

Condorcet-Consistent Choice Among Three Candidates

Felix Brandt, Chris Dong, and Dominik Peters

Abstract

A voting rule is a Condorcet extension if it returns a candidate that beats every other candidate in pairwise majority comparisons whenever one exists. Condorcet extensions have faced criticism due to their susceptibility to variable-election paradoxes, especially the reinforcement paradox [53] and the no-show paradox [42]. In this paper, we investigate the susceptibility of Condorcet extensions to these paradoxes for the case of exactly three candidates. For the reinforcement paradox, we establish that it must occur for every Condorcet extension when there are *at least eight* voters and demonstrate that certain refinements of maximin—a voting rule originally proposed by Condorcet [14]—are immune to this paradox when there are *at most seven* voters. For the no-show paradox, we prove that the *only* homogeneous Condorcet extensions immune to it are refinements of maximin. We also provide axiomatic characterizations of maximin and two of its refinements, Nanson’s rule and leximin, highlighting their suitability for three-candidate elections.

1 Introduction

Deciding between two candidates based on the preferences of multiple voters is straightforward and allows for simple and natural rules that satisfy virtually all desirable properties, such as majority rule or weighted threshold rules. However, *choosing from three or more candidates* is what leads to significant challenges and inevitable tradeoffs. Arrow’s impossibility and the Gibbard-Satterthwaite theorem are the best known among many results highlighting these difficulties [1, 27, 49]. These findings help to explain the plethora of voting rules proposed for elections with three or more candidates. In this paper, we address whether identifying a suitable rule becomes easier when focussing on the case of *exactly three candidates*.

The primary reason why plurality rule—the most common voting rule—struggles with three or more candidates is that it only takes into account the voters’ top choices and thereby, for example, ignores what a voter who favors candidate *a* thinks about candidates *b* and *c*. Once this information is taken into account, Condorcet [14] argued that the best choice is the candidate that the majority prefers, in the sense that every other candidate is judged to be worse by some majority of voters. Nowadays, such a candidate is called a *Condorcet winner*. Choosing a Condorcet winner is motivated by the observation that any other candidate can be overthrown by a coordinated majority of voters who all prefer the same candidate to the selected candidate. However, Condorcet recognized that a Condorcet winner need not exist since the majority relation may be cyclic. This has given rise to extensive efforts designing *Condorcet extensions*, voting rules that select the Condorcet winner whenever it exists and use some other procedure to choose a candidate in the remaining cases. Dozens of Condorcet extensions have been cataloged [22, 35, 6]. Condorcet also proposed a Condorcet extension of his own, clearly described for the case of three candidates:

“When the three [pairwise majority] views cannot exist together [because of a cycle], the adopted view results from the two [pairwise majority views] that are most probable [i.e., have the largest majority].”
Condorcet [14, p. 125]

How this idea should be extended to four or more candidates is only outlined in vague terms and

ultimately remains elusive. However, his proposal for three candidates is unambiguous and coincides with what many modern Condorcet extensions do in this case, including maximin, ranked pairs, beat path, split cycle, Kemeny’s rule, Dodgson’s rule, and Young’s rule.

While compelling, Condorcet’s principle to always return the Condorcet winner has also met some resistance. The two main lines of attack concern settings with a variable set of voters. In particular, Young and Levenglick [53] showed that no Condorcet extension satisfies *reinforcement*, a consistency condition demanding that if the same winner is returned for two different electorates, it should also be returned for the union of these electorates. In a similar spirit, Moulin [42] demonstrated that every Condorcet extension suffers from the *no-show paradox*, where a voter might benefit by abstaining from an election to achieve a more desirable outcome. The proofs of Young and Levenglick and Moulin are based on constructing collections of preference profiles on which the paradoxes cannot be avoided. Moulin’s proof uses profiles with at least four candidates and 25 voters. Later, a smaller proof requiring only 12 voters was found, and no proof with fewer voters exists [8]. Young and Levenglick’s proof for the reinforcement paradox uses at least three candidates and 13 voters.

In this paper, we show that refinements of maximin—the rule proposed by Condorcet—perform particularly well with respect to these paradoxes when there are exactly three candidates. Specifically, for up to seven voters, we exhibit a refinement of maximin that is immune to the reinforcement paradox. During a June 2022 lecture at the Institut Henri Poincaré (“Mathématiques et Démocratie” conference), Hervé Moulin asked what the smallest number of voters is for which Young and Levenglick’s result holds. We use SAT solvers to answer this question, showing that it holds even for eight voters. This bound is tight, and thus certain maximin refinements avoid the reinforcement paradox as much as possible. Regarding the no-show paradox, it is known that maximin with a fixed tie-breaking order avoids the paradox for three candidates [42]. We prove that the *only* homogeneous Condorcet extensions immune to the no-show paradox are refinements of maximin. Moreover, by adding continuity, we obtain an axiomatic characterization of maximin. We also characterize two refinements of maximin (Nanson’s rule and leximin) using natural monotonicity and invariance axioms.

The rest of the paper is structured as follows. In Section 2, we introduce our model and standard axioms. In Section 3, we provide an overview of Condorcet extensions and their inclusion relationships when only three candidates exist. In Appendix A, we present several impossibility results involving reinforcement, each using the minimum possible number of voters. In Section 4, we show that maximin refinements are the only homogeneous Condorcet extensions that circumvent the no-show paradox and provide characterizations of maximin, Nanson’s rule, and leximin.

2 Preliminaries

Let $A = \{a, b, c\}$ be the set of three *candidates* and let N^* be the (possibly infinite) set of potential voters. A non-empty and finite set $N \subseteq N^*$ is called an *electorate*. Given an electorate N , a (*preference*) *profile* $P = (\succ_i)_{i \in N}$ over N is a collection of linear orders, with $\succ_i$ denoting the preference of voter $i \in N$. For example, if $a \succ_i b \succ_i c$, then voter i prefers a the most, followed by b , with c preferred least. We sometimes write abc as a shorthand for this linear order; in Figure 1(a) and in other figures we show linear orders as columns. We will write $x \succeq_i y$ if $x = y$ or $x \succ_i y$.


Let P be a profile. For two candidates $x, y \in A$, we define the *margin* of x over y as

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


and we just write $m_{x,y}$ if P is clear from the context. Thus, if $m_{x,y} > 0$, then a strict majority of voters prefer x to y . Note that $m_{x,y} = -m_{y,x}$ and that $m_{x,x} = 0$. We can display this information in a *margin graph*, which is a weighted digraph on vertex set A where for each pair $\{x, y\}$ of vertices, we orient the edge as $x \rightarrow y$ in such a way that $m_{x,y} > 0$ (so that x beats y in a pairwise comparison) and give it

γ_1	γ_2	γ_3	γ_4	γ_5
a	a	b	c	c
b	b	c	a	a
c	c	a	b	b

(a) Preference profile.



(b) Margin graph.



(c) Ordinal margin graph.

Figure 1: An example of a preference profile over $N = \{1, \dots, 5\}$ and its margin graph and ordinal margin graph. The ordinal margin graph encodes the information that $m_{a,b} > m_{b,c} = m_{c,a} > 0$.

weight $m_{x,y}$. For pairs with $m_{x,y} = 0$, we do not draw an edge. See Figure 1(b) for an example of a margin graph. For our purposes, it is often not necessary to know the precise weight of an edge, and in those cases, we can draw the *ordinal margin graph*, where we preserve the direction of each edge and show which edges have strictly higher weights than others (by drawing the higher-weight edge with more thickness) and which edges have equal weight (by drawing the edges with the same thickness). Sometimes we use an empty arrowhead for the thinnest edges. See Figure 1(c) for an example of an ordinal margin graph.

For any profile P over N , the margins $(m_{x,y})_{x,y \in A}$ all have the same parity (which is the same as the parity of $|N|$). The well-known *McGarvey theorem* states that for every margin graph with same-parity weights, there is a profile P that induces it [38, 17].

A *social choice function* f maps each profile P (on any electorate $N \subseteq N^*$) to a non-empty set $f(P) \subseteq A$ of winners. It is *resolute* if $|f(P)| = 1$ for all profiles P . A social choice function f_1 *refines* another social choice function f_2 if $f_1(P) \subseteq f_2(P)$ for all profiles P .

Let us define four basic properties of a social choice function. The first two are standard symmetry conditions.

- A social choice function f is *anonymous* if its output does not depend on the identity of the voters. Formally, for all profiles P over N and all injections $\pi: N \rightarrow N^*$, we have $f(P) = f(\pi(P))$ where $\pi(P)$ is the profile over $\pi(N)$ where voter $\pi(i)$ has preferences $\succ_i$.
- A social choice function f is *neutral* if its output does not depend on the identities of the candidates. Formally, for all profiles P and all bijections $\sigma: A \rightarrow A$, we have $\sigma(f(P)) = f(\sigma(P))$, where $\sigma(P)$ refers to the profile obtained from P by replacing each voter i 's preference $x \succ_i y \succ_i z$ with $\sigma(x) \succ_i \sigma(y) \succ_i \sigma(z)$.

The next two properties involve the operation of copying a profile, which is always possible, provided that the set of potential voters is infinite. Thus, we will only invoke these axioms when $N^* = \mathbb{N}$. Let P be a profile over N with $|N| = n$. We say that a profile Q over electorate N' with $|N'| = t \cdot n$ is a *t-fold copy* of P if we can partition N' into t disjoint subelectorates $N_1, \dots, N_t$, each of size n , such that the restricted profile Q_{N_j} is a copy of P for each $j = 1, \dots, t$, in the sense that there is a bijection $\phi: N \rightarrow N_j$ such that each voter $i \in N$ has the same preferences as $\phi(i)$. We usually write “ tP ” to denote a profile that is a t -fold copy of P .

- A social choice function f is *homogeneous* if doubling the profile does not affect the output. Formally, f is homogeneous if for all profiles P and $2P$ such that $2P$ is a 2-fold copy of P , we have $f(P) = f(2P)$.
- A social choice function f is *continuous* if for all profiles P and P' , there is some $n' \in \mathbb{N}$ such that for all $n \geq n'$ it holds that $f(nP + P') \subseteq f(P)$ whenever nP is an n -fold copy of P defined on an electorate disjoint from P' .

