Exploring the No-Show Paradox for Condorcet Extensions Using Ehrhart Theory and Computer Simulations

Felix Brandt[1], **Johannes Hofbauer**[2], **Martin Strobel**[3],
[1],[2] Technische Universität München
[3] National University of Singapore
brandtf@in.tum.de, hofbauej@in.tum.de, martin.r.strobel@gmail.com

Abstract

Results from voting theory are increasingly used when dealing with collective decision making in computational multiagent systems. An important and surprising phenomenon in voting theory is the *No-Show Paradox (NSP)*, which occurs if a voter is better off by abstaining from an election. While it is known that certain voting rules suffer from this paradox in principle, the extent to which it is practical concern is not well understood. We aim at filling this gap by analyzing the likelihood of the NSP for three Condorcet extensions (Black's rule, MaxiMin, and Tideman's rule) under various preference models using Ehrhart theory as well as extensive computer simulations. We find that, for few alternatives, the probability of the NSP is negligible (less than 1% for four alternatives and all considered preference models, except for Black's rule). As the number of alternatives increases, the NSP becomes much more likely and which rule is most susceptible to abstention strongly depends on the underlying distribution of preferences.

1 Introduction

Results from voting theory are increasingly used when dealing with collective decision making in computational multiagent systems [see, e.g. Rothe, 2015, Brandt et al., 2016a, Endriss, 2017]. A large part of the voting literature studies paradoxes in which seemingly mild properties are violated by common voting rules. Moreover, there are a number of sweeping impossibilities, which entail that there exists no optimal voting rule that avoids all paradoxes. It is therefore important to evaluate and compare how severe these paradoxes are in real-world settings. In this paper, we employ sophisticated analytical and experimental methods to assess the frequency of the No-Show Paradox (NSP), which occurs if a voter is better off by abstaining from an election. It is well-known that all Condorcet extensions, a large class of attractive voting rules, suffer from this paradox and this is often used as an argument against Condorcet extensions. Our analysis covers three Condorcet extensions: Black's rule, MaxiMin, and Tideman's rule.

In principle, quantitative results on voting paradoxes can be obtained via three different approaches. The analytical approach uses theoretical models to quantify paradoxes based on certain assumptions about the voters' preferences such as the *impartial anonymous culture (IAC)* model, in which every preference profile is equally likely. Analytical results usually tend to be quite hard to obtain and are limited to simple—and often unrealistic—assumptions. The experimental approach uses computer simulations based on underlying stochastic models of how the preference profiles are distributed. Experimental results have less general validity than analytical results, but can be obtained for arbitrary distributions of preferences. Finally, the empirical approach is based on evaluating real-world data to analyze how frequently paradoxes actually occur or how frequently they would have occurred if certain voting rules had been used for the given preferences. Unfortunately, only very limited real-world data for elections is available.

We analytically study the NSP under the assumption of IAC via *Ehrhart theory*, which goes back to the French mathematician Eugène Ehrhart [Ehrhart, 1962]. The idea of Ehrhart theory is to model the space of all preference profiles as a discrete simplex and then count the number of integer points inside of the polytope defined by the paradox in question. The number of these integer points can be described by so-called quasi- or Ehrhart-polynomials, which can be computed with the help of computers. The computation of the quasi-polynomials that arise in our context is computationally very demanding, because the dimension of the polytopes grows super-exponentially in the number of alternatives and was only made possible by recent advances of the computer algebra system NORMALIZ [Bruns et al.]. We complement these results by extremely elaborate simulations using four common preference models in addition to IAC (IC, urn, spatial, and Mallows). In contrast to existing results, our analysis goes well beyond three alternatives.

2 Related Work

The NSP was first observed by Fishburn and Brams [1983] for a voting rule called *single-transferrable vote (STV)*. Moulin [1988] later proved that all Condorcet extensions are prone to the NSP; the corresponding bound on the number of voters was recently tightened by Brandt et al. [2017]. Similar results were obtained for weak preferences and stronger versions of the paradox [Pérez, 2001, Duddy, 2014]. The NSP was also transferred to other settings including set-valued voting rules [see, e.g., Pérez et al., 2010, Brandl et al., 2015a], probabilistic voting rules [see, e.g., Brandl et al., 2015b, 2017a] and random assignment rules [Brandl et al., 2017b].[1]

The frequency of the NSP was first studied by Ray [1986], who, in line with Fishburn and Brams's classic paper, analyzed situations where STV can be manipulated in elections with three alternatives. A similar goal was pursued by Lepelley and Merlin [2000] who quantified occurrences of the NSP assuming preferences are distributed according to IC or IAC. However, in contrast to the present paper, Lepelley and Merlin employed different statistical techniques to estimate the likelihood of multiple variants of the paradox and focused on score-based runoff rules in elections with three alternatives.

