and a selection process based on this is known as a *tournament solution* (Laslier 1997; Brandt, Brill, and Harrenstein 2016). Here, majority rule is embodied by “*beatpaths*”: For any chosen winner x , if there is a defeat $y \rightarrow x$, there should be a path of defeats leading from x back to y : $x \rightarrow z_1 \rightarrow \dots \rightarrow z_k = y$. These alternatives form the *Smith set* (a.k.a. *GETCHA* or *top cycle*) (Good 1971; Smith 1973; Schwartz 1986).

While tournament theory extends this concept by studying several refinements of the Smith set based solely on pairwise majority defeats, another strand of research focuses on the strength of these majority defeats, measured by the numerical majority margin (the number of voters ranking x over y minus the number of voters with opposite rankings) (Fischer, Hudry, and Niedermeier 2016). The tournament graph with edges weighted by this margin is called the *margin graph*.

A natural extension of the Smith set from the tournament graph to the margin graph is that for any defeat $y \rightarrow x$ of a winner x , there should be a “rebutting” beatpath from x to y , with each defeat in the path being at least as strong as $y \rightarrow x$. This allows defending the choice of x against claims of the form “a majority ranks y over x ” by pointing to a sequence of equally strong (or stronger) claims of the same form leading back to y . An alternative fulfilling this property is called *immune*. Since one can show that at least one immune alternative always exists in every election, this approach can be seen as a natural way to operationalize the concept of majority rule (Dung 1995; Heitzig 2002). The corresponding voting method, choosing all immune alternatives, is called *Split Cycle* (Holliday and Pacuit 2023a). This method has many appealing properties but it often chooses multiple alternatives as the winner. There exist several popular methods that always choose a subset of immune alternatives and are therefore refinements of *Split Cycle*, most notably *Ranked Pairs* (Tideman 1987) and *Beat Path* (Schulze 2011). Both typically select a single winner, except in rare cases of ties, and each satisfies a distinct set of desirable properties. The same is true for *Stable Voting* (Holliday and Pacuit 2023b), the most recently introduced refinement of *Split Cycle*.

All four methods—*Split Cycle*, *Ranked Pairs*, *Beat Path*, and *Stable Voting*—suffer from a weakness related to agenda manipulation: introducing a new alternative z might alter the winning set even if z is *Pareto-dominated* by an existing alternative y (i.e., all voters rank y over z).

Formally, these methods violate *independence of Pareto-dominated alternatives (IPDA)*. This property goes beyond the standard notion of Pareto efficiency (stating that Pareto-dominated alternatives should not be selected) and requires that Pareto-dominated alternatives should not influence the outcome at all. Without IPDA, it becomes possible to propose additional, Pareto-dominated alternatives in order to gain an advantage or disturb the voting process. As a result, IPDA is often considered a desirable property in social choice theory (Fishburn 1973; Richelson 1978; Ching 1996; Chebotarev and Shamis 1998; Gonzalez, Laruelle, and Solal 2019; Öztürk 2020; Brandl and Peters 2022). Recently, (Greaves and Cotton-Barratt 2023) explored IPDA in the context of decision making under moral uncertainty.

Our contribution In this paper, we introduce the novel social choice function *River*, a refinement of Split Cycle. River closely resembles Ranked Pairs in that both methods construct an acyclic subgraph of the margin graph rooted at the winner. However, River distinguishes itself through simplicity and interpretability: it selects a spanning tree rather than a dense subgraph, ensuring each alternative has a unique predecessor.

We show that River satisfies many desirable social-choice-theoretic axioms, including independence of Smith-dominated alternatives and independence of clones when equipped with a suitable tiebreaker. Most notably, we show that River satisfies IPDA, making it the only known refinement of Split Cycle satisfying this property.

Omitted proofs are in the full version of this paper.

2 Preliminaries

Let $[k] := \{1, \dots, k\}$. We consider elections with a set N of $n \geq 1$ voters with preferences over a set A of $m \geq 2$ alternatives. The preferences² of a voter $i \in N$ are given as a strict ranking, i.e., a linear order \succ_i over A and $x \succ_i y$ denotes that voter i ranks alternative x above y . A (preference) profile $\mathbf{P} = (\succ_1, \dots, \succ_n)$ is a list containing the preferences of all voters. We may denote its corresponding set of voters and alternatives by $N(\mathbf{P})$ and $A(\mathbf{P})$. For alternative $x \in A$, \mathbf{P}_{-x} denotes the restriction of \mathbf{P} to $A \setminus \{x\}$.

The *majority margin* of x over y according to \mathbf{P} is

$$m_{\mathbf{P}}(x, y) = |\{i \in N : x \succ_i y\}| - |\{i \in N : y \succ_i x\}|.$$

When $m_{\mathbf{P}}(x, y) > 0$, we say x *defeats* y , denoted $x \succ_{\mathbf{P}} y$, and when $m_{\mathbf{P}}(x, y) \geq 0$, we say x *weakly defeats* y , denoted $x \succeq_{\mathbf{P}} y$. The *margin graph* $\mathcal{M}_{\mathbf{P}}$ of a profile \mathbf{P} has vertex set A and an edge from x to y whenever $x \succeq_{\mathbf{P}} y$. Each edge $(x, y) \in E(\mathcal{M}_{\mathbf{P}})$ is weighted by its margin, $m_{\mathbf{P}}(x, y) \geq 0$. We refer to the edges $E(\mathcal{M}_{\mathbf{P}})$ of the margin graph as *weak majority edges*. Further, the *strict margin graph* $\mathcal{M}_{\mathbf{P}}^{>0}$ is the subgraph of $\mathcal{M}_{\mathbf{P}}$ containing only edges with $m_{\mathbf{P}}(x, y) > 0$. We call these *strict majority edges*. If clear from the context, we drop index \mathbf{P} from $x \succ y$, $m(x, y)$, \mathcal{M} , and $\mathcal{M}^{>0}$.