Young [52] introduced the continuity axiom for anonymous social choice functions, with the intuition that overwhelmingly large groups of voters should decide the outcome, while negligible, arbitrarily small groups should only be able to break ties. It is also known as the “overwhelming majorities axiom” [43].

3 Condorcet Extensions

We say that a candidate $x \in A$ is a *Condorcet winner* if $m_{x,y} > 0$ for all $y \in A \setminus \{x\}$, which means that x beats every other candidate in a pairwise majority comparison. A Condorcet winner may fail to exist (see Figure 1), but if it exists, it has to be unique. A social choice function f is a *Condorcet extension* or *Condorcet-consistent* if $f(P) = \{x\}$ whenever x is a Condorcet winner.

Numerous Condorcet extensions have been defined, though for the case of three candidates, many of these rules coincide. Of particular interest is the *maximin rule*, which selects the candidates whose worst margin is highest:

$$f_{\text{maximin}}(P) = \arg \max_{x \in A} \min_{y \in A \setminus \{x\}} m_{x,y}.$$

Thus, the winners under this rule never lose too badly against another candidate. It is sometimes also known as the *minimax* rule (since it minimizes the worst loss) or the *Simpson-Kramer* rule.

To define some other social choice functions, we need the concept of the *Borda score* of a candidate, defined as $\beta_x = \sum_{y \in A} m_{x,y}$. This is the net out-degree of x in the margin graph. The *Borda rule* selects the candidates with the highest Borda score.¹ Note that a Condorcet winner x , if one exists, has a strictly positive Borda score (because $m_{x,y} > 0$ for all $y \neq x$). In addition, because $m_{x,y} = -m_{y,x}$, we have $\sum_{x \in A} \beta_x = 0$, so the average Borda score of a candidate is 0, and thus a Condorcet winner always has an above-average Borda score. However, it need not have the *highest* Borda score. Still, Borda scores can be used to define some interesting Condorcet extensions:

- The Black [4] rule is the rule that selects the Condorcet winner if one exists and otherwise returns the candidates with the highest Borda score.
- The Nanson [44] rule repeatedly deletes all candidates whose Borda score is not positive. Borda scores are computed with respect to the remaining candidates $A' \subseteq A$, and the process continues until there is no candidate x with positive Borda score $\sum_{y \in A'} m_{x,y}$.
- The *leximin rule* selects candidates whose worst margin is highest, just like maximin. However, when there are several such candidates, it breaks the tie in favor of those with the higher second-worst margin, breaking any remaining ties using the third-worst margin, and so on.² This is a natural way to refine maximin, though it does not appear to have been studied in the academic literature.³ In the case of three candidates, the leximin rule is equivalent to the maximin rule with ties broken in favor of candidates with higher Borda score, a rule that was also discussed by Holliday and Pacuit [33, Appendix A]. To see this, note that if two candidates x and y have the same worst margin, then x has a higher second-worst margin than y if and only if the sum of these two margins is higher for x than for y .

¹For the standard, equivalent definition, see Appendix A.3

²Formally, for $x \in A$, let $m^x = (m_1^x, \dots, m_m^x)$ be a reordered vector of the majority margins $(m_{x,y})_{y \in A}$ with $m_1^x \leq \dots \leq m_m^x$. For $x, y \in A$ we write $m^x \triangleright m^y$ if and only if there is some $r \leq m$ such that $m_i^x = m_i^y$ for all $i < r$ and $m_r^x > m_r^y$. For a given profile, leximin returns the set of all x such that there is no $y \in A$ with $m^y \triangleright m^x$. We call $\triangleright$ the *lexicographic ordering*. It forms a linear order over $\mathbb{Z}^m$, so for leximin to return multiple candidates, these must have the same ordered vector of majority margins.

³There have been discussions of this rule on the election-methods mailing list (2010, 2011), and a preprint finds that leximin is more frequently resolute than other maximin refinements in simulations [16].




Figure 2: Hasse diagram of three-candidate Condorcet extensions where lower rules refine higher ones, and rules at the same node are identical. Copeland refers to Copeland⁰.

There are many other commonly studied Condorcet extensions. For brevity, we will not include definitions of all the rules we mention below since they can be found elsewhere (18, for variants of the uncovered set (UC); 30, for top cycle, Llull, Copeland⁰ (which we henceforth call Copeland), Baldwin, strict Nanson, ranked pairs, beat path, split cycle; 13, for Dodgson and Young; 32, for stable voting; 29 or Section 4 below, for the defensible set).

We will now study how all these Condorcet extensions relate to each other in the case of three candidates. In particular, we will be interested in which rules are equivalent (i.e., select the same output for all profiles) and which rules refine which other rules. Figure 2 shows a Hasse diagram of the refinement relation of these Condorcet extensions. Rules that appear in the same node of the diagram are equivalent. Note, in particular, the large cluster of equivalent Condorcet extensions in the middle: maximin is equivalent to split cycle, ranked pairs, beat path, as well as the rules of Kemeny, Dodgson, and Young. All omitted proofs can be found in the appendix.

Theorem 1. *For three candidates, Figure 2 shows the Hasse diagram of the refinement relation between select Condorcet extensions.*

In the absence of majority ties (i.e., when only Graphs A, C, G, I, and K of Table 1 are considered), even more rules coincide. The top cycle, all variants of the uncovered set, Schwartz, and Copeland are equivalent because they return all candidates when the majority graph is a three-cycle. Moreover, stable voting coincides with maximin, and Nanson’s rule coincides with leximin. When restricting attention to the generic case where, on top of the absence of majority ties, no two margins are equal

	Graph A	Graph B	Graph C	Graph D	Graph E	Graph F	Graph G	Graph H	Graph I	Graph J	Graph K	Graph L
	$m_{a,b}$ $= m_{b,c}$ $= m_{c,a}$ > 0	$m_{a,b}$ $= m_{b,c}$ $= m_{c,a}$ $= 0$	$m_{a,b}$ $> m_{b,c}$ $= m_{c,a}$ > 0	$m_{a,b}$ $> m_{b,c}$ $= m_{c,a}$ $= 0$	$m_{a,b}$ $= m_{c,b}$ $> m_{c,a}$ $= 0$	$m_{a,b}$ $> m_{c,b}$ $> m_{c,a}$ $= 0$	$m_{a,b}$ $> m_{b,c}$ $> m_{c,a}$ > 0	$m_{a,b}$ $> m_{b,c}$ $> m_{c,a}$ $= 0$	$m_{a,b}$ $= m_{b,c}$ $> m_{c,a}$ > 0	$m_{a,b}$ $= m_{b,c}$ $> m_{c,a}$ $= 0$	$m_{b,c}$ $> m_{a,b}$ $> m_{c,a}$ > 0	$m_{b,c}$ $> m_{a,b}$ $> m_{c,a}$ $= 0$
top cycle	$\{a, b, c\}$	$\{a, b, c\}$	$\{a, b, c\}$	$\{a, b, c\}$	$\{a, c\}$	$\{a, c\}$	$\{a, b, c\}$	$\{a, b, c\}$	$\{a, b, c\}$	$\{a, b, c\}$	$\{a, b, c\}$	$\{a, b, c\}$
UC McKelvey	$\{a, b, c\}$	$\{a, b, c\}$	$\{a, b, c\}$	$\{a, c\}$	$\{a, c\}$	$\{a, c\}$	$\{a, b, c\}$	$\{a, b, c\}$	$\{a, b, c\}$	$\{a, b, c\}$	$\{a, b, c\}$	$\{a, b, c\}$
UC Bordes Banks	$\{a, b, c\}$	$\{a, b, c\}$	$\{a, b, c\}$	$\{a, c\}$	$\{a, c\}$	$\{a, c\}$	$\{a, b, c\}$	$\{a, b\}$	$\{a, b, c\}$	$\{a, b\}$	$\{a, b, c\}$	$\{a, b\}$
UC Gillies	$\{a, b, c\}$	$\{a, b, c\}$	$\{a, b, c\}$	$\{a, c\}$	$\{a, c\}$	$\{a, c\}$	$\{a, b, c\}$	$\{a, c\}$	$\{a, b, c\}$	$\{a, c\}$	$\{a, b, c\}$	$\{a, c\}$
defensible set	$\{a, b, c\}$	$\{a, b, c\}$	$\{a, c\}$	$\{a, c\}$	$\{a, c\}$	$\{a, c\}$	$\{a, c\}$	$\{a, c\}$	$\{a, c\}$	$\{a, c\}$	$\{a\}$	$\{a\}$
Llull	$\{a, b, c\}$	$\{a, b, c\}$	$\{a, b, c\}$	$\{a, c\}$	$\{a, c\}$	$\{a, c\}$	$\{a, b, c\}$	$\{a\}$	$\{a, b, c\}$	$\{a\}$	$\{a, b, c\}$	$\{a\}$
Schwartz	$\{a, b, c\}$	$\{a, b, c\}$	$\{a, b, c\}$	$\{a, c\}$	$\{a, c\}$	$\{a, c\}$	$\{a, b, c\}$	$\{a\}$	$\{a, b, c\}$	$\{a\}$	$\{a, b, c\}$	$\{a\}$
UC Fishburn	$\{a, b, c\}$	$\{a, b, c\}$	$\{a, b, c\}$	$\{a, c\}$	$\{a, c\}$	$\{a, c\}$	$\{a, b, c\}$	$\{a\}$	$\{a, b, c\}$	$\{a\}$	$\{a, b, c\}$	$\{a\}$
Copeland	$\{a, b, c\}$	$\{a, b, c\}$	$\{a, b, c\}$	$\{a\}$	$\{a, c\}$	$\{a, c\}$	$\{a, b, c\}$	$\{a\}$	$\{a, b, c\}$	$\{a\}$	$\{a, b, c\}$	$\{a\}$
maximin	$\{a, b, c\}$	$\{a, b, c\}$	$\{a, c\}$	$\{a, c\}$	$\{a, c\}$	$\{a, c\}$	$\{a\}$	$\{a\}$	$\{a\}$	$\{a\}$	$\{a\}$	$\{a\}$
ranked pairs	$\{a, b, c\}$	$\{a, b, c\}$	$\{a, c\}$	$\{a, c\}$	$\{a, c\}$	$\{a, c\}$	$\{a\}$	$\{a\}$	$\{a\}$	$\{a\}$	$\{a\}$	$\{a\}$
beat path	$\{a, b, c\}$	$\{a, b, c\}$	$\{a, c\}$	$\{a, c\}$	$\{a, c\}$	$\{a, c\}$	$\{a\}$	$\{a\}$	$\{a\}$	$\{a\}$	$\{a\}$	$\{a\}$
split cycle	$\{a, b, c\}$	$\{a, b, c\}$	$\{a, c\}$	$\{a, c\}$	$\{a, c\}$	$\{a, c\}$	$\{a\}$	$\{a\}$	$\{a\}$	$\{a\}$	$\{a\}$	$\{a\}$
strict Nanson	$\{a, b, c\}$	$\{a, b, c\}$	$\{c\}$	$\{a, c\}$	$\{a, c\}$	$\{a, c\}$	$\{a\}$	$\{a\}$	$\{a\}$	$\{a\}$	$\{a\}$	$\{a\}$
stable voting	$\{a, b, c\}$	$\{a, b, c\}$	$\{a, c\}$	$\{a\}$	$\{a, c\}$	$\{a\}$	$\{a\}$	$\{a\}$	$\{a\}$	$\{a\}$	$\{a\}$	$\{a\}$
Nanson	$\{a, b, c\}$	$\{a, b, c\}$	$\{a\}$	$\{a\}$	$\{a, c\}$	$\{a, c\}$	$\{a\}$	$\{a\}$	$\{a\}$	$\{a\}$	$\{a\}$	$\{a\}$
leximin	$\{a, b, c\}$	$\{a, b, c\}$	$\{a\}$	$\{a\}$	$\{a, c\}$	$\{a\}$	$\{a\}$	$\{a\}$	$\{a\}$	$\{a\}$	$\{a\}$	$\{a\}$