The general idea to quantify voting paradoxes via IAC has been around since the formal introduction of this preference model by Gehrlein and Fishburn [1976]. Still, it took a good 30 years until the connection to Ehrhart theory [Ehrhart, 1962] was established by Lepelley et al. [2008]. We refer to Gehrlein and Lepelley [2011] for a more profound explanation of all details and an overview of results subsequently achieved. The step from three to four alternatives, i.e., from six to 24 dimensions, was only made possible through recent advances in computer algebra systems by De Loera et al. [2012] and Bruns and Söger [2015]. Brandt et al. [2016b] used a framework similar to ours to study the frequency of two single-profile paradoxes (the Condorcet Loser Paradox and the Agenda Contraction Paradox).

Plassmann and Tideman [2014] conducted computer simulations for various voting rules and paradoxes under a modified spatial model in the three-alternative case. To the best of our knowledge, this is—apart from Brandt et al. [2016b]—the only study of Condorcet extensions from a quantitative angle.

3 Preliminaries

Let A be a set of m alternatives and $N = \{1, \dots, n\}$ a set of voters. We assume that every agent $i \in N$ is endowed with a preference relation $\succ_i$ over the alternatives A. More formally, $\succ_i$ is a complete, asymmetric and transitive binary relation, $\succ_i \in A \times A$, which gives a strict ranking over the alternatives. If $x \succ_i y$, we say that i prefers x to y.

A *preference profile* $\succ$ is a tuple consisting of one preference relation per voter, i.e., $\succ = (\succ_1, \dots, \succ_n)$. By $\succ_{-i}$ we denote the preference profile resulting of voter i abstaining the election, $\succ_{-i} = (\succ_1, \dots, \succ_{i-1}, \succ_{i+1}, \dots, \succ_n)$.

For two alternatives $x, y \in A$ and a preference profile $\succ$ we define the *majority margin* $g_{xy}(\succ)$ as

$$g_{xy}(\succ) = |\{i \in N \colon x \succ_i y\}| - |\{i \in N \colon y \succ_i x\}|.$$

Whenever $\succ$ is clear from the context we only write g_{xy}. A *voting rule* is a function f mapping a preference profile $\succ$ to a single alternative, $f(\succ) \in A$.

[1]Interestingly, when considering set-valued or probabilistic voting rules, there are Condorcet extensions immune to the NSP under suitable assumptions [Brandl et al., 2015a, 2017a].

Condorcet Extensions. Alternative $x \in A$ is a *Condorcet winner* if it beats all other alternatives in pairwise majority comparisons, i.e., $g_{xy} > 0$ for all $y \in A \setminus \{x\}$. If a voting rule always selects the Condorcet winner whenever one exists, it is called a *Condorcet extension*. A wide variety of Condorcet extensions has been studied in the literature [see, e.g., Fishburn, 1977, Brandt et al., 2016a]. In this paper, we consider three Condorcet extensions: Black's rule, MaxiMin, and Tideman's rule. The main criteria for selecting these rules were *discriminability* (in order to minimize the influence of lexicographic tie-breaking), *simplicity* (to allow for Ehrhart analysis and because voters generally prefer simpler rules), and *efficient computability* (to enable rigorous and comprehensive simulations).[2] In the following, we briefly define the three rules.

Black's rule selects the Condorcet winner whenever one exists and otherwise returns a winner according to Borda's rule (Borda's rule itself is no Condorcet extension).

$$f_{\text{Black}}(\succ) \in \begin{cases} x & \text{if } x \text{ is a Condorcet winner in } \succ \\ \arg\max_{x \in A} \sum_{y \in A \setminus \{x\}} g_{xy} & \text{otherwise.} \end{cases}$$

The *MaxiMin* rule returns an alternative for which the worst pairwise majority comparison is maximal. Formally,

$$f_{\text{MaxiMin}}(\succ) \in \arg\max_{x \in A} \min_{y \in A \setminus \{x\}} g_{xy}.$$

Tideman's rule returns an alternative for which the sum of weighted pairwise majority defeats is minimal, i.e.,

$$f_{\text{Tideman}}(\succ) \in \arg\max_{x \in A} \sum_{y \in A \setminus \{x\}} \min(0, g_{xy}).^3$$

In order to obtain well-defined voting rules we employ lexicographic tie-breaking for all rules defined above. All presented voting rules can be computed in polynomial time and do not rely on the exact preference profile $\succ$ but only on the majority margins that can conveniently be represented by a skew-symmetric matrix or a weighted directed graph.

For the sake of illustration consider an example with $N = \{1, \dots, 7\}$ and $A = \{a, b, c, d\}$. $\succ$ shall be such that

$$a \succ_1 c \succ_1 d \succ_1 b, \quad b \succ_4 c \succ_4 d \succ_4 a, \quad d \succ_7 b \succ_7 a \succ_7 c$$

and $\succ_1 = \succ_2 = \succ_3$, $\succ_4 = \succ_5 = \succ_6$. The matrix of pairwise majority margins then evaluates to

$$(g_{xy})_{x,y \in A} = \begin{array}{c} \\ a \\ b \\ c \\ d \end{array} \begin{array}{c} \begin{array}{cccc} a & b & c & d \end{array} \\ \begin{pmatrix} 0 & -1 & 1 & -1 \\ 1 & 0 & 1 & -1 \\ -1 & -1 & 0 & 5 \\ 1 & 1 & -5 & 0 \end{pmatrix} \end{array}.$$

In the absence of a Condorcet winner, and due to lexicographic tie-breaking, we have $f_{\text{Black}}(\succ) = c$, $f_{\text{MaxiMin}}(\succ) = a$, and $f_{\text{Tideman}}(\succ) = b$.