¹The method was first described by one of the authors of this paper in a 2004 mailing list post (Heitzig 2004).

²In this paper, we do not consider strategic misrepresentation of preferences, which would require distinguishing between a voter's actual preferences and the ranking they choose on their ballot.

A *majority path* is a path in $\mathcal{M}^{>0}$, i.e., a sequence $p = (x_1, \dots, x_\ell)$ of distinct alternatives with $m(x_i, x_{i+1}) > 0$ for all $i \in [\ell - 1]$. The *strength* of p is its lowest margin, i.e.,

$$\text{strength}(p) = \min\{m(x_i, x_{i+1}) : i \in [\ell - 1]\}.$$

Analogously, *majority cycles* are closed paths in $\mathcal{M}^{>0}$, and their strength is the lowest margin on the cycle.

A *Condorcet winner* x is an alternative which defeats all other alternatives, i.e., $x \succ y$ for all $y \in A \setminus \{x\}$. Such an alternative does not always exist. A natural extension is the notion of dominant sets. A nonempty set $X \subseteq A$ is *dominant*, if $x \succ y$ for all $x \in X$ and $y \in A \setminus X$. There always exists at least one dominant set: the set of alternatives A .

2.1 Immunity against Majority Complaints

In the absence of a Condorcet winner, every alternative suffers at least one weak majority defeat, potentially eliciting a ‘majority complaint’. Certain alternatives can be defended against such complaints by showing the existence of majority paths from the defended alternative to each alternative defeating it. In the *unweighted* setting, the *Smith set* consists precisely of the alternatives that can be defended in this way.

Definition 2.1. The Smith set $Sm(\mathbf{P}) \subseteq A$ of a preference profile \mathbf{P} is the unique inclusion-minimal dominant subset.

Such resistance against majority complaints can be generalized to the weighted setting in a straightforward way by taking the strength of defeats into account (Heitzig 2002; Holliday and Pacuit 2023a).

Definition 2.2. Given a preference profile \mathbf{P} , an alternative $x \in A$ is called *immune*, if for every $y \in A \setminus \{x\}$ with $m(y, x) > 0$, there exists a majority path p in $\mathcal{M}^{>0}$ from x to y with $\text{strength}(p) \geq m(y, x)$.

2.2 Social Choice Functions

A *social choice function (SCF)* F maps a preference profile \mathbf{P} to a non-empty set $F(\mathbf{P}) \subseteq A$ of *winning* alternatives. Given two social choice functions F and G , we call F a *refinement* of G if $F(\mathbf{P}) \subseteq G(\mathbf{P})$ for all profiles \mathbf{P} .

The social choice function **Split Cycle (SC)** (Holliday and Pacuit 2023a) selects all immune alternatives as winners. Formally, for each majority cycle C in $\mathcal{M}^{>0}$, an edge in C with margin $\text{strength}(C)$ is called a *splitting edge*. Splitting edges are, therefore, exactly the edges with the lowest margin within the cycle. The *Split Cycle diagram*, denoted \mathcal{M}^{SC} , is the subgraph of $\mathcal{M}^{>0}$ obtained by removing the splitting edges of every cycle. As a result, \mathcal{M}^{SC} contains no majority cycles. The winners under Split Cycle are the alternatives with no incoming edge in \mathcal{M}^{SC} .

In this paper, we consider three refinements of Split Cycle: *Ranked Pairs* (Tideman 1987), *Beat Path* (Schulze 2011), and *Stable Voting* (Holliday and Pacuit 2023b).

Ranked Pairs (RP) Starting with an empty graph on vertices A , add edges from \mathcal{M} one at a time in order of decreasing margin, skipping any edge that would create a cycle. In case of ties in the margins, break them using a tiebreaker. The unique alternative without incoming edges in the resulting *Ranked Pairs diagram* \mathcal{M}^{RP} is the Ranked Pairs winner.




Figure 1: An election with 6 alternatives showing the behaviour of the different social choice functions. From left to right are the margin graph \mathcal{M} , the Split Cycle graph \mathcal{M}^{SC} with winning set $\{a, b, c\}$, the Stable Voting winner a with the deciding election without e , and the Ranked Pairs graph \mathcal{M}^{RP} with winner c , and a table of all beatpaths with Beat Path winner b .

Beat Path (BP) In $\mathcal{M}^{>0}$, compute the strongest majority path from x to y for all ordered pairs (x, y) of alternatives. If the strongest path from x to y has greater strength than the strongest path from y to x , then x is said to “beat” y . One can show that there always exists an unbeaten alternative. All unbeaten alternatives are Beat Path winners.

Stable Voting (SV) The winners are defined recursively. If $|A| = 1$, that alternative wins. Otherwise, list all weak majority edges $(x, y) \in \mathcal{M}$ where x is immune (a Split Cycle winner), sorted by decreasing margin with ties broken using a tiebreaker. In that order, return the first x that is a Stable Voting winner in the election without y . One can show that such an alternative always exists.

Since Ranked Pairs, Beat Path and Stable Voting are refinements of Split Cycle, they always select immune alternatives. Additionally, Split Cycle and Ranked Pairs offer a certificate for the immunity of the winner via “rebutting” paths in the graphs \mathcal{M}^{SC} and \mathcal{M}^{RP} . Figure 1 illustrates the four choice functions on an example profile with six alternatives.