Table 1: The outputs of rules that only depend on the ordinal margin graph.

(called *uniquely weighted* profiles by Holliday and Pacuit [31]), maximin is resolute and hence coincides with *all* its refinements. In Table 1, only Graph G and Graph K are generic in this sense, and we can see that the rules mentioned in the table collapse to a hierarchy of only three Condorcet extensions: the top cycle, the defensible set, and maximin.

4 Participation

Throughout this section, we assume that $N^* = \mathbb{N}$, so the set of potential voters is infinite. This means that it is always possible to copy a profile, and thus, the homogeneity and continuity axioms apply.

The participation axiom demands that voters are never better off by abstaining from an election. For a preference profile $P = (\succ_i)_{i \in N}$ and voter $j \in N$, we write $P_{-j} = (\succ_i)_{i \in N \setminus \{j\}}$ for the profile obtained from P by deleting j . For resolute rules, the definition of the participation axiom is straightforward: voter j should weakly prefer the outcome when j is present to the outcome when j is absent.

Definition 1 (Resolute participation, 42). A resolute social choice function satisfies *resolute participation* if for all profiles P and $j \in N$ with $f(P) = \{x\}$ and $f(P_{-j}) = \{y\}$, we have $x \succeq_j y$.

Moulin [42] has shown that when there are at least four candidates, then every resolute Condorcet extension violates resolute participation. This incompatibility holds for 12 or more voters [8]. On the other hand, Moulin [42] proved that for three candidates, maximin with a fixed tie-breaking order satisfies resolute participation.

For irresolute social choice functions, defining participation requires us to compare the sets of candidates $f(P)$ and $f(P_{-j})$ with respect to the preference relation $\succ_j$. This can be done in various ways. We will focus on the following well-studied variant that evaluates sets “optimistically”, that is, by comparing the

most-preferred candidates contained in each set. For a non-empty set $A' \subseteq A$ of candidates, we write $\max_{\succ_i} A'$ for the most-preferred candidate in A' according to $\succ_i$, so that $\max_{\succ_i} A' \succeq_i x$ for all $x \in A'$.

Definition 2 (Optimist participation). A social choice function satisfies *optimist participation* if for all profiles P and all $j \in N$, $\max_{\succ_j} f(P) \succeq_j \max_{\succ_j} f(P_{-j})$.

For four or more candidates, Jimeno et al. [34] showed that no Condorcet extension satisfies optimist participation [see, also, 8, Thm. 6]. For three candidates, maximin satisfies optimist participation. To see this, let P be a profile and $j \in N$ be a voter with preference $\succ_j$. Let g be the resolute social choice function obtained from maximin by breaking ties using $\succ_j$ as tie-breaking order. Then g satisfies resolute participation [42]. Thus,

$$\max_{\succ_j} f(P) = g(P) \succeq_j g(P_{-j}) = \max_{\succ_j} f(P_{-j}),$$

so that maximin satisfies optimist participation.⁴

To compare our results with existing results, let us define two additional participation-related conditions. A social choice function satisfies *positive involvement* [47] if whenever a voter joins a profile in which that voter's top candidate wins, then this candidate still wins. Formally, for all profiles P and all $j \in N$ with x the top-ranked candidate of j , we have that if $x \in f(P_{-j})$ then $x \in f(P)$. Pérez [47] also defined an analogous notion of *negative involvement*, which demands that if an agent's least-favorite candidate loses, then it still loses when the agent joins the electorate. We define an alternative version of this condition, *singleton negative involvement*, which says that if $f(P_{-j}) \neq \{x\}$ and x is bottom ranked by j , then $f(P) \neq \{x\}$.⁵ For three candidates, it turns out that the conjunction of positive involvement and singleton negative involvement is equivalent to optimist participation, which one can verify by case analysis on all possible combinations of sets $f(P_{-j}) \subseteq A$ and $f(P) \subseteq A$.

Proposition 1. *A social choice function satisfies optimist participation if and only if it satisfies positive involvement and singleton negative involvement.*

Pérez [46] noted that in order to satisfy positive involvement, a Condorcet extension must refine what Holliday [29] later called the defensible set. The *defensible set* of a profile P is

$$f_{\text{defensible}}(P) = \{x \in A : \text{for all } y \in A, \text{ there is } z \in A \text{ with } m_{z,y} \geq m_{y,x}\}.$$

We include a proof for completeness, noting that it is similar to an argument that was already used by Moulin [42, Claim (3)].⁶

Lemma 1 (Holliday [29], Lemma 2(1); Pérez [46], Lemma 3). *Let f be a Condorcet extension that satisfies positive involvement. Then f is a refinement of $f_{\text{defensible}}$, the defensible set.*

Proof. Let P be a profile. Let $x \notin f_{\text{defensible}}(P)$, so that there exists $y \in A$ with $m_{y,x} > m_{z,y}$ for all $z \in A$. Recall that all margins of a profile must have the same parity, so $m_{y,x} > m_{z,y}$ implies that we even have $m_{y,x} > m_{z,y} + 1$ for all $z \in A$. We need to show that $x \notin f(P)$. Assume for a contradiction that $x \in f(P)$. Add $m_{y,x} - 1$ voters with rankings $x \succ y \succ \dots$ to P to obtain the profile P^* . By positive involvement, x has to remain chosen by f after each addition of one voter, and thus, $x \in f(P^*)$. However, y is the Condorcet winner in P^* , because we have margins $m_{y,x}^* = 1$ and $m_{y,z}^* = -m_{z,y} - 1 + m_{y,x} > 0$. Thus, because f is a Condorcet extension, we have $x \notin f(P)$, a contradiction. $\square$

⁴In fact, maximin satisfies a stronger notion of participation based on the Fishburn preference extension [24]. In particular, writing $X = f_{\text{maximin}}(P)$ and $Y = f_{\text{maximin}}(P_{-j})$, it can be shown that we always have $X \succ_j Y \setminus X$ and $X \setminus Y \succ_j Y$, by applying standard results about the Fishburn preference extension (19, Theorem 3.4; 10, Section 3.3(ii)) and using the fact that maximin satisfies resolute participation for all ways of breaking ties using a fixed tie-breaking ordering.

⁵While similar in spirit, singleton negative involvement and negative involvement are logically incomparable. However, for resolute social choice functions, they are equivalent.

⁶The “converse” does not hold. Indeed, the strict Nanson rule and Baldwin's rule refine the defensible set (Theorem 1) but fail positive involvement even for three candidates: on the profile $2acb + 2bac + 2cba$, both rules select $\{a, b, c\}$; after adding a cab voter, we get the profile $2acb + 2bac + cab + 2cba$ where both rules select $\{a\}$, violating positive involvement.

4.1 Participation Requires Refining Maximin

As noted above, maximin satisfies optimist participation for three candidates. We will now show that maximin is, in fact, pre-eminent among Condorcet extensions having this property: all such Condorcet extensions must be refinements of maximin under the mild additional assumption of homogeneity.⁷

Theorem 2. *Let f be a homogeneous Condorcet extension that satisfies optimist participation. Then f is a refinement of maximin.*

Our proof uses [Lemma 1](#) to deduce that f must refine the defensible set. It will then use the extra strength of optimist participation compared to positive involvement (in particular, singleton negative involvement, see [Proposition 1](#)) to narrow down the space of candidates that f can select to conclude that f must refine maximin.

Proof. On profiles that induce a Condorcet winner, f refines maximin because it is a Condorcet extension. On profiles whose margin graph is a cycle with three equal non-negative margins, maximin selects all three candidates, and so f trivially refines maximin. All other profiles induce margin graphs that match, up to renaming the candidates, one of the four cases depicted in [Figure 3](#).⁸

The defensible set coincides with the maximin winners in Cases 1, 2, and 4. Because f satisfies positive involvement (by [Proposition 1](#)), it refines the defensible set by [Lemma 1](#), and therefore f refines maximin in Cases 1, 2, and 4. ([Figure 3](#) makes the argument of [Lemma 1](#) explicit by explaining what voters to add to arrive at a Condorcet winner and a contradiction to participation.)