[2]Note that other discriminating Condorcet extensions such as Kemeny's rule, Dodgson's rule, and Young's rule are NP-hard to compute [see, e.g., Brandt et al., 2016a].

[3]Tideman's rule is arguably the least well-known voting rule presented here. It was proposed by Tideman [1987] to efficiently approximate Dodgson's rule and is not to be confused with *ranked pairs* which is sometimes also called Tideman's rule. Also note that the 'dual' rule returning alternatives for which the sum of weighted pairwise majority *wins* is maximal is not a Condorcet extension.

Strategic Abstention. A voting rule f is *manipulable by strategic abstention* if there exist some N, A, and $\succ$ such that for some $i \in N$, $f(\succ_{-i}) \succ_i f(\succ)$. Given an occurrence of manipulability by strategic abstention, f is said to suffer from the *No-Show Paradox* (NSP) (for N, A, $\succ$). Slightly abusing notation, we also say that $\succ$ is prone to the NSP whenever f, N, and A are clear from the context. Black's rule, Maximin, and Tideman's rule, are Condorcet extensions and therefore manipulable by strategic abstention. The smallest examples for this require three, four, and four alternatives, respectively.

Stochastic Preference Models. When analyzing properties of voting rules, it is a common approach to sample preferences according to some underlying model. Various concepts to model preferences have been introduced over the years; we refer to, e.g., Critchlow et al. [1991] and Marden [1995] for a detailed discussion. We focus on three parameter-free models, *impartial culture (IC)* where each voter's preferences are drawn uniformly at random, *impartial anonymous culture (IAC)* where anonymous preference profiles are drawn uniformly at random, and the two-dimensional *spatial model* where we sample points in the unit square and their proximity determines the voters' preferences. Furthermore, we consider the *urn model* with parameter 10 and *Mallows' model* with $\phi = 0.8$.

The preference models we consider (such as IC, IAC, and the Mallows model) have also found widespread acceptance for the experimental analysis of voting rules within the multiagent systems and AI community [see, e.g., Aziz et al., 2013, Brandt and Seedig, 2014, Goldsmith et al., 2014, Oren et al., 2015, Brandt et al., 2016b].

4 Quantifying the No-Show Paradox

The goal in this paper is to quantify the frequency of the NSP, i.e., to investigate for how many preference profiles a voter is incentivized to abstain from an election. In order to achieve this goal, we employ an exact analysis via Ehrhart Theory and experimental analysis via sampled preference profiles.

4.1 Exact Analysis via Ehrhart Theory

The imminent strength of exact analysis is that it gives reliable theoretical results. On the downside, precise computation is only feasible for very simple preference models and even then only for small values of m. We focus on IAC and make use of Ehrhart theory.

First, note that an anonymous preference profile is completely specified by the number of voters sharing each of the $m!$ possible rankings on m alternatives. Hence, we can uniquely represent an anonymous profile by an integer point x in a space of $m!$ dimensions. We interpret x_i as the number of voters of type $\succ_i$, i.e., sharing preference ranking $\succ_i$. For fixed m, our goal is to describe all profiles that are prone to the NSP by using linear (in)equalities, i.e., as a polytope P_n.[4] Given that this is possible, the fraction of profiles prone to the NSP can be computed by dividing the number of integer points contained in P_n by the total number of profiles for n voters, i.e., the number of integer points x satisfying $x_i \geq 0$ for all $1 \leq i \leq m!$ and $\sum_{1 \leq i \leq m!} x_i = n$.

While the latter number is known to be $\binom{m!+n+1}{m!-1}$, the former can be determined using Ehrhart theory. Ehrhart [1962] shows that it can be found by so-called *Ehrhart-* or *quasi-polynomials* f—a collection of q polynomials f_i of degree d such that $f(n) = f_i(n)$ if $n \equiv i \mod q$. Obtaining f is possible via computer programs like LATTE [De Loera et al., 2004] or NORMALIZ [Bruns et al.]. Brandt et al. [2016b] give a more detailed description of the general methodology.

In order to illustrate this method, consider MaxiMin in elections with four alternatives under IAC. For the modeling we need to give linear constraints in terms of voter types—or equivalently majority margins—that describe polytopes containing all profiles prone to the NSP.

Recall the definition of MaxiMin from Section 3 and assume $f_{\text{MaxiMin}} = x$. For the NSP to occur, two intrinsic conditions have to be fulfilled: (i) There is a voter i such that $f_{\text{MaxiMin}}(\succ_{-i}) = y \neq x$ and (ii) for voter i, we have $y \succ_i x$. We find that for $A = \{a, b, c, d\}$, conditions (i) and (ii) entail that manipulation from a to b is only possible for $\succ_i$: c, b, a, d and $\succ_j$: d, b, a, c. It can be shown that no instance exists in which both voter types can influence the outcome in their favor. For the sake of this example, let us focus on $\succ_i$.