2.3 Tiebreaking and Uniquely Weighted Profiles

Computing winners for Ranked Pairs and Stable Voting involves ordering weak majority edges by their margin. With sufficiently many voters, *ties* between margins (i.e., two edges with the same margin) are rare, and we often restrict our attention to profiles where ties do not occur.

Definition 2.3. A preference profile \mathbf{P} is uniquely weighted if for all $v, w, x, y \in A_{\mathbf{P}}$ with $x \neq y$ and $(v \neq x \text{ or } w \neq y)$, we have $m_{\mathbf{P}}(x, y) \neq m_{\mathbf{P}}(v, w)$ and $m_{\mathbf{P}}(x, y) \neq 0$.

An SCF is called *resolute* if it always outputs a single alternative. Split Cycle is not resolute, even for uniquely weighted profiles (Figure 1), while Ranked Pairs, Beat Path, and Stable Voting are resolute on these profiles (Holliday and Pacuit 2023a). For preference profiles that are not uniquely weighted, *tiebreakers* are required to compute Ranked Pairs and Stable Voting winners.

Definition 2.4. A tiebreaker is a linear order over all edges in the margin graph of a preference profile.

Tiebreakers can either be specified directly or derived using a *tiebreaker function*.

Definition 2.5. A tiebreaker function takes a preference profile \mathbf{P} as input and returns a linear order over all edges in the margin graph of \mathbf{P} .

In the literature, the notion for *tiebreaker* and *tiebreaker function* are often used interchangeably.

A minimal condition that we require from tiebreaker functions is *consistency*, which demands that the ordering of pairs does not change when alternatives are removed.

Definition 2.6. A tiebreaker function is consistent, if an edge (a, b) precedes (c, d) in the tiebreaker on \mathbf{P} if and only if (a, b) precedes (c, d) in the tiebreaker for \mathbf{P}_{-x} , for any $x \notin \{a, b, c, d\}$.

A simple way to define a consistent tiebreaker function consists in sorting the edges (x, y) *lexicographically* based on a given linear order of the alternatives: first, by the source alternative x , and then among edges with the same source, by the target alternative y . The linear order of alternatives that is used as the basis for this lexicographic ordering of edges may depend on the profile \mathbf{P} ; e.g., it could be defined via the preferences \succ_i of a given voter $i \in N$.³ When discussing properties of SCFs, we specify whether the analysis assumes uniquely weighted profiles (where no tiebreaker is needed) or general profiles, in which case we may specify conditions on the tiebreaker (function).⁴

3 The River Method

Aiming to satisfy as many axiomatic properties as possible can lead to definitions of rather complex social choice functions that are hard to analyse and non-trivial to execute. For example, *Stable Voting* requires a recursive computation of Stable Voting and Split Cycle winners on smaller instances, reducing the profile one alternative at a time.

Moreover, even if a social choice function is guaranteed to only select immune alternatives, the “rebutting” paths that witness the immunity are sometimes nontrivial to find.

³A similar construction was used in a paper by Zavist and Tideman (1989). Note, however, that Zavist and Tideman (1989) use tiebreakers to rank *unordered* pairs, whereas our tiebreaker functions rank ordered pairs.

⁴For an SCF F that depends on tiebreaking, one can define a *parallel-universes tiebreaking* (PUT) variant of F that returns all alternatives that are chosen by F under *some* tiebreaker (Conitzer, Rognlie, and Xia 2009; Wang et al. 2019). Determining the winners of the PUT variant of Ranked Pairs is NP-complete (Brill and Fischer 2012). In contrast, it was recently shown that the PUT variant of River can be computed in polynomial time (Döring, Malanowski, and Neubert 2026). We do not consider parallel-universes tiebreaking in this paper.

For example, Ranked Pairs produces a typically very large, acyclic subgraph of the margin graph in which it is nontrivial to find rebutting paths without the help of a computer. In the full version of this paper, we present an example with 14 alternatives.

The main motivation behind *River* is simplicity. The method is (i) simple to explain (see below for the procedural definition), (ii) simple to compute (the winner can easily be calculated “by hand”), and (iii) gives rise to unique rebutting paths that are easy to spot in the resulting diagram.

Operating similarly to Ranked Pairs, *River* “splits” majority cycles by starting with an empty graph and iteratively adding weak majority edges in order of decreasing margin while maintaining acyclicity. In contrast to Ranked Pairs, for which acyclicity is the only criterion, *River* also avoids adding edges towards an already defeated alternative. This results in a tree of weak majority edges with the winner as the root.

Definition 3.1. *River* (RV) operates as follows, given a preference profile \mathbf{P} :

1. Order the weak majority edges $E(\mathcal{M}_{\mathbf{P}})$ by decreasing margin (using a tiebreaker if necessary).
2. Initialize an empty graph on the set of all alternatives A .
3. Process the edges of $\mathcal{M}_{\mathbf{P}}$ following the order defined in Step 1 and add an edge if this addition neither creates (Cy) a cycle, nor (Br) a branching (two in-edges for an alternative).
4. The *River* winner $RV(\mathbf{P})$ is the unique source in the resulting *River* diagram \mathcal{M}^{RV} .

We refer to the *cycle condition* by (Cy) and the *branching condition* by (Br). Note that *River* differs from Ranked Pairs only in (Br), which ensures \mathcal{M}^{RV} to be a tree.⁵ The consequences of this subtle change will be thoroughly discussed in the rest of the paper. It is easy to check that *River* winners are immune:

Proposition 3.2. *River* is a refinement of Split Cycle, i. e., $x \in RV(\mathbf{P})$ implies $x \in SC(\mathbf{P})$.