It remains to consider Case 3, which concerns profiles P that, up to renaming candidates, induce a cycle with $0 \leq m_{c,a} < m_{b,c} \leq m_{a,b}$. In P , the unique maximin winner is a . Thus, our task is to show that $f(P) = \{a\}$. We will show this by connecting this case to Case 1. Because $f_{\text{defensible}}(P) = \{a, c\}$, we get that $b \notin f(P)$ from [Lemma 1](#). Assume for a contradiction that $c \in f(P)$. Let $P' = 2P$ be a profile consisting of two copies of P . By homogeneity, we obtain $c \in f(P')$. Let $j = (m'_{a,b} - m'_{b,c})/2 + 1$. Adding j voters to P' with preferences $b \succ c \succ a$ yields a profile P^* with margins

$$m_{c,a}^* = m'_{c,a} + j, \quad m_{a,b}^* = m'_{a,b} - j, \quad \text{and} \quad m_{b,c}^* = m'_{b,c} + j.$$

Crucially, we have $m'_{c,a} + 2 < m'_{b,c}$ and by choice of j we have $m_{c,a}^* < m_{a,b}^* < m_{b,c}^*$. Thus, P^* is a profile matching Case 1. Following our prior analysis, this implies $f(P^*) = \{a\}$. This is a contradiction to optimist participation because $\{a\}$ is a worse outcome for voters with preferences $b \succ c \succ a$ than an outcome containing c as a winner. (Note this step uses just singleton negative involvement.) $\square$

[Theorem 2](#) can be reinterpreted for resolute social choice functions, as optimist and resolute participation are equivalent for resolute social choice functions. Thus, we get the following corollary for resolute Condorcet extensions.

Corollary 1. *Let f be a homogeneous and resolute Condorcet extension that satisfies resolute participation. Then f is a refinement of maximin.*

⁷Our proof of [Theorem 2](#) actually uses the weaker “inclusion homogeneity” [[33](#), Definition 3.6] requiring only $f(P) \subseteq f(2P)$. In fact, we only need that for every profile P , there exists *at least one* profile $2P$ that is a 2-fold copy of P and such that $f(P) \subseteq f(2P)$. In the absence of anonymity, this is a weaker notion. Indeed, there are some non-anonymous refinements of maximin that satisfy participation and this weakened homogeneity, but not full homogeneity. An example is maximin with ties broken in favor of the most-preferred candidate of the voter $i \in N \subseteq N^* = \mathbb{N}$ with the smallest label. This satisfies homogeneity as long as the copy $2P$ has the same smallest-labelled voters as in P .

⁸Comparing to the graphs shown in [Table 1](#), we see that Case 1 corresponds to Graphs K and L; Case 2 corresponds to Graphs E and F; Case 3 (left) corresponds to Graphs G and H; Case 3 (right) corresponds to Graphs I and J; Case 4 corresponds to Graphs C and D.




Figure 3: A proof sketch of **Theorem 2**. In each panel, arrows from x to y indicate that x has a non-negative majority margin against y : bold arrows represent the largest margins, normal arrows smaller but strictly positive margins, and dashed arrows the smallest margins, which may be zero. Between panels, the arrow from A to B explains why the choice of a candidate in panel A would contradict the choices established in panel B .

Inspecting the rules shown in the Hasse diagram of **Figure 2**, we see that **Theorem 2** immediately implies that many rules fail optimist participation. However, there are some important rules that do satisfy optimist participation (and also the strengthened version described in **Footnote 4**): these are maximin itself (and the many other rules equivalent to it for three candidates), stable voting, Nanson's rule, and leximin. In addition, if we use a fixed tie-breaking order to make any of these rules resolute, they satisfy (optimist/resolute) participation. On the other hand, the strict Nanson rule fails optimist participation (and even positive involvement, see **Footnote 6**) despite being an anonymous and neutral maximin refinement.

Theorem 2 shows that every social choice function satisfying certain axioms must refine maximin. A theorem with the same conclusion was recently obtained by Holliday and Pacuit [33]. They introduce the axioms of *weak positive responsiveness*, which requires that if $z \in f(P)$ and if P' is obtained from P through some voter changing her preferences from $x \succ y \succ z$ to $z \succ x \succ y$, then $f(P') = \{z\}$, and the axiom of *block preservation* which requires that if Q is a profile in which each of the $3! = 6$ possible preferences occurs exactly once, then $f(P) \subseteq f(P + Q)$.

Theorem 3 (Holliday and Pacuit [33], Theorem 3.8 and Proposition 3.10). *Let f be an anonymous, neutral, and homogeneous Condorcet extension that satisfies positive involvement, weak positive responsiveness, and block preservation. Then f is a refinement of maximin.*

Note that our **Theorem 2** does not require the axioms of anonymity, neutrality, weak positive responsiveness, and block preservation but instead uses singleton negative involvement (see **Proposition 1**).

In the following sections, we will characterize several rules as the unique maximin refinements satisfying

certain axioms. Each of these results can be combined with either the axioms of [Theorem 2](#) (or those of [Theorem 3](#)) to obtain axiomatic characterizations within the class of all social choice functions.

4.2 Characterization of Leximin

Leximin refines maximin and is remarkably resolute because it only returns a tie when two candidates are “locally indistinguishable” within the margin graph (see [Footnote 2](#)). This allows it to satisfy a demanding monotonicity-style axiom called positive responsiveness. We say that a profile P' is obtained from a profile P by *improving x relative to y* if the ranking of all candidates stays fixed except that one or more voters who rank y immediately above x in P instead rank x immediately above y in P' . A social choice function f satisfies *positive responsiveness* (May [37]; Barberà [2], Definition 7) if whenever $x \in f(P)$ and P' is obtained from P by improving $x \in f(P)$ relative to $y \in A$, we have $f(P') = \{x\}$. We show that leximin is the unique maximin refinement that satisfies this property among maximin refinements that are neutral and that are pairwise. A social choice function f is *pairwise* if it only depends on the margin graph, i.e., for all P and P' such that $m_{x,y}(P) = m_{x,y}(P')$ for all $x, y \in A$, we have $f(P) = f(P')$.

Lemma 2. *Leximin is the only neutral and pairwise refinement of maximin that satisfies positive responsiveness.*

As we explained, combined with [Theorem 2](#), we obtain an axiomatic characterization.

Corollary 2. *Leximin is the only homogeneous, neutral, and pairwise Condorcet extension that satisfies optimist participation and positive responsiveness.*

4.3 Characterization of Nanson’s rule

Leximin is very decisive and only reports ties in very specific situations. However, one may argue that leximin neglects to select some candidates that deserve to win. In particular, it may fail to select a candidate even if it is “almost” a Condorcet winner in the sense that, in pairwise majority comparisons, it is never beaten but does beat some other candidates. For example, in Graph F of [Table 1](#), candidate c meets this description but is not selected by leximin, which selects only $\{a\}$. In contrast, Nanson’s rule selects $\{a, c\}$. Note that Graph F is the only case where leximin and Nanson’s rules differ. Graphs D and E also have intermediate Condorcet winners ($\{a\}$ and $\{a, c\}$, respectively).

The concept of candidates that are almost Condorcet winners in this sense was first considered by Buckley and Westen [12] and Gehrlein [25]. They were later termed “intermediate Condorcet winners” by Barberà and Bossert [3]. Formally, a candidate x is an *intermediate Condorcet winner* if $m_{x,y} \geq 0$ for all $y \in A$, and at least one of these inequalities is strict. Barberà and Bossert [3] say that f is a *strong Condorcet extension* if $f(P)$ is equal to the set of intermediate Condorcet winners whenever this set is non-empty. Note that whenever a candidate is a Condorcet winner, it is also the unique intermediate Condorcet winner. In Graph F in [Table 1](#), a and c are intermediate Condorcet winners, and thus a strong Condorcet extension must select both of them. Since leximin only selects a it is not a strong Condorcet extension (the same is true for stable voting). However, Nanson’s rule is a strong Condorcet extension.⁹

Given our previous result about leximin, Nanson’s rule must violate positive responsiveness. However, it does satisfy a slight weakening that is concerned with voters who rank two winning candidates next to each other; if such a voter flips those candidates, then the tie should be broken in favor of the strengthened candidate. Thus, a social choice function f satisfies *tie-break positive responsiveness* if whenever $x, y \in f(P)$ and P' is obtained from P by improving x relative to y , then $f(P') = \{x\}$. This is a weaker property than positive responsiveness.

⁹For more than 3 candidates, Nanson’s rule is not a strong Condorcet extension. In the profile $2abcd + 4adbc + 1bcda + 1cdab + 4cdba$, a is the unique intermediate Condorcet winner but Nanson selects $\{c\}$.

Lemma 3. *Nanson’s rule is the only neutral and pairwise refinement of maximin that satisfies tie-break positive responsiveness and is a strong Condorcet extension.*

As before, combined with [Theorem 2](#), we obtain an axiomatic characterization.

Corollary 3. *Nanson’s rule is the only homogeneous, neutral, and pairwise strong Condorcet extension that satisfies optimist participation and tie-break positive responsiveness.*

4.4 Characterization of Maximin

Lastly, we characterize maximin itself. We do so using the continuity axiom, which has been used in other characterizations to rule out tie-breaking mechanisms [e.g., 52].¹⁰ One can view leximin and Nanson as applying a certain tie-breaking mechanism to maximin, and indeed these two rules fail continuity. In fact, *all* of maximin’s strict refinements fail continuity. This is true for any number of candidates and may thus be of independent interest.

Lemma 4. *Maximin is continuous. If f is continuous and a refinement of maximin, then f is equal to maximin.*

Again, combined with [Theorem 2](#), we obtain an axiomatic characterization.

Corollary 4. *Maximin is the only homogeneous and continuous Condorcet extension that satisfies optimist participation.*

The axioms in [Corollary 4](#) are independent: leximin satisfies all axioms except continuity, Borda’s rule satisfies all axioms except Condorcet-consistency, and the top cycle satisfies all axioms except optimist participation. Without homogeneity, we can construct an artificial rule that is identical to maximin, except for the profile $3abc + 3bca + 2cab + 1acb$, where the rule returns $\{a, c\}$ instead of $\{a\}$. Clearly, homogeneity is violated (doubling this profile leads to outcome $\{a\}$). All other properties are inherited from maximin. Checking that optimist participation is satisfied can be done by case analysis.

5 Conclusion

We have investigated whether the search for a desirable Condorcet extension becomes easier when focussing on the special case of three candidates. Our results highlight the maximin rule and two of its refinements (Nanson’s rule and leximin) as being particularly robust to common criticisms of Condorcet extensions. Indeed, we showed that they are axiomatically characterized by their immunity to the no-show paradox, together with other desirable properties such as positive responsiveness. These conclusions could motivate advocating for their adoption in real-world elections.