A first analysis shows that a's highest defeat has to be against d while b's highest defeat necessarily is against c with $g_{ad} = g_{bc}$,[5] and any other defeat of b lower by at least two. This gives rise to a first set of essential constraints.[6]

$$\begin{aligned} &g_{ad} = g_{bc}, && g_{ad} \leq 0, \\ &g_{ab} \geq g_{ad}, && g_{ba} \geq g_{ad} + 2 \qquad \text{(basis)} \\ &x_i \geq 1 \end{aligned}$$

At this point, we distinguish between $g_{cd} = 0$, $g_{cd} \leq -1$, and $g_{cd} \geq 1$. In case $g_{cd} = 0$, we trivially only need bounds on the defeats of c against a and d against b:

$$g_{cd} = 0, \qquad g_{ca} \leq g_{ad}, \qquad g_{db} \leq g_{ad} \qquad \text{(A)}$$

If $g_{cd} \leq -1$, c's highest defeat could be against a, d, or both. We consequently need a case distinction to accommodate for these possibilities.

$$\begin{aligned} &g_{cd} \leq -1, && g_{db} \leq g_{ab} && && \text{(B)} \\ &g_{cd} \leq g_{ad}, && g_{ca} \leq g_{ad} && && \text{(B.1)} \\ &g_{cd} \leq g_{ad}, && g_{ca} \geq g_{ad} + 1, && g_{ac} \geq g_{ad} && \text{(B.2)} \\ &g_{cd} \geq g_{ad} + 1, && g_{ca} \leq g_{ad} && && \text{(B.3)} \end{aligned}$$

For $g_{cd} \geq 1$ and an almost symmetric reasoning with reversed arguments for c and d we obtain (C), (C.1), (C.2), and (C.3).

Finally, the total set of profiles admitting a manipulation from a to b by i can be described by seven polytopes making use of the constraints developed above. We obtain

- $P_1 =$ (basis) + (A),

[4] More precisely, P_n is a *dilated polytope* depending on n, $P_n = nP = \{n\vec{x} \colon \vec{x} \in P\}$.

[5] Theoretically, we only require $g_{ad} - 1 \leq g_{bc} \leq g_{ad}$. As either all g_{xy} are even or all g_{xy} are odd, this collapses to $g_{ad} = g_{bc}$.

[6] Some inequalities are omitted to remove redundancies when taken together with later constraints.

- $P_2 = \text{(basis)} + \text{(B)} + \text{(B.1)}$, $P_3 = \text{(basis)} + \text{(B)} + \text{(B.2)}$, $P_4 = \text{(basis)} + \text{(B)} + \text{(B.3)}$,
- $P_5 = \text{(basis)} + \text{(C)} + \text{(C.1)}$, $P_6 = \text{(basis)} + \text{(C)} + \text{(C.2)}$, and $P_7 = \text{(basis)} + \text{(C)} + \text{(C.3)}$.[7]

Since we also need to consider $\succ_j$ and all other combinations of alternatives we undergo a similar reasoning 24 times which amounts to a total of 168 disjoint polytopes. Note that for the lexicographic tie-breaking and different types of manipulators, there are no exact symmetries that could be exploited to reduce this number.

This approach is substantially more involved than using Ehrhart theory for other paradoxes, e.g., the Condorcet Loser Paradox [Brandt et al., 2016b], because of three reasons.

(i) An occurrence of the NSP requires the presence of a certain type of voter.
(ii) Preference profiles for which different types of voters are able to manipulate must be counted only once.[8]
(iii) Possible manipulations not only rely on the winning alternative itself but on all majority margins that have to adhere to different constraints.

4.2 Experimental Analysis

In contrast to exact analysis, the experimental approach relies on simulations to grasp the development of different phenomena under varying conditions. On the upside, this usually allows for results for more complex problems or a larger scale of parameters, both of which might be prohibitive for exact calculations. At the same time we however face the problem that we need a huge number of simulations per setting to get sound estimates which in turn often requires a high-performance computer and a lot of time. Also, there remains the risk that even a vast amount of simulations fails to capture one specific, possibly crucial, effect.

Regarding the pivotal question of our paper, the frequency of the NSP for various voting rules, we sample preference profiles for different combinations of n and m using the modeling assumptions explained in Section 3.

5 Results and Discussion

In this section we present our results obtained by both exact analysis and computer simulations.

5.1 Results under IAC

We first focus on MaxiMin with four alternatives, as our modeling in Section 4.1 allows for an exact analysis of the NSP. The fraction of profiles prone to the NSP is depicted in Figure 1 together with an experimental analysis of the same question. There are four interesting observations to be made.

First, we note that the results obtained by simulation almost perfectly match the exact calculations which can be seen as strong evidence for the correctness of both. It additionally stands to reason that this accordance with the exact numbers also holds for larger m or even different rules, which is most useful for cases where determining the corresponding Ehrhart polynomials or even the modeling via polytopes is infeasible.