Proof. Let \mathbf{P} be a preference profile. We show the contraposition of the claim, i. e., $x \notin SC(\mathbf{P})$ implies $x \notin RV(\mathbf{P})$. If $x \notin SC(\mathbf{P})$, there is an alternative $y \in A \setminus \{x\}$ with $(y, x) \in E(\mathcal{M}^{SC})$. Hence, the edge (y, x) is not a splitting edge of any cycle in $\mathcal{M}_{\mathbf{P}}^{\geq 0}$. If $(y, x) \in E(\mathcal{M}^{RV})$, then $x \notin RV(\mathbf{P})$ follows. So, assume $(y, x) \notin E(\mathcal{M}^{RV})$. If the edge was rejected because of (Br), then there must be another edge $(z, x) \in E(\mathcal{M}^{RV})$ and hence $x \notin RV(\mathbf{P})$. If the edge was rejected because of (Cy), it must be that (y, x) closes a cycle with edges of margin at least $m(y, x)$, which is a contradiction to (y, x) not being a splitting edge. \square

Figure 2 shows the *River* diagram for the margin graph from Figure 1 (with $RV(\mathbf{P}) \neq RP(\mathbf{P})$) containing notably fewer edges than the Ranked Pairs diagram. In fact, \mathcal{M}^{RV} always contains exactly $m - 1$ edges, whereas \mathcal{M}^{RP} may

⁵The name “*River*” refers to the fact that (ideal) rivers may merge, but a river typically does not branch again but rather just flows into the lowest of several possible riverbeds.




Figure 2: The *River* diagram \mathcal{M}^{RV} for the margin graph of Figure 1 with $RV(\mathbf{P}) = \{a\}$.

	SC	SV	RP	BP	RV
Anonymity	✓	✓*	✓*	✓	✓*
Neutrality	✓	✓*	✓*	✓	✓*
Monotonicity	✓	✗	✓	✓	✓
Condorcet Winner	✓	✓	✓	✓	✓
Condorcet Loser	✓	✓	✓	✓	✓
Smith criterion	✓	✓	✓	✓	✓
Pareto efficiency	✓	✓	✓	✓	✓
ISDA	✓	✓ ^o	✓	✓	✓*
IPDA	✗	✗	✗	✗	✓*
Indep. of clones	✓	?	✓*	✓	✓*

Table 1: Overview of the properties for Split Cycle and refinements. “✓” indicates the property is satisfied for general profiles. “✓*” indicates the property is satisfied for general profiles with a suitable tiebreaker (not necessarily the same for each). “✓^o” indicates the property is satisfied for uniquely weighted profiles. Finally, “✗” indicates the property is violated even for uniquely weighted profiles. For all results not concerning *River* or IPDA, we refer to Holliday and Pacuit (2023a,b).

contain all $\binom{m}{2}$ edges. Note also that *River*, Ranked Pairs and Split Cycle clearly justify their choice through the respective diagrams, unlike Stable Voting and Beat Path.

River, like Ranked Pairs and Stable Voting, requires a tiebreaker to be resolute for non-uniquely weighted preference profiles, and is resolute for uniquely weighted profiles.

4 Axiomatic Properties

In this section, we analyse the axiomatic properties of *River* and compare them to the other introduced social choice functions. An overview of our results is provided in Table 1.

We start by observing that *River* satisfies all basic axioms required of a reliable social choice function. For formal definitions and proofs, we refer to the full version of this paper.

Anonymity and *neutrality* require that all voters, respectively alternatives, are treated equally. Both properties are naturally satisfied by *River* for uniquely weighted preference profiles. In profiles that are not uniquely weighted, however,

the tiebreaker might invalidate anonymity or neutrality.⁶

Monotonicity demands that if support for a winning alternative increases (i. e., some voters rank it higher without changing the relative order of the other alternatives), this alternative must remain a winner. River satisfies monotonicity, as do Split Cycle, Ranked Pairs, and Beat Path. Maybe surprisingly, Stable Voting violates monotonicity.

River, like the other social choice functions, always selects the Condorcet winner if one exists, and never selects a Condorcet loser. Recall the *Smith set* (Section 2.1) as a generalization of a Condorcet winner. The *Smith criterion* states that the set of winners has to be chosen from the Smith set. This property is implied by selecting only immune alternatives and thus satisfied by all Split Cycle refinements.

Given a preference profile \mathbf{P} and two alternatives $x, y \in A$, y *Pareto-dominates* x if every voter ranks y over x , i.e., $y \succ_i x$ for all $i \in N$. *Pareto efficiency* requires that Pareto-dominated alternatives are never chosen. Since Split Cycle satisfies Pareto efficiency, so do all its refinements.

4.1 Independence of Smith-Dominated Alternatives

Next, we consider independence criteria. These properties prescribe that the winning set of a social choice function should not be affected by the presence (or absence) of alternatives that are in some sense “inferior”. Independence properties can also be interpreted as safeguards against agenda manipulation: introducing an inferior alternative into an election should not disturb the set of chosen candidates.

We begin with independence of Smith-dominated alternatives. Recall that the Smith set is the smallest dominating subset of alternatives. Every alternative *not* in the Smith set is called *Smith-dominated*. Independence of Smith-dominated alternatives requires that the winners do not change when a Smith-dominated alternative is removed.

Definition 4.1. A social choice function F is independent of Smith-dominated alternatives (*ISDA*) if for any preference profile \mathbf{P} and $x \in A \setminus Sm(\mathbf{P})$, we have $F(\mathbf{P}) = F(\mathbf{P}_{-x})$.

