

Analyzing the Practical Relevance of Voting Paradoxes via Ehrhart Theory, Computer Simulations, and Empirical Data

Felix Brandt
TU München
Munich, Germany
brandtf@in.tum.de

Christian Geist
TU München
Munich, Germany
geist@in.tum.de

Martin Strobel
TU München
Munich, Germany
martin.r.strobel@gmail.com

ABSTRACT

Results from social choice theory are increasingly used to argue about collective decision making in computational multiagent systems. A large part of the social choice literature studies voting paradoxes in which seemingly mild properties are violated by common voting rules. In this paper, we investigate the likelihood of the Condorcet Loser Paradox (CLP) and the Agenda Contraction Paradox (ACP) using Ehrhart theory, computer simulations, and empirical data. We present the first analytical results for the CLP on four alternatives and show that our experimental results, which go well beyond four alternatives, are in almost perfect congruence with the analytical results. It turns out that the CLP—which is often cited as a major flaw of some Condorcet extensions such as Dodgson’s rule, Young’s rule, and MaxiMin—is of no practical relevance. The ACP, on the other hand, frequently occurs under various distributional assumptions about the voters’ preferences. The extent to which it is real threat, however, strongly depends on the voting rule, the underlying distribution of preferences, and, somewhat surprisingly, the parity of the number of voters.

Keywords

Social choice theory; voting; Ehrhart theory

1. INTRODUCTION

Results from social choice theory are increasingly used to argue about collective decision making in computational multiagent systems (see, e.g., [12, 8, 30, 9]). A large part of the social choice literature studies voting 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. As a consequence, much of the research in social choice theory is concerned with whether a paradox can appear for a given voting rule or not. However, it turns out that some paradoxes—while possible in principle—will almost never occur in practice.

An extreme example of this phenomenon was recently revealed for the voting rule *TEQ*. Due to its unwieldy recursive definition, it was unknown for more than 20 years

whether *TEQ* satisfies any of a number of very basic desirable properties. In 2011, Brandt et al. [7] have shown that *TEQ* violates all of these properties. However, their proof is non-constructive and only shows the existence of astronomically large counterexamples requiring about 10^{136} alternatives. While there are smaller computer-generated counterexamples [4], computer experiments have shown that these counterexamples are *extremely* rare and that *TEQ* satisfies the desirable properties for all practical purposes [6]. These findings motivated us to provide analytical, experimental, and empirical justifications for such statements.

In this paper, we study two voting paradoxes. The first is the well-known Condorcet loser paradox (CLP), which occurs when a voting rule selects the Condorcet loser, an alternative that loses against every other alternative in pairwise majority contests. Perhaps surprisingly, this paradox affects some Condorcet extensions, i.e., voting rules that are guaranteed to select an alternative that *wins* against every other alternative in pairwise majority contests. Common affected Condorcet extensions are Dodgson’s rule, Young’s rule, and MaxiMin [18]. The second paradox, called agenda contraction paradox (ACP), occurs when removing losing alternatives changes the set of winners. There are only few voting rules that do not suffer from this paradox, one of them being the essential set. In fact, all common voting rules that violate the CLP also violate the ACP.

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. 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.

Our main results are as follows.

Using Ehrhart theory, we compute upper bounds for the CLP as well as the exact probabilities under which the CLP occurs for MaxiMin when there are four alternatives and preferences are distributed according to the Impartial Anonymous Culture (IAC) distribution. This approach also

Appears in: *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2016)*, J. Thangarajah, K. Tuyls, C. Jonker, S. Marsella (eds.), May 9–13, 2016, Singapore.

Copyright © 2016, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

yields the exact limit probabilities (for CLP and ACP) when the number of voters goes to infinity. To the best of our knowledge, these are the first analytical results for the CLP on four alternatives (which is the minimal number of alternatives for which the voting rules we consider exhibit the CLP).

For both the CLP and the ACP, we thoroughly analyze a variety of other settings with more alternatives and other stochastic preference models using computer simulations. For those settings in which the analytical approach is also feasible, our results are in almost perfect congruence with the analytical results. This is strong evidence for the accuracy of our simulation results.