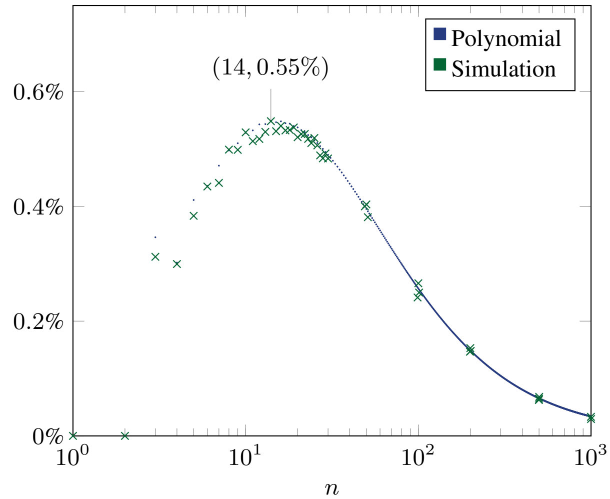

Figure 1: Profiles prone to the NSP for MaxiMin and $m = 4$

We see that the maximum is attained at 14 voters with 0.55% of all profiles suffering from the NSP. Hence, we can argue that for elections with four alternatives, the NSP hardly causes a problem, independently of the number of voters.

Furthermore, we note that the probability for the NSP to occur converges to zero as n goes to ∞; a behavior that holds true for all voting rules considered and all fixed m. Intuitively, this is to be expected as for larger electorates, a single voter's power to sway the result diminishes. This first idea can be confirmed by considering the respective modeling via polytopes. Each modeling will contain at least one equality constraint, e.g. in (basis) of f_{MaxiMin} in Section 4.1. Consequently, the polytopes describing profiles for which a manipulation is possible are of dimension at most $m! - 1$ meaning the number of those profiles can be described by a polynomial of n of degree at most $m! - 1$ [see also Ehrhart, 1962]. The total number of profiles, on the other hand, can equivalently be determined via a polynomial of degree $m!$ giving that the fraction of profiles prone to the NSP is upper-bounded by $O(1/n)$. Following the intuitive argument, similar behavior is to be expected for all reasonable preference models.

Finally, note that the computed numbers result from Ehrhart polynomials with period $q = 55\,440$, i.e., no two values in Figure 1 are computed via the same polynomial. It is thus even more remarkable that they form such a regular curve.

With respect to Black's rule we can also observe that the computed numbers are exactly in line with experimental results for $m = 3$. For $m = 4$, determining the Ehrhart polynomials for both Black's as well as Tideman's rule proved to be infeasible, even when using a custom-tailored version of

[7] We choose this informal notation for the sake of readability. It is to be understood in a way that P_1 is the polytope described by (in)equalities labelled (basis) as well as (A). We additionally assume for all polytopes that the sum of voters per type adds up to n and each type consists of a nonnegative amount of voters.

[8] As a matter of fact, this cannot occur for MaxiMin. It is, however, relevant for, e.g., Black's rule.

NORMALIZ and employing a high-performance cluster.[9] For all rules, $m \geq 5$ appears to be out of scope for years to come.

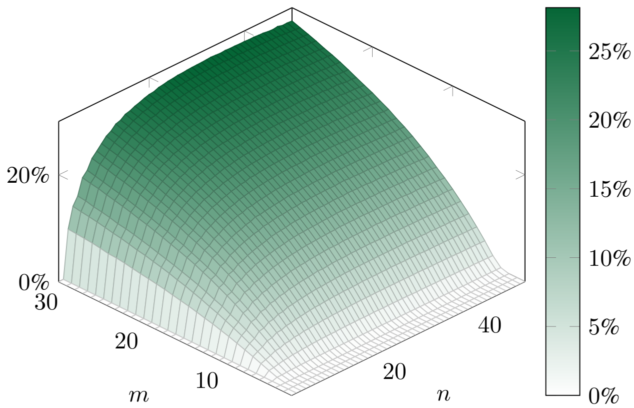

Figure 2: Profiles prone to the NSP for MaxiMin

We therefore rely on simulations to grasp how often the NSP can occur for different combinations of n and m up to 50 voters and 30 alternatives. Our results can be found in Figures 2 to 4. The following observations can be made.

To begin with, the relatively low fraction of profiles prone to the NSP for MaxiMin and $m = 4$ increases dramatically as m grows. This leads to the fact that for only 20 alternatives a rough fifth of all profiles admit a manipulation by abstention for a medium count of voters—a number too large to discard the NSP as merely a theoretical problem.