Split Cycle, Ranked Pairs, and Beat Path satisfy ISDA (Holliday and Pacuit 2023a). Stable Voting satisfied ISDA for uniquely weighted profiles, but violates it in general (Holliday and Pacuit 2023b). We show that River satisfies ISDA.

Theorem 4.2. River satisfies ISDA for general preference profiles when equipped with a consistent tiebreaker.

Proof sketch. Let $x \in A \setminus Sm(\mathbf{P})$. By definition of the Smith set, we have: (i) $S := Sm(\mathbf{P}) = Sm(\mathbf{P}_{-x})$; (ii) no alternative in $A \setminus S$ has a majority edge to any alternative in S ; and (iii) no majority cycle contains vertices from both S and $A \setminus S$.

Hence, no edge between vertices of S can be rejected by River because of edges involving $A \setminus S$ (neither in \mathbf{P} nor in \mathbf{P}_{-x}): by (ii), there is no majority edge from

⁶Ranked Pairs and Stable Voting encounter the same issue. Generally, anonymity and neutrality conflict with resoluteness; e.g., consider a profile with $A_{\mathbf{P}} = \{x, y\}$ and $m_{\mathbf{P}}(x, y) = 0$.

$A \setminus S$ to S that could trigger a (*Br*) rejection, and by (iii), there is no majority cycle involving a vertex from $A \setminus S$ that could trigger a (*Cy*) rejection. Formally, one can show by induction over the margin edges that for any $s_1, s_2 \in S$, $(s_1, s_2) \in E(\mathcal{M}^{RV}(\mathbf{P}))$ if and only if $(s_1, s_2) \in E(\mathcal{M}^{RV}(\mathbf{P}_{-x}))$. \square

4.2 Independence of Pareto-Dominated Alternatives

Next, we consider alternatives that are “inferior” because they are *Pareto-dominated*. The resulting property can be defined analogously to ISDA.⁷ Independence of Pareto-dominated alternatives was already studied by Fishburn (Fishburn 1973) under the name “reduction condition.”

Definition 4.3. A social choice function F is independent of Pareto-dominated alternatives (*IPDA*) if for any preference profile \mathbf{P} and $x, y \in A_{\mathbf{P}}$ with $m_{\mathbf{P}}(y, x) = |N_{\mathbf{P}}|$, we have $F(\mathbf{P}) = F(\mathbf{P}_{-x})$.

Observe that constructing a Pareto-dominated alternative is not a complex endeavour. One can choose any of the current alternatives as a blueprint and construct a copy alternative which is worse for everyone.

Surprisingly, Split Cycle, Ranked Pairs, Beat Path and Stable Voting all violate IPDA, even for uniquely weighted preferences with at most 5 alternatives.⁸

Theorem 4.4. Split Cycle, Ranked Pairs, Beat Path and Stable Voting do not satisfy IPDA, not even for uniquely weighted profiles.

Proof. We present the counterexamples as margin graphs in Figure 3. The corresponding preference profiles $\mathbf{P}_1, \mathbf{P}_2, \mathbf{P}_3$, and \mathbf{P}_4 can be found in the full version of this paper. In all profiles, b Pareto-dominates a .

Split Cycle: Observe that $c \in SC(\mathbf{P}_1)$, because its only incoming edge (d, c) is the smallest-margin edge of cycle (c, a, d, c) . However, $c \notin SC((\mathbf{P}_1)_{-a})$, because (c, a, d, c) was the only cycle that (d, c) was contained in.

Ranked Pairs: Observe that $RP(\mathbf{P}_2) = \{d\}$. Ranked Pairs adds (b, a) , (b, c) , (a, c) , (d, a) , skips (c, d) because of the cycle (c, d, a, c) , adds (d, a) , and skips (d, b) . However, $RP((\mathbf{P}_2)_{-a}) = \{b\}$, because without a , Ranked Pairs adds (b, c) and (c, d) , skipping (d, b) .

Beat Path: It can be verified by computing the beatpath strengths that $BP(\mathbf{P}_3) = \{d\}$, but $BP((\mathbf{P}_3)_{-a}) = \{c\}$.

Stable Voting: It can be verified that $SV(\mathbf{P}_4) = \{d\}$, but $SV((\mathbf{P}_4)_{-a}) = \{c\}$. \square

River satisfies IPDA due to its unique way of resolving cycles in the margin graph. We prove this claim for uniquely weighted preference profiles in Theorem 4.6, and extend it in Theorem 4.8 to general profiles.

⁷IPDA is logically independent from ISDA, because neither does Pareto-dominance imply Smith-dominance, nor vice versa (Fishburn 1973).

⁸The example for Stable Voting is due to Wesley Holliday (personal communication, 2024).




Figure 3: Margin graphs for the preference profiles used in the proof of Theorem 4.4.

To build towards the first result, we start with the following observation related to the concept of covering: An alternative y is said to *cover* another alternative x if $m(y, x) > 0$ and $m(y, z) \geq m(x, z)$ for all $z \in A \setminus \{x, y\}$. We observe that Pareto domination implies covering.

Observation 4.5. *Given a preference profile \mathbf{P} , if y Pareto-dominates x , then y covers x .*

Intuitively, if y Pareto-dominates x , then $m(y, x) = n$ and y is ranked above x by every voter. Thus, whenever $x \succ_i z$, we also have $y \succ_i x \succ_i z$. This observation will be instrumental in the proof of the following claim.