It turns out that the CLP—which is often cited as a major flaw of some Condorcet extensions—is of no practical relevance. The maximum probability under all preference models we studied is 2.2% (for MaxiMin, three voters, four alternatives, and IAC). In more realistic settings, it is much lower. For Dodgson’s rule, it never exceeds 0.01%. We did not find any occurrence of the paradox in real-world data, neither in the PREFLIB library [25] nor in millions of elections based on data from the Netflix Prize [3].

The ACP, on the other hand, frequently occurs under various distributional assumptions about the voters’ preferences. The extent to which it is real threat, however, strongly depends on the voting rule, the underlying distribution of preferences, and the parity of the number of voters. If the number of voters is much larger than the number of alternatives, less discriminating voting rules seem to fare better than more discriminating ones. For example, when there are 1,000 voters and four alternatives, the probability for the ACP under Copeland’s rule and IAC is 9% while it occurs with a probability of 33% for Borda’s rule. When there are fewer voters, the parity of the number of voters plays a surprisingly strong role. For example, if there are 6 alternatives, the ACP probability for Copeland’s rule is 44% for 50 voters, but only 26% for 51 voters. These results are in line with the empirical data we analyzed.

2. RELATED WORK

There is a huge body of research on the quantitative study of voting paradoxes. Gehrlein [19] focusses on the non-existence of Condorcet winners, arguably the most studied voting paradox. Gehrlein and Lepelley [20], on the other hand, provide an overview of many paradoxes and, in particular, analyze the influence of group coherence. In addition, Gehrlein and Lepelley [20] survey different tools and techniques that have been applied over the years for the quantitative study of voting paradoxes.

The analytical study of voting paradoxes under the assumption of IAC is most effectively done via Ehrhart theory, which goes back to the year 1962 and the French mathematician Eugène Ehrhart [16]. Interestingly, parts of these results have been reinvented (in the context of social choice) by Huang and Chua [22] in 2000, before Ehrhart’s original work was independently rediscovered for social choice by Wilson and Pritchard [34] and Lepelley et al. [23] more than forty years later.

Current research on the probability of voting paradoxes under IAC is based on algorithms that build upon Ehrhart’s results, such as the algorithm developed by Barvinok [2]. For many years, these approaches were limited to cases with

three or fewer alternatives. Recent advances in software tools and mathematical modeling enabled the study of elections with four alternatives. Bruns and Söger [10] and Schürmann [31] provide such results for Condorcet’s paradox, the Condorcet efficiency of plurality and the similarity between plurality and plurality with runoff. Schürmann [31] further shows how symmetries in the formulation of the paradoxes can be exploited to facilitate the corresponding computations.

For the CLP (sometimes also referred to as “Borda’s paradox”) many quantitative results are known [20], which are, however, limited to simple voting rules and scoring rules in particular. These results also include some empirical evidence for the paradox under plurality ([20], p.15) and suggest that it is an unlikely yet possible problem in practice. Interestingly, the CLP for *Condorcet extensions* has—to the best of our knowledge—only been considered by Plassmann and Tideman [28]. However, they restrict their analysis to the three-alternative case and find that the CLP never occurs, which is unsurprising since provably four alternatives are required for the Condorcet extensions they considered.

The ACP appears to have received less attention in the quantitative literature on voting paradoxes. Some limit probabilities for scoring rules were obtained by Gehrlein and Fishburn (see [20], p. 282–284). Fishburn [17] experimentally studied a variant of this paradox called “winner turns loser paradox” for Borda’s rule under Impartial Culture. For Condorcet extensions, Plassmann and Tideman [28] considered another variant of the ACP under a spatial model, but again limit their experiments to three alternatives. These few results already seem to indicate that the ACP might occur even under realistic assumptions. However, there are no results for more than three alternatives, Condorcet extensions, and the ACP in its full generality.

The preference models we consider (such as IC, IAC, and the Mallows- ϕ model) have also found widespread acceptance for the experimental analysis of voting rules within the multiagent systems and AI community (see, e.g., [1, 5, 21, 27]).