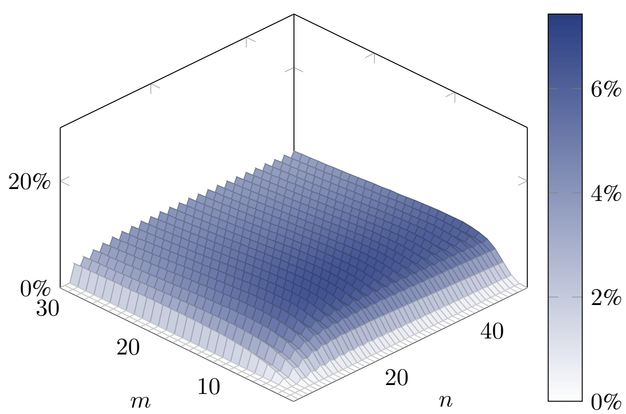

Figure 3: Profiles prone to the NSP for Black's rule

Especially when considering Black's and Tideman's rule we see that the parity of n crucially influences the results. However, the parity of n affects the fractions in opposite directions: higher fractions occur for Black's rule when n is even, in contrast to Tideman's rule where this happens when n is odd. For Black's rule, this is most probably due to the fact that there are more suitable profiles close to having a Condorcet winner ($g_{xy} = 0$) than profiles close to not having one ($g_{xy} = 1$).[10] For Tideman's rule we currently lack a convincing explanation for the observed behavior, mostly because it is hard to intuitively grasp when exactly a preference profile is manipulable. Regarding MaxiMin, the parity of n seems to have little effect on the numbers. More detailed analysis shows that this appearance is deceptive; when manipulating towards a lexicographically preferred alternative fractions are higher for even n while the contrary holds for manipulations towards a lexicographically less preferred alternative. In sum, these two effects approximately cancel each other out.

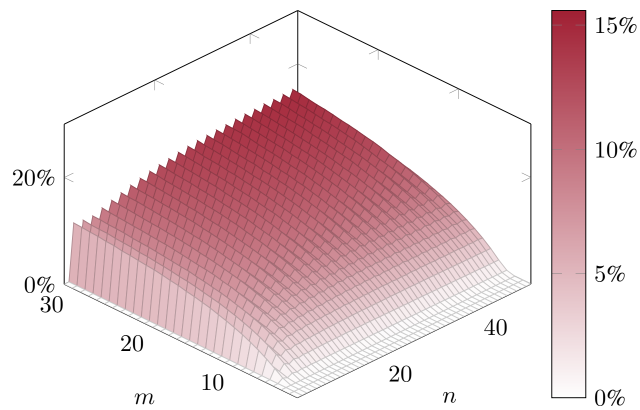

Figure 4: Profiles prone to the NSP for Tideman's rule

The flawless smoothness and regularity of Figures 2 to 4 are due to 10^6 runs per data point. This large number allows for all 95% confidence intervals to be smaller than 0.2%. Our simulations were conducted on a XeonE5-2697 v3 with 2 GB memory per job and took 35 to 48 hours for *each* data point. Since there are 1 500 data points per plot, the total runtime for all three figures easily accumulates to ten years on a single-processor machine.

5.2 Comparing Different Preference Models

In order to get an impression of the frequency of the NSP under different preference models we fix the number of alternatives to be $m = 4$ or $m = 30$ and sample 10^6 profiles for increasing n up to 1 000 or 200, respectively.[11] Figure 5 gives the fraction of profiles prone to the NSP using either MaxiMin, Black's, or Tideman's rule.

A close inspection of these graphs allows for multiple conclusions. First, we see that in particular Black's rule shows a severe dependency on the parity of n. For better illustration, we depict two lines per preference model to highlight this effect; which line stands for odd and which for even n is easiest checked using their corresponding point of intersection with the x-axis which is either 1, 2, or 3 throughout. Apart from

[9] For Black's rule we find that the polynomial would be of period $q \approx 2.7 \cdot 10^7$ corresponding to a mid two-digit GB file size.

[10] For Black's rule, manipulation is only possible either towards or away from a Condorcet winner since Borda's rule is immune to strategic abstention.

[11] For increasing m the computations quickly become very demanding. The values for $m = 30$ and $n > 99$ are determined with 50 000 runs each only. The size of all 95% confidence intervals is however still within 0.5%.