Nanson [44] gave a rather simple description of his rule for three candidates: each voter assigns 2 points to his most preferred candidate and 1 point to his second most preferred candidate; all candidates whose score exceeds the total number of voters face off in a runoff election. Leximin is a simple refinement of Nanson where a tie in the runoff is broken using the scores from the first round. In the generic case (which applies when the number of voters is large), maximin, Nanson, and leximin all coincide.

Studies on the frequency of voting paradoxes complement our results by showing that maximin (and its refinements) not only do well for small but also for large numbers of voters when there are three

¹⁰To see why continuity rules out tie-breaking, suppose candidates a and b are tied, with the tie broken in favor of a . When copying the profile n times, there will still be a tie, broken in favor of a (assuming homogeneity). However, if we add a profile P' where b is a clear winner, then in $nP + P'$, candidate b will be the strongest. This is a failure of continuity: $f(nP + P') = \{b\} \not\subseteq \{a\} = f(P)$.




Figure 4: Fraction of anonymous profiles in which SCFs return more than one winner, computed using Ehrhart theory [see, e.g., 7] and the Normaliz package [11].

candidates. Courtin et al. [15] analyze the frequency of the reinforcement paradox of various Condorcet extensions using Monte Carlo simulations and find that “although all frequencies are small, they are smaller for [maximin].” Heilmaier [28] proves that when the number of voters goes to infinity, maximin only suffers from the reinforcement paradox for 0.37% of all pairs of anonymous profiles in which the winners coincide, which is lower than the corresponding numbers for Black’s rule and plurality with runoff. Plassmann and Tideman [48] analyze the frequency of voting paradoxes based on data generated using a spatial model, which they argue most accurately describes real-world preference profiles for three candidates. They conclude that “the Black rule and the Nanson rule encounter most paradoxes and ties less frequently than the other rules do, especially in elections with few voters.” As we have discussed, the leximin rule produces ties even less frequently than Nanson’s rule. We can quantify this effect by computing the fraction of anonymous profiles on which different rules are non-resolute, as a function of the number of voters. The results are shown in Figure 4 and show that leximin even outperforms Black’s rule.¹¹

In conclusion, we believe that maximin and its refinements are very attractive for three-candidate elections and are compelling choices for adoption in real-world applications.¹² Similar conclusions were drawn by Nurmi [45], Felsenthal and Nurmi [21], and Lepelley and Smaoui [36]. We emphasize that our arguments do not extend beyond three-candidate elections. When there are four or more candidates, the no-show paradox cannot be avoided. In addition, there are many different ways of extending the three-candidate maximin rule to handle additional candidates, and not all of them are equally desirable. For example, for four or more candidates, maximin may return candidates that are last-ranked by a majority of voters [see, e.g., 20, 9], while other generalizations like split cycle avoid this problem.

¹¹This can be explained by observing that for every social choice function f that is pairwise, neutral, and never selects Pareto-dominated candidates, it holds that if leximin is irresolute on a profile P , then f is also irresolute on P (see Graphs A, B, and F in Table 1). So leximin is resolute whenever possible.

¹²Early attempts to implement Nanson’s rule in real-world elections were met with limited success [see 39, for a historical account].

Acknowledgments

This material is based on work supported by the Deutsche Forschungsgemeinschaft under grants BR 2312/11-2 and BR 2312/12-1 and by the Agence Nationale de la Recherche under grant ANR-22-CE26-0019 (CITIZENS). The authors are grateful to former students Florian Grundbacher, Keyvan Kardel, and Christian Stricker for constructing some of the examples used in this paper.

References

- [1] Kenneth J. Arrow. *Social Choice and Individual Values*. New Haven: Cowles Foundation, 1st edition, 1951. 2nd edition 1963.
- [2] Salvador Barberà. The manipulation of social choice mechanisms that do not leave “too much” to chance. *Econometrica*, 45(7):1573–1588, 1977.
- [3] Salvador Barberà and Walter Bossert. Intermediate Condorcet winners and losers. Technical Report 4320762, SSRN, 2024.
- [4] Duncan Black. On the rationale of group decision-making. *Journal of Political Economy*, 56(1): 23–34, 1948.
- [5] Florian Brandl, Felix Brandt, Dominik Peters, and Christian Stricker. Distribution rules under dichotomous preferences: Two out of three ain’t bad. In *Proceedings of the 22nd ACM Conference on Economics and Computation (ACM-EC)*, pages 158–179, 2021.
- [6] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
- [7] Felix Brandt, Christian Geist, and Martin Strobel. Analyzing the practical relevance of voting paradoxes via Ehrhart theory, computer simulations, and empirical data. In *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 385–393, 2016.
- [8] Felix Brandt, Christian Geist, and Dominik Peters. Optimal bounds for the no-show paradox via SAT solving. *Mathematical Social Sciences*, 90:18–27, 2017. Special Issue in Honor of Hervé Moulin.
- [9] Felix Brandt, Marie Matthäus, and Christian Saile. Minimal voting paradoxes. *Journal of Theoretical Politics*, 34(4):527–551, 2022.
- [10] Felix Brandt, Christian Saile, and Christian Stricker. Strategyproof social choice when preferences and outcomes may contain ties. *Journal of Economic Theory*, 202:105447, 2022.
- [11] Winfried Bruns, Bogdan Ichim, Christof Söger, and Ulrich von der Ohe. Normaliz. Algorithms for rational cones and affine monoids. Available at <https://www.normaliz.uni-osnabrueck.de>, 2024.
- [12] James J. Buckley and Thomas E. Westen. The probability of the voter’s paradox for an even number of voters. *Journal of Interdisciplinary Modeling and Simulation*, 2:185–210, 1979.
- [13] Ioannis Caragiannis, Edith Hemaspaandra, and Lane A. Hemaspaandra. Dodgson’s Rule and Young’s Rule. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 5. Cambridge University Press, 2016.
- [14] Marquis de Condorcet. *Essai sur l’application de l’analyse à la probabilité des décisions rendues à la pluralité des voix*. Imprimerie Royale, 1785. Facsimile published in 1972 by Chelsea Publishing Company, New York.
- [15] Sébastien Courtin, Boniface Mbih, and Issofa Moyouwou. Are Condorcet procedures so bad according to the reinforcement axiom? *Social Choice and Welfare*, 42(4):927–940, 2014.
- [16] Richard B. Darlington. Minimax is the best electoral system after all. Technical report, <https://arxiv.org/abs/1606.04371>, 2016.
- [17] Bernard Debord. Caractérisation des matrices des préférences nettes et méthodes d’agrégation associées. *Mathématiques et sciences humaines*, 97:5–17, 1987.
- [18] John Duggan. Uncovered sets. *Social Choice and Welfare*, 41(3):489–535, 2013.

- [19] Bora Erdamar and M. Remzi Sanver. Choosers as extension axioms. *Theory and Decision*, 67(4): 375–384, 2009.
- [20] Dan S. Felsenthal. Review of paradoxes afflicting various voting procedures where one out of m candidates ($m \geq 2$) must be elected. In *Electoral Systems*, pages 19–91. Springer-Verlag, 2012.
- [21] Dan S. Felsenthal and Hannu Nurmi. Monotonicity violations by Borda’s elimination and Nanson’s rules: A comparison. *Group Decision and Negotiation*, 27(4):637–664, 2018.
- [22] Peter C. Fishburn. Condorcet social choice functions. *SIAM Journal on Applied Mathematics*, 33(3): 469–489, 1977.
- [23] Peter C. Fishburn. Discrete mathematics in voting and group coherence. *SIAM Journal of Algebraic and Discrete Methods*, 5(2):263–275, 1984.
- [24] Peter Gärdenfors. On definitions of manipulation of social choice functions. In Jean-Jacques Laffont, editor, *Aggregation and Revelation of Preferences*. North-Holland, 1979.
- [25] William V. Gehrlein. Condorcet’s paradox. *Theory and Decision*, 15(2):161–197, 1983.
- [26] Christian Geist and Dominik Peters. Computer-aided methods for social choice theory. In Ulle Endriss, editor, *Trends in Computational Social Choice*, chapter 13, pages 249–267. AI Access, 2017.
- [27] Allan Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41(4):587–601, 1973.
- [28] Sophie Heilmaier. Optimal voting rules for few candidates. Master’s thesis, Technical University of Munich, 2020.
- [29] Wesley H. Holliday. An impossibility theorem concerning positive involvement in voting. *Economics Letters*, 236:111589, 2024.
- [30] Wesley H. Holliday and Eric Pacuit. Measuring violations of positive involvement in voting. In *Proceedings of the 18th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, EPTCS 335, pages 189–209, 2021.
- [31] Wesley H. Holliday and Eric Pacuit. Split cycle: A new Condorcet consistent voting method independent of clones and immune to spoilers. *Public Choice*, 197(1–2):1–62, 2023.
- [32] Wesley H. Holliday and Eric Pacuit. Stable voting. *Constitutional Political Economy*, 34:421–433, 2023.
- [33] Wesley H. Holliday and Eric Pacuit. An extension of May’s theorem to three alternatives: axiomatizing minimax voting. Technical report, <https://arxiv.org/abs/2312.14256>, 2023.
- [34] José L. Jimeno, Joaquín Pérez, and Estefanía García. An extension of the Moulin no show paradox for voting correspondences. *Social Choice and Welfare*, 33(3):343–459, 2009.
- [35] Jean-François Laslier. *Tournament Solutions and Majority Voting*. Springer-Verlag, 1997.
- [36] Dominique Lepelley and Hatem Smaoui. Comparing two ways for eliminating candidates in three-alternative elections using sequential scoring rules. *Group Decision and Negotiation*, 28: 787–804, 2019.
- [37] Kenneth May. A set of independent, necessary and sufficient conditions for simple majority decisions. *Econometrica*, 20(4):680–684, 1952.
- [38] David C. McGarvey. A theorem on the construction of voting paradoxes. *Econometrica*, 21(4): 608–610, 1953.
- [39] Iain McLean. E. J. Nanson, social choice and electoral reform. *Australian Journal of Political Science*, 31(3):369–386, 1996.
- [40] Iain McLean and Fiona Hewitt. *Condorcet: Foundations of Social Choice and Political Theory*. Edward Elgar Publishing, 1994.
- [41] Hervé Moulin. *Axioms of Cooperative Decision Making*. Cambridge University Press, 1988.
- [42] Hervé Moulin. Condorcet’s principle implies the no show paradox. *Journal of Economic Theory*, 45(1):53–64, 1988.
- [43] Roger B. Myerson. Axiomatic derivation of scoring rules without the ordering assumption. *Social Choice and Welfare*, 12(1):59–74, 1995.
- [44] Edward J. Nanson. Methods of election. *Transactions and Proceedings of the Royal Society of Victoria*, 19:197–240, 1883.