3. MODELS AND DEFINITIONS

Let A be a set of m alternatives and $N = \{1, \dots, n\}$ a set of voters. Each voter is equipped with a (*strict*) *preference relation* \succ_i , i.e., a connex,¹ transitive, and asymmetric binary relation on A . We read $x \succ_i y$ as voter i (strictly) preferring alternative x to alternative y .

A (*preference*) *profile* (or an *election*) is an n -tuple of preference relations and will be denoted by $R := (\succ_1, \dots, \succ_n)$. We will sometimes consider the restriction of \succ_i to a subset of alternatives $B \subseteq A$, called an *agenda*. Such a restriction will be denoted by $R|_B := (\succ_1|_B, \dots, \succ_n|_B)$.

3.1 Stochastic Preference Models

In this paper we consider five of the most common stochastic preference models. These models vary in their degree of realism. Impartial culture (IC) and impartial anonymous culture (IAC), for example, are usually considered as rather unrealistic. However, the simplicity of these models enables

¹A binary relation \succ_i on A is connex if $x \succ_i y \vee y \succ_i x$ for all $x \neq y \in A$. One may alternatively define \succ_i as the irreflexive component of a complete, antisymmetric, and transitive relation \succcurlyeq_i .

the use of analytical tools that cannot be applied to the other models. IC and IAC typically yield higher probabilities for paradoxes than other preference models and can therefore be seen as worst-case estimates (see, e.g., [29]). We only give informal definitions here; for more extensive treatments see, e.g., Critchlow et al. [14] and Marden [24].

Impartial culture The most widely-studied distribution is the so-called *impartial culture (IC)*, under which every possible preference relation has the same probability of $\frac{1}{m!}$. Thus, every preference profile is equally likely to occur.

Impartial anonymous culture In contrast to IC the *impartial anonymous culture (IAC)* is not based on the probabilities of individual preferences but on the probabilities of whole profiles. Under IAC one assumes that each possible *anonymous* preference profile on n voters is equally likely to occur. A more formal definition is given in Section 4.1.

Mallows- ϕ model In *Mallows- ϕ model*, the distance to a reference ranking (or ground truth) is measured by means of the Kendall-tau distance² and a parameter ϕ is used to indicate the dispersion. The case of $\phi = 1$ means absolute dispersion and coincides with IC, the case $\phi = 0$ corresponds to no dispersion and every voter always picks the “true” ranking. We chose $\phi = 0.8$ to simulate voters with relatively bad estimates, which leads to situations in which paradoxes are more likely to occur.

Pólya-Eggenberger urn model In the *Pólya-Eggenberger urn model*, each possible preference relation is represented by a ball in an urn from which individual preferences are drawn. After each draw, the chosen ball is put back and $\alpha \in \mathbb{N}_0$ new balls of the same kind are added to the urn. While the urn model subsumes both impartial culture ($\alpha = 0$) and impartial anonymous culture ($\alpha = 1$), we set $\alpha = 10$ to obtain a reasonably realistic interdependence of individual preferences.

Spatial model In the *spatial model* alternatives and agents are placed in a multi-dimensional space uniformly at random and the agents’ preferences are then determined by the Euclidean distances to the alternatives (closer alternatives are preferred to more distant ones). The spatial model is considered particularly realistic in political science where the dimensions are interpreted as different aspects of the alternatives [32]. We chose the simple case of two dimensions for our analysis.³

3.2 Voting Rules

A *voting rule* is a function f that maps a preference profile to a non-empty *set of winners*.

For a preference profile R , let $g_{xy} := |\{i \in N : x \succ_i y\}| - |\{i \in N : y \succ_i x\}|$ denote the *majority margin* of x against y . A very influential concept in social choice is the notion of a Condorcet winner, an alternative that wins against any other alternative in a pairwise majority contest. Alternative x is a *Condorcet winner (CW)* of a profile R if $g_{xy} > 0$ for all $y \in A \setminus \{x\}$. Conversely, alternative x is a *Condorcet loser (CL)* if $g_{yx} > 0$ for all $y \in A \setminus \{x\}$. Neither CWs nor CLs

²The Kendall-tau distance counts the number of pairwise disagreements.