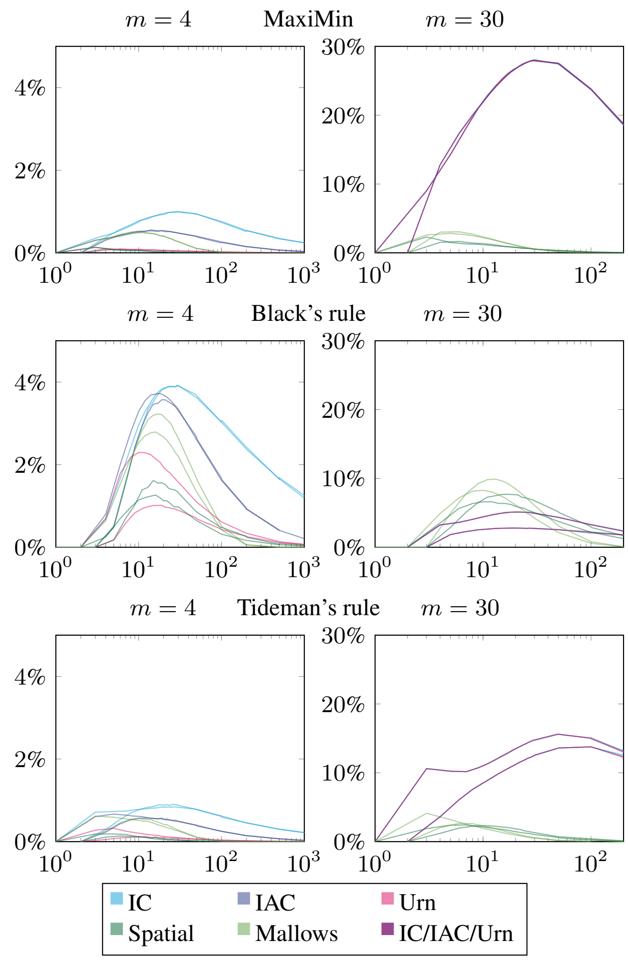

Figure 5: Profiles prone to the NSP for different rules, fixed m, and increasing n; two lines per preference model depending on the parity of n; IC, IAC and the urn model collapse for $m = 30$, resulting in a pink line.

explanations given earlier, it is not completely clear why differences are more prominent for some voting rules, why we sometimes see higher percentages for odd n and other times for even n, or why for some instances there is a large discrepancy for one preference model but hardly any for another.

IC and IAC are often criticized for being unrealistic and only giving worst-case estimates. This criticism is generally confirmed by our experiments which show that the highest fractions of profiles is prone to the NSP when they are sampled according to IC or IAC. A notable exception is Black's rule for 30 alternatives, where a different effect prevails: For many alternatives and comparably few voters situations in which a Condorcet winner (almost) exists appear less frequently under IC or IAC than under the other preference models. In absence thereof, Black's rule collapses to Borda's rule which is immune to the NSP. Note that were we to conduct a dual experiment with fixed n and increasing m, the fraction of profiles prone to the NSP using Black's rule and IC or IAC would converge to zero for similar reasons.

We moreover see that IC, IAC, and the urn model exhibit identical behavior for $m = 30$. The right-hand side of Figure 5 therefore seems to only feature three preference models, even though all five are depicted. This may be surprising at first but is to be expected since IC and IAC can equivalently be seen as urn models with parameters 0 and 1, respectively. For 30! $\approx 2.7 \cdot 10^{32}$ voter types the difference between parameters 0, 1, and 10 is simply too small for a visible difference.

Finally, Maximin appears to fare exceptionally bad with respect to the NSP and IC, IAC, and the urn model while a contrary behavior is visible for the spatial and Mallows model. Though generally in line with expectations, we currently do not have an explanation for the magnitude of this effect.

5.3 Empirical Analysis

We have also analyzed the NSP for empirical data obtained from real-world elections. Unfortunately, such data is generally relatively rare and imprecise and often only fragmentarily available. We have checked all 315 strict profiles contained in the PrefLib library [Mattei and Walsh, 2013] for occurrences of the NSP. Our evaluation shows that two profiles admit a manipulation by abstention when Black's rule is used and that no manipulation is possible for MaxiMin and Tideman's rule. While this suggests a low susceptibility to the NSP in real-world elections, much more data would be required to allow for meaningful conclusions.

6 Conclusion

We analyzed the likelihood of the NSP for three Condorcet extensions (Black's rule, MaxiMin, and Tideman's rule) under various preference models using Ehrhart theory as well as extensive computer simulations and some empirical data. Our main results are as follows.

- When there are few alternatives, the probability of the NSP is negligible (less than 1% for $m = 4$, MaxiMin, Tideman's rule, and all considered preference models; less than 4% for Black's rule).
- When there are 30 alternatives and preferences are modeled using IC, IAC, and the urn model, Black's rule is least susceptible to the NSP ($< 6\%$), followed by Tideman's rule ($< 16\%$) and Maximin ($< 29\%$).
- For 30 alternatives and the spatial and Mallows model, this ordering is reversed. Maximin is least susceptible ($< 4\%$), followed by Tideman's rule ($< 5\%$) and Black's rule ($< 10\%$).
- The parity of the number of voters significantly influences the manipulability of Black's and Tideman's rules. Black's rule is more manipulable for an even number of voters whereas Maximin is more manipulable for an odd number of voters.
- Whenever analysis via Ehrhart theory is feasible, the results are perfectly aligned with our simulation results, highlighting the accuracy of the experimental setup.
- Only two (out of 315) strict preference profiles in the PrefLib database are manipulable by strategic abstention (both manipulations only occur for Black's rule, but not for MaxiMin and Tideman's rule).

References

H. Aziz, S. Gaspers, N. Mattei, N. Narodytska, and T. Walsh. Ties matter: Complexity of manipulation when tie-breaking with a random vote. In *Proc. of 27th AAAI Conference*, pages 74–80. AAAI Press, 2013.

F. Brandl, F. Brandt, C. Geist, and J. Hofbauer. Strategic abstention based on preference extensions: Positive results and computer-generated impossibilities. In *Proc. of 24th IJCAI*, pages 18–24. AAAI Press, 2015a.

F. Brandl, F. Brandt, and J. Hofbauer. Incentives for participation and abstention in probabilistic social choice. In *Proc. of 14th AAMAS Conference*, pages 1411–1419. IFAAMAS, 2015b.