- [45] Hannu Nurmi. On Nanson’s method. In *Democracy in the Modern World - Essays for Tatu Vanhanen*, volume 260 of *Acta Universitatis Tamperensis*, pages 199–210. University of Tampere, 1989.
- [46] Joaquín Pérez. Incidence of no show paradoxes in Condorcet choice functions. *Investigaciones Economicas*, 19:139–154, 1995.
- [47] Joaquín Pérez. The strong no show paradoxes are a common flaw in Condorcet voting correspondences. *Social Choice and Welfare*, 18(3):601–616, 2001.
- [48] Florenz Plassmann and Nicolaus Tideman. How frequently do different voting rules encounter voting paradoxes in three-candidate elections? *Social Choice and Welfare*, 42(1):31–75, 2014.
- [49] Mark A. Satterthwaite. Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2): 187–217, 1975.
- [50] James S. Weber. How many voters are needed for paradoxes? *Economic Theory*, 20:341–355, 2002.
- [51] H. Peyton Young. An axiomatization of Borda’s rule. *Journal of Economic Theory*, 9(1):43–52, 1974.
- [52] H. Peyton Young. Social choice scoring functions. *SIAM Journal on Applied Mathematics*, 28(4): 824–838, 1975.
- [53] H. Peyton Young and Arthur B. Levenglick. A consistent extension of Condorcet’s election principle. *SIAM Journal on Applied Mathematics*, 35(2):285–300, 1978.
- [54] William S. Zwicker. Introduction to the theory of voting. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 2. Cambridge University Press, 2016.

Felix Brandt
 Technische Universität München
 Munich, Germany
 Email: brandtf@in.tum.de

Chris Dong
 Technische Universität München
 Munich, Germany
 Email: chris.dong@tum.de

Dominik Peters
 CNRS, Université Paris Dauphine
 Paris, France
 Email: mail@dominik-peters.de

Appendix

We now present our results regarding reinforcement and then provide all missing proofs.

A Reinforcement

In this appendix, we study the reinforcement axiom, which demands that candidates who win for two disjoint electorates should be precisely the winners for the union of these electorates [51, 52]. Suppose that $P = (\succ_i)_{i \in N}$ and $P' = (\succ_i)_{i \in N'}$ are preference profiles defined over electorates N and N' that are disjoint: $N \cap N' = \emptyset$. We can, therefore, merge these profiles into a single profile $P + P' = (\succ_i)_{i \in N \cup N'}$ defined on the electorate $N \cup N'$.

Definition 3 (Reinforcement). A social choice function f satisfies *reinforcement* if for all profiles P and P' on disjoint electorates, we have $f(P + P') = f(P) \cap f(P')$ whenever $f(P) \cap f(P') \neq \emptyset$.

In other words, if there is a candidate a who wins in both P and P' , then a also wins in the combined profile $P + P'$, and all other winners in $P + P'$ must also win in both P and P' .

Young and Levenglick [53, Theorem 2] famously showed that reinforcement is in conflict with Condorcet-consistency. Their result is only phrased in terms of *weak* Condorcet winners, but they briefly mention a variant of this result that can be turned into a 13-voter impossibility for Condorcet extensions [see, e.g., 41, Theorem 9.2]. We will show that the incompatibility remains intact even when there are fewer than 13 voters.

A.1 Simple Proof for Nine Voters


We begin by giving a very simple proof of the incompatibility for 9 voters. It can be obtained by essentially optimizing the existing proofs of Young and Levenglick [53] and Moulin [41] with respect to the number of required voters. This proof was found by Keyvan Kardel and also appears in the handbook chapter by Zwicker [54, Prop. 2.5].

A remarkable feature of this proof is that it only uses a weak form of reinforcement: A social choice function satisfies *subset-reinforcement* if for all P_1 and P_2 defined on disjoint electorates, we have $f(P_1) \cap f(P_2) \subseteq f(P_1 + P_2)$. This property was also discussed by Young and Levenglick [53].

Theorem 4. *Every Condorcet extension violates subset-reinforcement when $|N^*| \geq 9$.*

Proof. Let f be a Condorcet extension satisfying subset-reinforcement. Let N_1 and N_2 be two disjoint electorates consisting of 6 and 3 voters, respectively. Without loss of generality, let $N_1 = \{1, 2, 3, 4, 5, 6\}$ and $N_2 = \{7, 8, 9\}$. Consider the following profile P_1 with 6 voters, which is a “double Condorcet cycle”:

$\succ_1$	$\succ_2$	$\succ_3$	$\succ_4$	$\succ_5$	$\succ_6$
a	a	b	b	c	c
b	b	c	c	a	a
c	c	a	a	b	b



Because $f(P_1) \neq \emptyset$, we may assume without loss of generality that $a \in f(P_1)$ (otherwise relabel the candidates). Now consider the following profile P_2 with 3 voters:

$\succ_7$	$\succ_8$	$\succ_9$
a	a	c
c	c	a
b	b	b

In this profile, a is the Condorcet winner, so $f(P_2) = \{a\}$ because f is a Condorcet extension. Note that P_1 and P_2 are defined over disjoint electorates, and that $a \in f(P_1) \cap f(P_2)$. Because f satisfies subset-reinforcement, we have $f(P_1) \cap f(P_2) \subseteq f(P_1 + P_2)$, and therefore $a \in f(P_1 + P_2)$. However, in the combined profile $P_1 + P_2$, c is the Condorcet winner:

$\succ_1$	$\succ_2$	$\succ_3$	$\succ_4$	$\succ_5$	$\succ_6$	$\succ_7$	$\succ_8$	$\succ_9$
a	a	b	b	c	c	a	a	c
b	b	c	c	a	a	c	c	a
c	c	a	a	b	b	b	b	b

Thus, since f is a Condorcet extension, we have $f(P_1 + P_2) = \{c\}$, a contradiction. $\square$

Theorem 4 also holds for “superset-reinforcement”, which requires that for all profiles P and P' , we have $f(P) \cap f(P') \supseteq f(P + P')$ whenever $f(P) \cap f(P') \neq \emptyset$. To see this, we can slightly adapt the proof by noting that superset-reinforcement implies that $f(P_1 + P_2) \subseteq \{a\}$.

Note that the proof of **Theorem 4** only invokes (subset-)reinforcement for pairs of profiles where at least one of them has a Condorcet winner. This is a particularly compelling application of reinforcement since it has the flavor of participation: if x is a winner in a profile and we add some new voters to it who agree that x is the strongest candidate (in the sense that x is the Condorcet winner with respect to the new voters), then x should still be a winner.

A.2 Proof using Anonymity

Are 9 voters necessary to obtain the incompatibility between reinforcement and being a Condorcet extension? As it turns out, the result holds even for 8 voters.

Fact 1. Every Condorcet extension violates reinforcement when $|N^*| \geq 8$.

We discovered this fact using a SAT solver,¹³ and we have no human-readable proof for it. The minimal unsatisfiable subsets of clauses (a measure of proof complexity) that we were able to extract argue about hundreds or even thousands of profiles.

However, we can give a proof if we assume anonymity in addition. Anonymity is particularly natural in a variable-electorate model and allows us to apply reinforcement without having to ensure that electorates are disjoint.

Theorem 5. Every anonymous Condorcet extension violates reinforcement when $|N^*| \geq 8$.

¹³SAT solving is a computational method that has been leveraged to obtain impossibility theorems in social choice [see, e.g., 26]. We have also used it to find the proofs of **Theorem 5** and **Theorem 6**. To check **Fact 1**, our Boolean formula did not list all non-anonymous profiles for up to 8 voters since there are $7^8 > 5$ million of them. Instead, similarly to the approach of Brandl et al. [5, Section 5.2], it was sufficient to list all profiles that can be obtained via voter and candidate permutations from the profiles used in the proof of **Theorem 5**. To encode reinforcement, add a variable v_{P_1, P_2} for each pair P_1, P_2 of profiles indicating whether $f(P_1) \cap f(P_2) \neq \emptyset$; then for all $x \in A$ add clauses $(x \in f(P_1) \wedge x \in f(P_2)) \rightarrow v_{P_1, P_2}$ and $(v_{P_1, P_2} \wedge x \in f(P_1 \cup P_2)) \rightarrow x \in f(P_i)$ for $i = 1, 2$, and $(v_{P_1, P_2} \wedge x \in f(P_1) \wedge x \in f(P_2)) \rightarrow x \in f(P_1 \cup P_2)$.

A.3 Proof using Anonymity and Neutrality

When additionally requiring neutrality, the impossibility can be proven using only 5 voters.

Theorem 6. *Every anonymous and neutral Condorcet extension violates reinforcement when $|N^*| \geq 5$.*

The reinforcement axiom was originally introduced by Young [51, 52] to give an axiomatic characterization of the class of *scoring rules*. A scoring rule is a social choice function defined by a scoring vector $(s_1, s_2, s_3) \in \mathbb{R}^3$. Given a profile P , it then computes the score of each candidate by letting each voter assign s_1 points to their top-ranked candidate, s_2 points to the second-ranked candidate, and s_3 points to their last-ranked candidate. Then the scoring rule selects the candidates with the highest score. A scoring rule is *strictly monotonic* if $s_1 > s_2 > s_3$ and it is *weakly monotonic* if $s_1 \geq s_2 \geq s_3$ and $s_1 > s_3$. Borda’s rule is a scoring rule with $s_1 = 2$, $s_2 = 0$, and $s_3 = -2$, where we have chosen the scores so they agree with our earlier definition of the Borda score of a candidate as $\beta_x = \sum_{y \in A} m_{x,y}$. Another classic scoring rule is the *plurality rule* with $s_1 = 1$ and $s_2 = s_3 = 0$, choosing as winners those candidates who are top-ranked most frequently. Young [52] showed that if the set of potential voters N^* is infinite, then an anonymous, neutral, and continuous social choice function satisfies reinforcement if and only if it is a scoring rule. Dropping continuity from the axioms, Young [52] characterized the class of *composite* scoring rules where additional scoring rules are used to break ties (for example, choosing the plurality winners with the highest Borda score).