³In a related study, Brandt and Seedig [5] have found that the number of dimensions does not seem to have a large impact on the results as long as it is at least two.

necessarily exist, but whenever they do they are unique. A voting rule f is called a *Condorcet extension* if $f(R) = \{x\}$ whenever x is the CW in R .

In the following paragraphs we briefly introduce the voting rules considered in this paper.

Borda’s Rule Under *Borda’s rule* each alternative receives from 0 to $|A| - 1$ points from each voter (depending on the position the alternative is ranked in). The alternatives with highest accumulated score win.

MaxiMin The *MaxiMin* rule is only concerned with the highest defeat of each alternative in a pairwise majority contest. It yields all alternatives as winners which have the maximal value of $\min_{y \in A} g_{xy}$.

Young’s Rule *Young’s rule* yields all alternatives that can be made a CW by removing a minimal number of voters.

Dodgson’s Rule *Dodgson’s rule* selects all alternatives that can be made a CW by a minimal number of pairwise swaps of adjacent alternatives in the individual preference relations.

Essential Set Consider the symmetric two-player zero-sum game G given by the skew-symmetric matrix with entries g_{xy} for all pairs of alternatives x, y . The *essential set* is the set of all alternatives that are played with positive probability in some mixed Nash equilibrium of G .⁴

Except for Borda’s rule, all presented voting rules are in fact Condorcet extensions. While Borda’s rule, MaxiMin, and the essential set can be computed efficiently, Young’s rule and Dodgson’s rule have been shown to be complete for parallel access to NP. The essential set is one of the few voting rules that do suffer from neither the CLP nor the ACP, and is merely included as a reference. For more formal definitions and computational properties of these rules, we refer to Brandt et al. [9].

3.3 Voting Paradoxes

In this paper we focus on two voting paradoxes whose occurrence can be determined given a voting rule f and a preference profile R .

Let f be a voting rule. Formally, a (*voting*) *paradox* is a characteristic function that maps a preference profile to 0 or 1. In the latter case, we say the paradox *occurs* for voting rule f at profile R .

The *Condorcet Loser Paradox (CLP)* occurs when a voting rule selects the CL as a winner.

DEFINITION 1. Given a voting rule f the Condorcet loser paradox CLP_f is defined as

$$CLP_f(R) = \begin{cases} 1 & \text{if } f(R) \text{ contains a CL} \\ 0 & \text{otherwise.} \end{cases}$$

The *agenda contraction paradox (ACP)* occurs when reducing the set of alternatives, by eliminating unchosen alternatives, influences the outcome of an election.

DEFINITION 2. Given a voting rule f the agenda contraction paradox ACP_f is a paradox defined as

$$ACP_f(R) = \begin{cases} 1 & \text{if } f(R|_B) \neq f(R) \text{ for some } B \supseteq f(R) \\ 0 & \text{otherwise.} \end{cases}$$

⁴These mixed equilibria are also known as *maximal lotteries* in probabilistic social choice.

4. QUANTIFYING VOTING PARADOXES

In this section we present the three general approaches for quantifying voting paradoxes: the analytical approach via Ehrhart theory, the experimental approach via computer simulations, and the empirical approach via real-world data.

4.1 Exact Analysis via Ehrhart Theory

Anonymous preference profiles only count the number of voters for each of the $m!$ possible rankings on m alternatives. An anonymous preference profile can hence be viewed as an integer point in a space of $d := m!$ dimensions. Formally, the set $S_{m,n}$ of anonymous preference profiles on m alternatives with n voters can be identified with the set of all integer points $z = (z_1, \dots, z_{m!}) \in \mathbb{Z}^{m!}$ which satisfy

$$z_i \geq 0 \text{ for all } i \in \{1, \dots, m!\}, \text{ and } \sum_{i=1}^{m!} z_i = n.$$

Under IAC each anonymous preference profile is assumed to be equally likely to occur. Hence, in order to determine the probability of a paradox *under IAC* it is enough to compute the number of points belonging to preference profiles in which the paradox occurs and compare them to the total number of points in $S_{m,n}$, which is known to be $|S_{m,n}| = \binom{m!+n-1}{m!-1}$.⁵