Theorem 4.6. *River satisfies IPDA for uniquely weighted preference profiles.*

Proof. Let \mathbf{P} be a uniquely weighted preference profile and let $x, y \in A_{\mathbf{P}}$ such that y Pareto-dominates x . Let $E = E(\mathcal{M}^{RV}(\mathbf{P}))$ and $E_{-x} = E(\mathcal{M}^{RV}(\mathbf{P}_{-x}))$ denote the edge sets of the River diagrams with and without x , respectively.

First, we show that $(y, x) \in E$. Since $m(y, x) = n$ and the preference profile is uniquely weighted, the edge (y, x) is processed as the first edge. Hence, adding (y, x) can neither form a cycle nor branching, and (y, x) is added to E .

Next, we show that $(x, z) \notin E$ for all $z \in A \setminus \{x\}$. Assume towards contradiction that (x, z) is added to E . Then (Br) was not fulfilled, and thus $(y, z) \notin E$. Since y covers x and \mathbf{P} is uniquely weighted, $m(y, z) > m(x, z)$ and (y, z) is processed before (x, z) . If (y, z) was not added because of (Br) , i.e., there is some $(z', z) \in E$ with higher margin, then (Br) would also reject (x, z) . This would be a contradiction. Therefore, (y, z) must be rejected by (Cy) . This means, there must be a path $(z, p_1, \dots, p_k, y) \in E$ with strength larger than $m(y, z)$. But this path would form a cycle with (y, x) and (x, z) : $(z, p_1, \dots, p_k, y, x, z)$. Thus, (Cy) would reject (x, z) , a contradiction. We conclude that $(x, z) \notin E$.

So in \mathbf{P} , no edge adjacent to x is considered in E apart from (y, x) , which cannot be part of any cycle, since x has no outgoing edges in E . Removing x from the election does not influence the margin of any edge not containing x , and all edges incident to x are simply removed. Therefore, in E_{-x} , we have an edge $(z, z') \in E_{-x}$ if and only if $(z, z') \in E$.

It follows that all edges in E apart from (y, x) are in E_{-x} , and thus $RV(\mathbf{P}) = RV(\mathbf{P}_{-x})$. \square

For non-uniquely weighted preference profiles, River satisfies IPDA when equipped with any tiebreaker that “respects” Pareto dominance.

Definition 4.7. *A consistent tiebreaker is called Pareto-consistent if, whenever y Pareto-dominates x , for all $z \notin \{x, y\}$, the edge (y, z) is ranked higher than the edge (x, z) .*

Note that Pareto-consistent tiebreaker always exists and can be constructed in a straightforward way, along the lines described in Section 2.3: Fix any voter’s preference order and sort the edges lexicographically according to that order. Since the voter’s preferences respect Pareto dominance (by definition of the Pareto-dominance relation), the resulting ranking of edges is Pareto-consistent.

Theorem 4.8. *River satisfies IPDA when equipped with a Pareto-consistent tiebreaker.*

In the full version of this paper, we also show that we cannot drop the Pareto-consistency requirement by providing an example where River fails to satisfy IPDA for a tiebreaker that is not Pareto-consistent.

We can extend Theorem 4.8 by considering a more permissive notion of dominance that we call *quasi-Pareto dominance*. Instead of requiring that y is preferred over x by every voter (i.e., $m(y, x) = n$), we merely require that y covers x and that the margin of y over x is no weaker than any other margin involving x .

Definition 4.9. *Let \mathbf{P} be a profile and $x, y \in A_{\mathbf{P}}$. We say y quasi-Pareto-dominates x if (1) y covers x and (2) for all $z \in A_{\mathbf{P}} \setminus \{x, y\}$, $m(y, x) \geq m(z, x)$ and $m(y, x) \geq m(x, z)$.*

It is easy to check that Pareto domination implies quasi-Pareto domination and that the latter is an acyclic relation. This leads to a strengthened version of IPDA, namely *independence of quasi-Pareto-dominated alternatives* (IQDA). We show that River satisfies IQDA when it is equipped with quasi-Pareto-consistent tiebreakers, which are defined analogously to Pareto-consistent tiebreakers. The formal definition of such tiebreakers and the proof of the following theorem can be found in the full version of this paper.

Theorem 4.10. *River satisfies IQDA for uniquely weighted preference profiles, and for general profiles when equipped with a quasi-Pareto-consistent tiebreaker.*

4.3 Independence of Clones

The concept of clones was introduced by Tideman (1987) to model “nearly identical” alternatives.

Definition 4.11. Let \mathbf{P} be a preference profile. A set $C \subseteq A$ with $2 \leq |C| < m$ is a set of clones in \mathbf{P} , if for all $c, c' \in C$, $x \in A \setminus C$ and $i \in N$ it holds that $c \succ_i x \Leftrightarrow c' \succ_i x$ and $x \succ_i c \Leftrightarrow x \succ_i c'$. We denote the set of all clone sets in \mathbf{P} by $\mathbb{C}(\mathbf{P})$.

Given the notion of clones, Tideman proposed the notion of *independence of clones* (IoC) to formalize the idea that clones should not hinder each other's performance in an election, nor should adding or removing a clone affect unrelated alternatives.

Definition 4.12. An SCF F satisfies independence of clones (IoC) if for every profile \mathbf{P} with a set of clones C and $c \in C$:

1. $C \cap F(\mathbf{P}) \neq \emptyset \Leftrightarrow C \setminus \{c\} \cap F(\mathbf{P}_{-c}) \neq \emptyset$, and
2. for all $x \in A \setminus C$: $x \in F(\mathbf{P}) \Leftrightarrow x \in F(\mathbf{P}_{-c})$.