Now, because every (composite) scoring rule satisfies anonymity, neutrality, and reinforcement, **Theorem 6** implies that every (composite) scoring rule fails Condorcet-consistency on some profile with at most 5 voters.

This is a new proof of a classic result, first established by Condorcet [14, pages clxxvij–clxxviii, translated by 40, pages 137–138] who constructed an 81-voter profile ($30abc + acb + 10cab + 29bac + 10bca + cba$) on which a is the Condorcet winner but every strictly monotonic scoring rule uniquely elects b . Fishburn [23] gives an example with 7 voters ($3abc + 2bca + bac + cab$) that has the same property. For weakly monotonic scoring rules, Moulin [41, Theorem 9.1] gave a 17-voter example ($6abc + 3cab + 4bac + 4bca$) such that a is the Condorcet winner and every weakly monotonic scoring rule uniquely elects b . Moulin suggested that this example is minimal, but in fact an 11-voter example with the same property is minimal ($4abc + 3bca + 2bac + 2cab$, computed by Christian Stricker). Because in these examples, b is the unique scoring winner in each case, we can deduce that no refinement of a (weakly) monotonic scoring rule can be a Condorcet extension. Weber [50, Theorem 5] gives a related example.

The examples discussed in the previous paragraph are all “universal counterexamples”, while the proof of **Theorem 6** is based on reasoning across several profiles. However, by an observation of Florian Grundbacher, the profile $P_2 + P_3$ in that proof ($2abc + bac + bca + cab$) can also function as a universal counterexample, showing that no scoring rule (including non-monotonic or composite ones) can always *uniquely* select the Condorcet winner. This is because in this profile, a is the Condorcet winner, but a and b are “rank-indistinguishable”: they both appear twice in first rank, twice in second rank, and once in third rank. Hence, any scoring rule assigns a and b exactly the same score, so that a can never have the uniquely highest score.

Note that these universal counterexamples do not quite imply **Theorem 6**, because such an implication would depend on the characterization of Young [52] which requires an infinite set of potential voters, while **Theorem 6** only requires 5 voters.

A.4 Lower Bounds

In this section, we show that the numbers of voters in our incompatibility theorems for reinforcement are minimal. To start, **Theorem 6** contains a bound of 5 voters. This is the best possible since Black’s

rule, stable voting, and leximin are anonymous, neutral, and satisfy reinforcement for up to 4 voters. Interestingly, all other Condorcet extensions discussed in [Section 3](#) fail reinforcement even for 4 voters.¹⁴

For our other bounds, we will define an artificial rule that satisfies reinforcement for up to 7 voters and subset-reinforcement for up to 8 voters, which shows that the bounds of [Theorem 5](#) and [Theorem 4](#) are the best possible. Given [Theorem 6](#), this artificial rule necessarily fails neutrality. It is a refinement of maximin that breaks ties according to an intricate scoring system which we found by solving an integer linear program. This scoring system is non-neutral and specifies for each linear order a number of points that it assigns to its top and its second choice. In addition, for the linear order abc , the number of points differs depending on whether the number of voters in the input profile is in $\{2, 4, 6, 8\}$ (Case 1) or not (Case 2). Our rule then returns the maximin winner that obtained the highest score.

Case 1		Case 2					
a	11	a	18	a	10	b	18
b	8	b	13	c	7	a	11
c	0	c	0	b	0	c	0
						b	7
						c	0
						a	0
						c	13
						a	5
						b	0
						a	0

Using a computer, one can verify that this Condorcet extension satisfies reinforcement up to 7 voters and subset-reinforcement up to 8 voters.¹⁵ It is also anonymous (but not neutral) and satisfies participation as well as monotonicity, which states that improving a winner x in some voters' preference relation while keeping everything else intact should not remove x from the choice set. It fails homogeneity, as one can find a profile with 8 voters in which a and c are the weak Condorcet winners and c is selected due to the tie-breaking, but doubling the profile makes a the unique winner because different scores are used for 16 voters.

B Missing proofs

We now provide the omitted proofs.

Theorem 1. *For three candidates, [Figure 2](#) shows the Hasse diagram of the refinement relation between select Condorcet extensions.*

Proof. Except for Black's and Baldwin's rules (which we will consider separately), all of the social choice functions included in the diagram depend only on the ordinal margin graph. This means that if two profiles induce the same ordinal margin graph, then the output is the same for the two profiles [29]. For most rules, this can be seen straightforwardly from their definitions. In fact, many of the Condorcet extensions in the diagram only depend on the orientation of the edges in the margin graph and not their relative weights. For Kemeny's, Dodgson's, and Young's rules, this follows because they are equivalent to maximin for three candidates (15, working paper version, Theorem 1; see also 28, Theorem 3.2). For the Nanson and strict Nanson rules, this follows because, for three candidates, the ordinal margin graph contains enough information to determine whether a candidate has below-average Borda score.¹⁶ Note that all these rules are Condorcet extensions, and they are neutral. Therefore, these rules can be fully

¹⁴Consider the profile $acb + (abc + bca + cab)$, on which every anonymous and neutral Condorcet extension satisfying reinforcement must select $\{a\}$ (this can be seen from the indicated decomposition into two subprofiles). This profile corresponds to Graph D in [Table 1](#), which indicates that most rules select $\{a, c\}$ on this profile; this is also true for Baldwin. Copeland and Nanson fail reinforcement for the profile $(abc + bac) + (abc + bca)$, where they select $\{a, b\}$ on the first subprofile, $\{b\}$ on the second subprofile, but $\{a, b\}$ on the whole profile.

¹⁵Code is available at <https://gist.github.com/DominikPeters/eeee2456ee4a4cdf2c5aa302b24af7a6>.

¹⁶Nanson's rule eliminates all candidates whose Borda score is at most 0 (or strictly below 0 for the strict Nanson rule). For three candidates, a candidate's Borda score is at most 0 if and only if it either does not have outgoing edges or it has one outgoing edge that is lighter than its incoming edge. It is strictly below 0 if it either has no outgoing edges and an incoming edge or it has one outgoing edge that is strictly lighter than its incoming edge.

specified by listing their outputs on the 12 possible non-isomorphic ordinal margin graphs in which there is no Condorcet winner. We have produced such a list in Table 1, which immediately establishes most of the relations displayed in Figure 2.

For Black's rule, note that it will never select a candidate that only has incoming edges but no outgoing edges (since such a candidate has negative Borda score). Inspecting Table 1, we see that, when there is no Condorcet winner, Banks' rule returns all candidates for which this is not the case (among others); hence Black's rule refines Banks (which in turn refines UC McKelvey and the top cycle). On the profile $3abc + bca + 4cab$, Black's rule selects $\{a\}$ while UC Gillies selects $\{b, c\}$; hence Black's rule is not refined by and does not refine any other rule in the diagram (which are all refinements of UC Gillies).

For Baldwin's rule, we can go through the 12 ordinal margin graphs in Table 1 to determine what it may output on each of them and thereby deduce that it refines the defensible set (and consequentially also UC Gillies, UC McKelvey, and the top cycle).¹⁷ On the profile $4acb + 5bac + 3cab + 5cba$, Baldwin selects $\{a\}$ while Black's rule, leximin, and strict Nanson (the leaves of the diagram) all select $\{c\}$, so Baldwin is not refined by any other rule in the diagram. On the other hand, on the profile $1abc + 3bca + 4cab$, Baldwin's rule selects $\{b, c\}$ while Banks' rule selects $\{a, c\}$, so Baldwin does not refine rules other than the ones indicated in the diagram. $\square$

Lemma 2. *Leximin is the only neutral and pairwise refinement of maximin that satisfies positive responsiveness.*

Proof. We first show that leximin satisfies all axioms. It is neutral and pairwise by definition. To see that it satisfies positive responsiveness, note that if P' is obtained from P by moving up $x \in f_{\text{leximin}}(P)$ then the margins of x only improve, i.e., $m^x(P') \triangleright m^x(P)$, while for any $z \neq x$ the margins can only become worse, i.e., $m^z(P) \trianglerighteq m^z(P')$. (See Footnote 2 for the definition of the relation $\triangleright$.) Let $z \in A \setminus \{x\}$. Then because x is selected by leximin, we have $m^x(P) \trianglerighteq m^z(P)$. Thus, $m^x(P') \triangleright m^x(P) \trianglerighteq m^z(P) \trianglerighteq m^z(P')$. Hence $z \notin f_{\text{leximin}}(P')$ and so $f_{\text{leximin}}(P') = \{x\}$ as desired.

For uniqueness, let f be a maximin refinement satisfying neutrality, pairwiseness, and positive responsiveness. By pairwiseness, note that f is also anonymous. We need to show that $f = f_{\text{leximin}}$. By pairwiseness, it suffices to consider each possible margin graph, find some profile P that induces it, and show that $f(P) = f_{\text{leximin}}(P)$. So consider any margin graph. If maximin selects a unique winner when given a profile inducing that margin graph, we are done since then all maximin refinements select the same choice set. All other margin graphs correspond to Graphs A to F in Table 1.

For Graph A (the Condorcet cycle with three edges each with positive weight t) take $P = t \cdot (abc + bca + cab)$ and for Graph B (all pairs of candidates are majority tied) take the profile P consisting of one copy of each of the six preference orders. Anonymity and neutrality imply $f(P) = \{a, b, c\}$ which equals $f_{\text{leximin}}(P)$.

For Graphs C and D, we have a majority graph with $0 \leq m_{c,a} = m_{b,c} < m_{a,b}$. Let P be a profile inducing this graph (which exists by McGarvey's theorem). Let $t = \frac{1}{2}(m_{a,b} - m_{b,c}) > 0$. We may assume that there are at least t many voters with preferences $c \succ a \succ b$ in P (otherwise we can add those voters to P and also add t voters with opposite preferences $b \succ a \succ c$ which leads to a profile inducing the same margin graph). Let P' be the profile obtained from P by selecting t voters with preferences $c \succ a \succ b$ and replace their preference by $c \succ b \succ a$. Then on P' , all majority margins are equal, i.e., $0 \leq m_{c,a}(P') = m_{a,b}(P') = m_{b,c}(P')$. This corresponds to Graphs A or B, and thus $f(P') = \{a, b, c\}$. Note that P can be viewed as being obtained from P' by improving $a \in f(P')$ relative to b . Thus, by positive responsiveness, we obtain $f(P) = \{a\} = f_{\text{leximin}}(P)$.