In this framework, many paradoxes X can be described with the help of linear constraints, i.e., the set of points belonging to the event can be described with the help of (in)equalities, a polytope. For variable n , this approach then describes a *dilated polytope* $P_n = nP := \{n\vec{x} : \vec{x} \in P\}$. Hence, we know that the probability of a paradox X_n under IAC is given by:

$$\mathbb{P}(X_n) = \frac{|nP \cap \mathbb{Z}^d|}{|S_{m,n}|}$$

and we can determine the probability of (many) voting paradoxes under IAC by evaluating the function $L(P, n) := |nP \cap \mathbb{Z}^d|$, which describes the number of integer points inside the dilation nP . This can be done with the help of Ehrhart theory. Ehrhart [16] was the first to show that $L(P, n)$ can be described by special functions, called quasi- or Ehrhart-polynomials. A function $f : \mathbb{Z} \rightarrow \mathbb{Q}$ is a *quasi-polynomial* of degree d and period q if there exists a list of q polynomials $f_i : \mathbb{Z} \rightarrow \mathbb{Q}$ ($0 \leq i < q$) of degree d such that $f(n) = f_i(n)$ if $n \equiv i \pmod{q}$.

Quasi-polynomials can be determined with the help of computer programs such as LATTE (De Loera et al. [15]) or NORMALIZ (Bruns et al. [11]). Unfortunately, the computation of our quasi-polynomials is computationally very demanding, especially because the dimension of the polytopes grows super-exponentially in the number of alternatives. This limits analytical results under IAC to rather small numbers of alternatives. To the best of our knowledge, NORMALIZ is the only program which is able to compute polytopes corresponding to elections with up to *four* alternatives. And even NORMALIZ is not always able to compute the whole quasi-polynomial, but sometimes we had to

⁵For most preference models other than IAC this approach does not work. While for specific combinations of (simple) distributions and voting rules there are some highly tailor-made computations in the literature (cf. Section 2), these have to be redesigned for each individual setting.

Paradox	Voting rule(s)	Result
CLP	Condorcet extensions	upper bound ($\forall n \in \mathbb{N}$)
	MaxiMin	probability ($\forall n \in \mathbb{N}$)
	Tideman's rule	limit prob. ($n \rightarrow \infty$)
ACP	MaxiMin	limit prob. ($n \rightarrow \infty$)

Table 1: Theoretical results obtained via Ehrhart theory (for four alternatives and under IAC)

resort to computing the leading coefficients only of the polynomial, which fortunately suffices for determining the limit probability of a paradox when the number of voters goes to infinity.

An overview of our analytical findings obtained in this way is provided in Table 1.

Finding a Quasi-polynomial for MaxiMin

As an example for the method just described, we consider the $\text{CLP}_{\text{MaxiMin}}$ in four-alternative elections under IAC, the probabilities of which can be computed from a quasi-polynomial with degree 23 and a period of 5,040.⁶

In order to determine the polynomial, we first need to describe the corresponding polytope with equalities and inequalities. Recall the definition of MaxiMin from Section 3.2:

$$f_{\text{MaxiMin}}(R) := \arg \max_{x \in A} \min_{y \in A} g_{xy}.$$

For $\text{CLP}_{\text{MaxiMin}}(R) = 1$ the CL of R has to have the lowest highest defeat. Formally, there is $x \in A$ such that for all $y \in A \setminus \{x\}$,

$$g_{yx} > 0, \text{ and} \tag{1}$$

$$\max_{z \in A \setminus \{x\}} g_{zx} \leq \max_{z \in A \setminus \{y\}} g_{zy}. \tag{2}$$