F. Brandl, F. Brandt, and J. Hofbauer. Welfare maximization entices participation. 2017a. Working Paper.

F. Brandl, F. Brandt, and J. Hofbauer. Random assignment with optional participation. In *Proc. of 16th AAMAS Conference*, pages 326–334. IFAAMAS, 2017b.

F. Brandt and H. G. Seedig. On the discriminative power of tournament solutions. In *Proc. of 1st AAMAS Workshop on Exploring Beyond the Worst Case in Computational Social Choice (EXPLORE)*, 2014.

F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2016a.

F. Brandt, C. Geist, and M. Strobel. Analyzing the practical relevance of voting paradoxes via Ehrhart theory, computer simulations, and empirical data. In *Proc. of 15th AAMAS Conference*, pages 385–393. IFAAMAS, 2016b.

F. Brandt, C. Geist, and D. 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.

W. Bruns and C. Söger. The computation of generalized Ehrhart series in Normaliz. *Journal of Symbolic Computation*, 68(2):75–86, 2015. Effective Methods in Algebraic Geometry.

W. Bruns, B. Ichim, T. Römer, R. Sieg, and C. Söger. Normaliz. Algorithms for rational cones and affine monoids. Available at `https://www.normaliz.uni-osnabrueck.de`.

D. E. Critchlow, M. A. Fligner, and J. S. Verducci. Probability models on rankings. *Journal of Mathematical Psychology*, 35:294–318, 1991.

J. A. De Loera, R. Hemmecke, J. Tauzer, and R. Yoshida. Effective lattice point counting in rational convex polytopes. *Journal of Symbolic Computation*, 38(4):1273–1302, 2004.

J. A. De Loera, B. Dutra, M. Köppe, S. Moreinis, G. Pinto, and J. Wu. Software for exact integration of polynomials over polyhedra. *ACM Communications in Computer Algebra*, 45(3/4):169–172, 2012.

C. Duddy. Condorcet's principle and the strong no-show paradoxes. *Theory and Decision*, 77(2):275–285, 2014.

E. Ehrhart. Sur les polyedres rationnels homothetiques á n dimensions. *Compte-Rendus de l'académie des Sciences*, 254(4):616, 1962.

U. Endriss, editor. *Trends in Computational Social Choice*. AI Access, 2017.

P. C. Fishburn. Condorcet social choice functions. *SIAM Journal on Applied Mathematics*, 33(3):469–489, 1977.

P. C. Fishburn and S. J. Brams. Paradoxes of preferential voting. *Mathematics Magazine*, 56(4):207–214, 1983.

W. V. Gehrlein and P. C. Fishburn. Concorcet's paradox and anonymous preference profiles. *Public Choice*, 26(1):1–18, 1976.

W. V. Gehrlein and D. Lepelley. *Voting Paradoxes and Group Coherence*. Studies in Choice and Welfare. Springer-Verlag, 2011.

J. Goldsmith, J. Lang, N. Mattei, and P. Perny. Voting with rank dependent scoring rules. In *Proc. of 28th AAAI Conference*, pages 698–704. AAAI Press, 2014.

D. Lepelley and V. Merlin. Scoring run-off paradoxes for variable electorates. *Economic Theory*, 17(1):53–80, 2000.

D. Lepelley, A. Louichi, and H. Smaoui. On Ehrhart polynomials and probability calculations in voting theory. *Social Choice and Welfare*, 30(3):363–383, 2008.

J. I. Marden. *Analyzing and Modeling Rank Data*. Number 64 in Monographs on Statistics and Applied Probability. Chapman & Hall, 1995.

N. Mattei and T. Walsh. PrefLib: A library for preference data. In *Proc. of 3rd ADT*, volume 8176 of *LNCS*, pages 259–270. Springer, 2013.

H. Moulin. Condorcet's principle implies the no show paradox. *Journal of Economic Theory*, 45(1):53–64, 1988.

J. Oren, Y. Filmus, and C. Boutilier. Efficient vote elicitation under candidate uncertainty. In *Proc. of 23rd IJCAI*, pages 309–316. AAAI Press, 2015.

J. Pérez. The Strong No Show Paradoxes are a common flaw in Condorcet voting correspondences. *Social Choice and Welfare*, 18(3):601–616, 2001.

J. Pérez, J. L. Jimeno, and E. García. No show paradox in condorcet k-voting procedure. *Group Decision and Negotiation*, 21(3):291–303, 2010.

F. Plassmann and N. Tideman. How frequently do different voting rules encounter voting paradoxes in three-candidate elections? *Social Choice and Welfare*, 42(1):31–75, 2014.

D. Ray. On the practical possibility of a 'no show paradox' under the single transferable vote. *Mathematical Social Sciences*, 11(2):183–189, 1986.

J. Rothe, editor. *Economics and Computation: An Introduction to Algorithmic Game Theory, Computational Social Choice, and Fair Division*. Springer, 2015.

T. N. Tideman. Independence of clones as a criterion for voting rules. *Social Choice and Welfare*, 4(3):185–206, 1987.