¹⁷For Graphs A and B, the defensible set contains all candidates, so Baldwin refines it. For Graphs C, D, E, and F, b is the Borda loser, and so is eliminated first, so Baldwin refines the defensible set $\{a, c\}$. For Graphs G and H, the Borda loser is either b or c ; if b is eliminated then c wins; if c is eliminated then a wins, refining the defensible set $\{a, c\}$. For Graphs I and J, c is eliminated, followed by b , so a wins, refining the defensible set $\{a, c\}$. For Graphs K and L, c is eliminated, followed by b , so a wins, refining the defensible set $\{a\}$.

For Graph E, where the two edges both have weight $2t$, consider the profile $P = t \cdot (acb + cab)$ inducing it. Note that $b \notin f_{\text{maximin}}(P)$ and thus $b \notin f(P)$. By anonymity and neutrality, we obtain $f(P) = \{a, c\} = f_{\text{leximin}}(P)$.

For Graph F, we have $m_{a,b} > m_{c,b} > m_{c,a} = 0$. Let $t = \frac{1}{2}(m_{a,b} - m_{c,b}) > 0$ and let P be a profile inducing Graph F with at least t voters having preference $a \succ b \succ c$. By changing t of these votes to $b \succ a \succ c$, we reach a profile P' on which the new majority margins satisfy $m_{a,b}(P') = m_{c,b}(P')$, which is of type Graph E, and thus $f(P') = \{a, c\}$. Note that P can be viewed as being obtained from P' by improving $a \in f(P')$ relative to b . Thus, by positive responsiveness, we obtain $f(P) = \{a\} = f_{\text{leximin}}(P)$, concluding the proof. $\square$

Lemma 3. *Nanson's rule is the only neutral and pairwise refinement of maximin that satisfies tie-break positive responsiveness and is a strong Condorcet extension.*

Proof. We first show that Nanson's rule satisfies these axioms. It refines maximin by [Theorem 1](#) and is pairwise and neutral as Borda scores only depend on the majority margins and are neutral. That Nanson's rule is a strong Condorcet extension can be seen from [Table 1](#). Tie-break positive responsiveness is satisfied due to the following case distinction: Let P be a profile. If $|f_{\text{Nanson}}(P)| = 1$, then tie-break positive responsiveness does not apply to P . If $|f_{\text{Nanson}}(P)| = 2$, say $f_{\text{Nanson}}(P) = \{a, c\}$, then a and c have positive Borda score in the first round and the resulting majority comparison is a tie. When we improve, say, a relative to c , this increases the Borda score of a and decreases the one of c (and keeps the Borda score of b the same). Thereafter, a may be the unique candidate with positive Borda score in the first round, or a and c both retain positive Borda score, but then a wins the majority comparison against c . Hence, in either case, $f_{\text{Nanson}}(P') = \{a\}$, as required by tie-break positive responsiveness. If $|f_{\text{Nanson}}(P)| = 3$, then the rule stopped without eliminating any candidates, which means that all candidates have a Borda score of zero in P . Hence, improving some x relative to any y makes x the unique candidate with strictly positive Borda score, and thus $f_{\text{Nanson}}(P') = \{x\}$, as required by tie-break positive responsiveness.

For uniqueness, let f be neutral and pairwise maximin refinement that is a strong Condorcet extension and that satisfies tie-break positive responsiveness. We go through all possible margin graphs and show that f coincides with Nanson's rule. Whenever maximin chooses only one candidate, it is clear that all refinements return the same choice set. Considering [Table 1](#), we see that only Graphs A to F remain.

For Graphs A and B (Condorcet cycles with three equally and non-negatively weighted edges), pairwise and neutrality imply that $f(P) = \{a, b, c\} = f_{\text{Nanson}}(P)$ as in the proof of [Lemma 2](#). The argument for Graphs C and D is identical to the corresponding argument in the proof of [Lemma 2](#) because tie-break positive responsiveness suffices. For Graphs E and F, strong Condorcet-consistency requires that $f(P) = \{a, b\} = f_{\text{Nanson}}(P)$. This concludes the proof that f is equal to Nanson's rule. $\square$

Lemma 4. *Maximin is continuous. If f is continuous and a refinement of maximin, then f is equal to maximin.*

Proof. In this proof, a candidate $x^* \neq x$ is said to be an *opponent* of x w.r.t. profile P if $m_{x,x^*}(P) \leq m_{x,z}(P)$ for all $z \neq x$. With this terminology, the definition of the maximin rule becomes $f_{\text{maximin}}(P) = \{x \in A : m_{x,x^*}(P) \geq m_{y,y^*}(P) \text{ for all } y \in A\}$.

We first claim the following:

for any profiles P and P' and any candidate x , if x^* is an opponent of x w.r.t. P and among these minimizes $m_{x,x^*}(P')$, then for all large enough n , x^* is also an opponent of x w.r.t. $nP + P'$. (1)

To prove this, we have to show that $m_{x,x^*}(nP + P') \leq m_{x,z}(nP + P')$ for all $z \neq x$. First, let z not be an opponent of x . Then, $m_{x,x^*}(P) < m_{x,z}(P)$, and hence for large enough n , we have

$$m_{x,x^*}(nP + P') = nm_{x,x^*}(P) + m_{x,x^*}(P') < nm_{x,z}(P) + m_{x,z}(P') = m_{x,z}(nP + P').$$

Otherwise, z is an opponent of x and hence $m_{x,x^*}(P) = m_{x,z}(P)$. By choice of x^* , we obtain

$$m_{x,x^*}(nP + P') = nm_{x,x^*}(P) + m_{x,x^*}(P') \leq nm_{x,z}(P) + m_{x,z}(P') = m_{x,z}(nP + P'),$$

proving the claim.

To show that maximin satisfies continuity, let P and P' be two profiles and write $W = f_{\text{maximin}}(P)$. For each $z \in A$, denote by z^* an opponent of z w.r.t. P which among them minimizes $m_{z,z^*}(P')$. If $W = A$, then trivially $f_{\text{maximin}}(nP + P') \subseteq W$ and we are done. Otherwise, let $x \in W$ and $y \notin W$, and so $m_{x,x^*}(P) > m_{y,y^*}(P)$. By choosing n large enough, we also obtain that $m_{x,x^*}(nP + P') > m_{y,y^*}(nP + P')$. By (1), y is not chosen by maximin in $nP + P'$. Hence, for n large enough, we have $f_{\text{maximin}}(nP + P') \subseteq W$.

For the second claim, let f be a refinement of maximin that is continuous. Let P be a profile and write $W = f_{\text{maximin}}(P)$; we need to show that $f(P) = W$. If $|W| = 1$, this is true since f refines maximin. Otherwise, write $W = \{x_1, \dots, x_k\}$ with $k \geq 2$. Without loss of generality, it suffices to show that $x_1 \in f(P)$.

Suppose first that $m_{x_i,z}(P) \geq 0$ for all $x_i \in W$ and all $z \in A$. Then for all $x_i, x_j \in W$, we have $m_{x_i,x_j}(P) \geq 0$ and $m_{x_j,x_i}(P) \geq 0$. Since $m_{x_i,x_j}(P) = -m_{x_j,x_i}(P)$, it follows that $m_{x_i,x_j}(P) = 0$. Now let P' be a profile consisting of a single voter who ranks x_1 on top. For every $n \geq 1$, the profile $nP + P'$ has x_1 as its Condorcet winner because it now strictly beats all other candidates. Because f is a Condorcet extension (being a maximin refinement) and because it is continuous, we have $\{x_1\} = f(nP + P') \subseteq f(P)$ and therefore $x_1 \in f(P)$, as desired.

Otherwise, $m_{x_k,y}(P) < 0$ for some $x_k \in W$ and some $y \neq x_k$. For each i , let x_i^* be an opponent of x_i w.r.t. P . Because $m_{x_k,y}(P) < 0$, we have $m_{x_k,x_k^*}(P) < 0$. Thus by definition of maximin, we have $m_{x_i,x_i^*}(P) < 0$ for all $x_i \in W$. Using McGarvey's theorem, we now construct a profile P' with the following margin graph: for each $x_i \in W \setminus \{x_1\}$, set $m_{x_i,x_i^*}(P') = -2$ and $m_{x_i^*,x_i}(P') = 2$; all other majority margins are set to zero. To see that this is well-defined (and we do not assign different weights to the same edge), note that if we had $(x_i, x_i^*) = (x_j^*, x_j)$ for some i and j , then $m_{x_i,x_i^*}(P) = m_{x_j^*,x_j}(P) = -m_{x_j,x_j^*}(P) > 0$, contradicting $m_{x_i,x_i^*}(P) < 0$. Now, for each $x_i \in W$, the margin $m_{x_i,x_i^*}(P') = -2$ is the smallest margin appearing in P' and thus x_i^* is an opponent x_i w.r.t. P' which minimizes $m_{x_i,x_i^*}(P')$. Thus for large enough n , due to (1), we find that for every $x_j \in W$, the candidate x_j^* is an opponent of x w.r.t. $nP + P'$. Note that because maximin satisfies continuity, we have $f_{\text{maximin}}(nP + P') \subseteq f_{\text{maximin}}(P) = W$ for large enough n . In addition, for all $x_j \neq x_1$, we have

$$m_{x_1,x_1^*}(nP + P') = nm_{x_1,x_1^*}(P) + 0 = nm_{x_j,x_j^*}(P) > nm_{x_j,x_j^*}(P) - 2 = m_{x_j,x_j^*}(nP + P'),$$

and hence $x_j \notin f_{\text{maximin}}(nP + P')$. Thus $f_{\text{maximin}}(nP + P') = \{x_1\}$. Because f is a maximin refinement and continuous, we obtain $\{x_1\} = f(nP + P') \subseteq f(P)$ for n large enough, and therefore $x_1 \in f(P)$, as desired. $\square$