Now let $A = \{a, b, c, d\}$ and assume $x = d$. We then have that $g_{ad}, g_{bd}, g_{cd} > 0$, which implies $\max_{z \in A \setminus \{d\}} g_{zd} > 0$. Furthermore,

$$\max_{z \in A \setminus \{y\}} g_{zy} > 0 \text{ for all } y \in \{a, b, c\},$$

from which it follows that either $g_{ab}, g_{bc}, g_{ca} > 0$ or $g_{ba}, g_{cb}, g_{ac} > 0$. In both cases there is a majority cycle between a , b , and c . Due to symmetry we can choose one direction of the cycle arbitrarily and assume $g_{ab}, g_{bc}, g_{ca} > 0$. Then,

$$\max_{z \in A \setminus \{a\}} g_{za} = g_{ca}, \max_{z \in A \setminus \{b\}} g_{zb} = g_{ab}, \text{ and } \max_{x \in A \setminus \{c\}} g_{xc} = g_{ac}.$$

Condition (1) is already represented in the form of linear inequalities. In order to model condition (2) we determine $\max_{z \in A \setminus \{d\}} g_{zd}$ and distinguish cases for the seven possible outcomes. The inequalities for the case $\max_{z \in A \setminus \{d\}} g_{zd} = \{g_{ad}\}$ are

$$g_{ad} - g_{bd} > 0 \quad \text{and} \quad g_{ad} - g_{cd} > 0.$$

⁶In theory, the analysis can be adapted to also cover more complex rules (e.g., Dodgson's and Young's rule, which involve solving an ILP). It is unclear, however, how one would translate their definitions to linear inequalities.

Condition (2) furthermore yields

$$g_{ca} - g_{ad} \geq 0, \quad g_{ab} - g_{ad} \geq 0, \quad \text{and} \quad g_{bc} - g_{ad} \geq 0.$$

Each case belongs to a different polytope and the polytopes are pairwise distinct, so we can compute each quasi-polynomial separately and later combine them to one. To get the final polynomial we have to multiply by eight for the four different possible choices of a CL and the two possible directions of the majority cycle. This then enables us to efficiently evaluate the *exact* probabilities for any number of voters. The results are depicted in Figure 2. The leading coefficient of the quasi-polynomial can also be used to determine the limit probability which is given by

$$\begin{aligned} & \mathbb{P}(\text{CLP}_{\text{MaxiMin}} = 1 \mid m = 4, n \rightarrow \infty) \\ &= 8 \cdot \frac{485052253637930099}{6443662124777472000000} \approx \underline{\underline{0.06\%}}. \end{aligned}$$

4.2 Experimental Analysis

As we will see, simulating elections with the help of computers is a viable way of achieving very good approximations for the probabilities we are looking for. It even turns out that the results of our simulations are almost indistinguishable from the theoretical result obtained via Ehrhart theory (with the exception of the limit case, which cannot be realized via simulations).

More specifically, the experimental approach works as follows: a profile source creates random preference profiles according to a specific preference model. The profiles are then used to compute the winner(s) according to a given voting rule and to determine if the paradox occurs. Any such experiment is carried out for each pair of n and m and repeated frequently. In many cases in which we covered a wide range of voters, we did not consider every possible value of n but, more economically, only simulated the values: 1–30, 49–51, 99–101, 199–201, 499–501, 999–1,001.

In contrast to many other studies, we are concerned about the statistical significance of our experimental results. Thus, we also computed 99%-confidence intervals for each data point we generated. To this end, we used the `binofit` function in MATLAB which is based on the standard approach by Clopper and Pearson [13]. It shows that, based on our sampling rate of 10^5 and 10^6 , respectively, the 99%-confidence intervals are pleasantly small. Hence, even though they are depicted in all of the figures throughout this paper, sometimes it can be difficult to recognize them.

4.3 Empirical Analysis

The most valuable quantification of voting paradoxes would be their actual frequency in real-world elections. As mentioned before, real-world election data is generally relatively sparse, incomplete, and inaccurate. This makes empirical research on this topic rather difficult. Otherwise, the empirical approach strongly resembles the experimental approach.

For this paper we used two sources of empirical data. First, we used the 314 profiles with strict order preferences from the PREFLIB library [25]. Second, we had access to the 54,650 preference profiles over four alternatives without a CW which belong to the roughly 11 million four-alternative elections which Mattei et al. [26] derived from the Netflix Prize data [3]. Non-existence of Condorcet winners is a prerequisite for the paradoxes we study.