Split Cycle and Beat Path satisfy IoC for general preference profiles (Holliday and Pacuit 2023a), whereas it is unknown whether Stable Voting satisfies IoC. For Ranked Pairs, Zavist and Tideman (1989) have shown that IoC is satisfied as long as it is equipped with tiebreaker that satisfies a property they refer to as “impartiality.”

River and independence of clones. We show an analogous result (albeit with a more complicated proof): River satisfies independence of clones when equipped with a suitable tiebreaker function. The required property is defined in the following definition.

Definition 4.13. A tiebreaker \prec_A is called clone-consistent if it is consistent and for every clone set $C \subseteq A$ the following holds: (1) for all $x \in A$, the set of edges $\{(x, y) : y \in A\}$ and the set of edges $\{(x, c) : c \in C\}$ occur as contiguous blocks in \prec_A ; and (2) the set of edges $\{(c, y) : c \in C, y \in A\}$ occurs as a contiguous block in \prec_A .

Like in the case of Pareto-consistency (Definition 4.7), a clone-consistent tiebreaker can be obtained by lexicographically sorting the edges according to the preference order of a given voter. Therefore, every tiebreaker function that uses the preferences \succ_i of a fixed voter $i \in N$ as the basis for lexicographically ranking edges, always returns tiebreakers that are both Pareto-consistent and clone-consistent.

Intuitively, a clone-consistent tiebreaker ensures that all edges between clones in C and other alternatives $x \in A \setminus C$ are considered by River consecutively. As a result, River treats the clones of C as a single unit, and the edges added to the River diagram remain structurally the same even when adding or removing a clone member. Due to space constraints, we provide only a proof sketch and refer to the full version of this paper for a precise treatment.

Theorem 4.14. River satisfies independence of clones for general preference profiles when equipped with a clone-consistent tiebreaker.

Proof sketch. During the construction process, the River diagram is always a disjoint union of trees. It is sufficient to only keep track of those trees' node sets B and roots $w(B)$: (x, y) is added iff $B(x) \neq B(y)$ and $w(B(y)) = y$. Given a clone set $C = \{\gamma, \dots\}$, one can show that at all times, either (i) one B has $B \cap C \neq \emptyset$ and $w(B) \notin C$ and all other

have $B \cap C = \emptyset$, or (ii) each B has either $B \cap C = \emptyset$ or $w(B) \in C$; and at most times, this structure is essentially the same whether γ is (a) present or (b) not. Because of the edge ordering implied by the clone-consistent tiebreaker, scenarios (a) and (b) can only deviate temporarily when edges of the same margin are processed. In the end, case (i) means the same non-clone wins in (a) and (b), and case (ii) means some (possibly different) clone wins in (a) and (b). \square

That clone edges are processed “en suite” is crucial, see the counterexample for non-clone-consistent tiebreakers in the full version of this paper.

Composition consistency Inspired by the recent paper by Berker et al. (2025), we investigated whether River satisfies a strengthening of IoC known as *composition consistency* (Laffond, Lainé, and Laslier 1996). We found that this is not the case.

Theorem 4.15. River does not satisfy composition consistency, not even with a clone-consistent tiebreaker.

For the definition of composition consistency and the proof of Theorem 4.15, see the full version of this paper.

5 Discussion

We introduced *River*, a novel single-winner voting method based on majority margins that constructs a spanning tree of the margin graph rooted in the winner. A key feature of River is its simplicity in both explanation and manual computation. The spanning tree structure provides a clear and easy-to-interpret certificate for the winner.

River only selects immune alternatives and is, therefore, a refinement of Split Cycle. Like other refinements of Split Cycle, River satisfies desirable axiomatic properties, including Condorcet-consistency, independence of Smith-dominated alternatives, and independence of clones (with suitable tiebreaking). It also satisfies independence of Pareto-dominated alternatives, which is not satisfied by any of the other methods we discussed.

River is similar to Ranked Pairs in that it builds a directed graph that “explains” the selection of the winner. The advantage of River over Ranked Pairs is that its explanation is more concise and easier to understand. The simplicity of River also has computational benefits: As Döring, Malanowski, and Neubert (2026) have recently shown, River winners can be efficiently computed even under parallel-universes tiebreaking (PUT), a neutral way of dealing with ties that renders Ranked Pairs computationally intractable (Brill and Fischer 2012). These advantages are a direct consequence of River's tree structure.

Several open questions remain. For instance, it would be interesting to check whether River satisfies axiomatic properties beyond the ones studied in this paper and whether one can design tiebreakers that make it composition-consistent. Moreover, a comprehensive *empirical* analysis would reveal how River compares to other voting rules on realistic preference profiles. Preliminary simulations suggest that River may be viewed as mediating between Ranked Pairs and Beat Path since it agrees with each of them more often than they agree with each other.

Acknowledgments

We would like to thank the anonymous reviewers for their extraordinarily detailed feedback and many helpful suggestions, which significantly improved this paper.

This work was supported by a Structural Democracy Fellowship through the Brooks School of Public Policy at Cornell University and by the Federal Ministry of Research, Technology and Space under the funding code “KI-Servicezentrum Berlin-Brandenburg” 16IS22092. Responsibility for the content of this publication remains with the authors.

References

- Arrow, K. J.; Sen, A. K.; and Suzumura, K., eds. 2002. *Handbook of Social Choice and Welfare*, volume 1. North-Holland.
- Askell, A.; Bai, Y.; Chen, A.; Drain, D.; Ganguli, D.; Henighan, T.; Jones, A.; Joseph, N.; Mann, B.; DasSarma, N.; et al. 2021. A general language assistant as a laboratory for alignment. *arXiv preprint arXiv:2112.00861*.
- Berker, R. E.; Casacuberta, S.; Robinson, I.; Ong, C.; Conitzer, V.; and Elkind, E. 2025. From Independence of Clones to Composition Consistency: A Hierarchy of Barriers to Strategic Nomination. In *Proceedings of the 26th ACM Conference on Economics and Computation*. ACM.
- Brandl, F.; and Peters, D. 2022. Approval voting under dichotomous preferences: A catalogue of characterizations. *Journal of Economic Theory*, 205: 105532.
- Brandt, F.; Brill, M.; and Harrenstein, P. 2016. Tournament Solutions. In Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D., eds., *Handbook of Computational Social Choice*, chapter 3. Cambridge University Press.
- Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A., eds. 2016. *Handbook of Computational Social Choice*. Cambridge University Press.
- Brill, M.; and Fischer, F. 2012. The Price of Neutrality for the Ranked Pairs Method. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI)*, 1299–1305. AAAI Press.
- Chebotaev, P. Y.; and Shamis, E. 1998. Characterizations of scoring methods for preference aggregation. *Annals of Operations Research*, 80: 299–332.
- Ching, S. 1996. A simple characterization of plurality rule. *Journal of Economic Theory*, 71(1): 298–302.
- Conitzer, V.; Rognlie, M.; and Xia, L. 2009. Preference functions that score rankings and maximum likelihood estimation. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*, 109–115. AAAI Press.
- Dung, P. M. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial intelligence*, 77(2): 321–357.
- Döring, M.; Malanowski, J.; and Neubert, S. 2026. Cost-Free Neutrality for the River Method. In *Proceedings of the 40th AAAI Conference on Artificial Intelligence (AAAI)*. Forthcoming.
- Fischer, F.; Hudry, O.; and Niedermeier, R. 2016. Weighted Tournament Solutions. In Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D., eds., *Handbook of Computational Social Choice*, chapter 4. Cambridge University Press.
- Fishburn, P. C. 1973. *The Theory of Social Choice*. Princeton University Press.
- Fürnkranz, J.; and Hüllermeier, E. 2003. Pairwise preference learning and ranking. In *Machine Learning: ECML 2003*, 145–156. Springer.
- Gonzalez, S.; Laruelle, A.; and Solal, P. 2019. Dilemma with approval and disapproval votes. *Social Choice and Welfare*, 53: 497–517.
- Good, I. J. 1971. A Note on Condorcet Sets. *Public Choice*, 10(1): 97–101.
- Greaves, H.; and Cotton-Barratt, O. 2023. A Bargaining-Theoretic Approach to Moral Uncertainty. *Journal of Moral Philosophy*.
- Heitzig, J. 2002. Social choice under incomplete, cyclic preferences. *arXiv preprint math/0201285*.
- Heitzig, J. 2004. Hello again – and a new method for you! Online posting on the election-methods mailing list. Archive URL: <http://lists.electorama.com/pipermail/election-methods-electorama.com/2004-April/078029.html>, accessed 2025-11-11.
- Holliday, W. H.; and Pacuit, E. 2023a. Split Cycle: a new Condorcet-consistent voting method independent of clones and immune to spoilers. *Public Choice*, 197: 1–62.
- Holliday, W. H.; and Pacuit, E. 2023b. Stable Voting. *Constitutional Political Economy*, 421–433.
- Köpf, A.; Kilcher, Y.; von Rütte, D.; Anagnostidis, S.; Tam, Z. R.; Stevens, K.; Barhoum, A.; Nguyen, D.; Stanley, O.; Nagyfi, R.; et al. 2024. Openassistant conversations-democratizing large language model alignment. *Advances in Neural Information Processing Systems*, 36.
- Laffond, G.; Lainé, J.; and Laslier, J.-F. 1996. Composition-Consistent Tournament Solutions and Social Choice Functions. *Social Choice and Welfare*, 13(1): 75–93.
- Laslier, J.-F. 1997. *Tournament Solutions and Majority Voting*. Springer-Verlag.
- May, K. 1952. A Set of Independent, Necessary and Sufficient Conditions for Simple Majority Decisions. *Econometrica*, 20(4): 680–684.
- Mishra, A. 2023. AI Alignment and Social Choice: Fundamental Limitations and Policy Implications. *arXiv preprint arXiv:2310.16048*.
- Öztürk, Z. E. 2020. Consistency of scoring rules: A reinvestigation of composition-consistency. *International Journal of Game Theory*, 49: 801–831.
- Richelson, J. 1978. A characterization result for the plurality rule. *Journal of Economic Theory*, 19(2): 548–550.
- Schulze, M. 2011. A new monotonic, clone-independent, reversal symmetric, and Condorcet-consistent single-winner election method. *Social choice and Welfare*, 36: 267–303.

- Schwartz, T. 1986. *The Logic of Collective Choice*. Columbia University Press.
- Smith, J. H. 1973. Aggregation of Preferences with Variable Electorate. *Econometrica*, 41(6): 1027–1041.
- Tideman, T. N. 1987. Independence of clones as a criterion for voting rules. *Social Choice and Welfare*, 4: 185–206.
- Wang, J.; Sikdar, S.; Sheperd, T.; Zhao, Z.; Jiang, C.; and Xia, L. 2019. Practical Algorithms for STV and Ranked Pairs with Parallel Universes Tiebreaking. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, 2189–2196.
- Zavist, T. M.; and Tideman, T. N. 1989. Complete independence of clones in the ranked pairs rule. *Social Choice and Welfare*, 6(2): 167–173.
- Zwicker, W. S. 2016. Introduction to the Theory of Voting. In Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D., eds., *Handbook of Computational Social Choice*, chapter 2. Cambridge University Press.