5. CONDORCET LOSER PARADOX

In this section we present our findings on the CLP. We conclude that—even though the CLP is possible in principle—it is so unlikely that it cannot be used as a serious argument against any of the Condorcet extensions we considered.

5.1 An Upper Bound

Before analyzing the CLP for concrete voting rules, we discuss an upper bound valid for all Condorcet extensions. For a Condorcet extension to choose the CL a profile obviously has to satisfy two conditions. First, there has to exist a CL in the profile. Second, no CW may exist in the profile. In the case of four-alternative elections—which is the first interesting case—we can compute the quasi-polynomial via Ehrhart theory and hence know the exact probabilities for any number of voters. The derivation and presentation of the quasi-polynomial, which has degree 23 and contains 24 polynomials, is omitted due to space constraints. The resulting probabilities for up to 1 000 voters—and a comparison with the results of an experimental analysis—can be obtained from Figure 1. The value of the limit probability is approximately 8%.

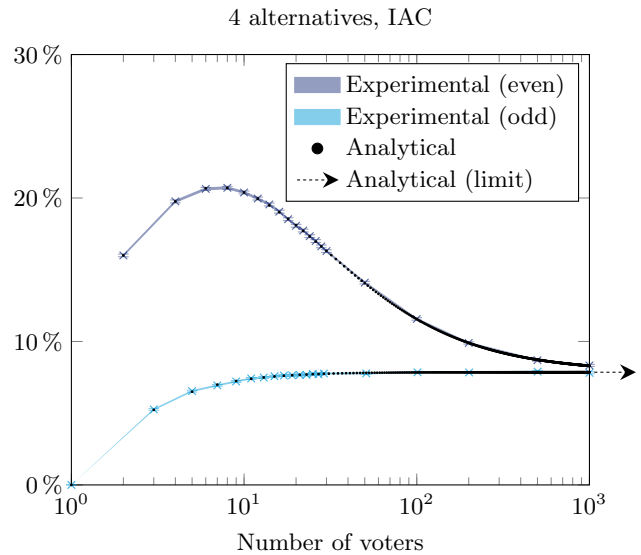


Figure 1: Probability of the event that a Condorcet extension could choose a CL in four-alternative elections under IAC

Especially for small even numbers of voters, where the probability is around 20%, the upper bound is too high to discard the CLP for Condorcet extensions altogether, and even the limit probability of 8% is relatively large. Also, for an increasing number of alternatives this problem does not vanish (for elections with 50 and 51 one voters and up to 100 alternatives the probabilities range between 5% and 25%).⁷

Note that differences between odd and even number of voters were to be expected since even numbers allow for majority ties, which have significant consequences for the paradoxes; this effect decreases for larger electorates. In the

⁷These upper bounds turn out to be relatively independent from the underlying preference distribution (among the models we considered, cf. Section 5.3).

specific case under consideration, the upper bound is generally higher for an even number of voters because the much higher likelihood of not having a CW more than counterbalances the lower likelihood of having a CL.

5.2 Results under IAC

Despite the high upper bounds from the previous section, the picture is quite clear for concrete Condorcet extensions: even under IAC, the risk of the considered Condorcet extensions selecting the CL is very low, as shown in Figure 2 and Figure 3 for four-alternative elections. The highest probability was found for $CLP_{MaxiMin}$ with 2.2% for three voters (CLP_{Young} with about 0.9%). The limit probability of $CLP_{MaxiMin}$, with 0.06% is so low that for sufficiently large electorates it would occur in only one out of 10,000 elections. The same seems to hold for the limit probability for CLP_{Young} . The probability of $CLP_{Dodgson}$ is even significantly lower, with a maximum of about 0.01% in elections with 9,999 voters. We could determine the limit probability of 0.01% only for an approximation of Dodgson’s rule by Tideman [33], which seems to be close to that for Dodgson’s rule, based on our experimental data.

When increasing the number of alternatives the probabilities drop even further. For elections with more than ten alternatives they reach a negligibly small level of less than 0.005% for all considered rules and in no simulations with twelve or more alternatives we could find any occurrence of the paradox.

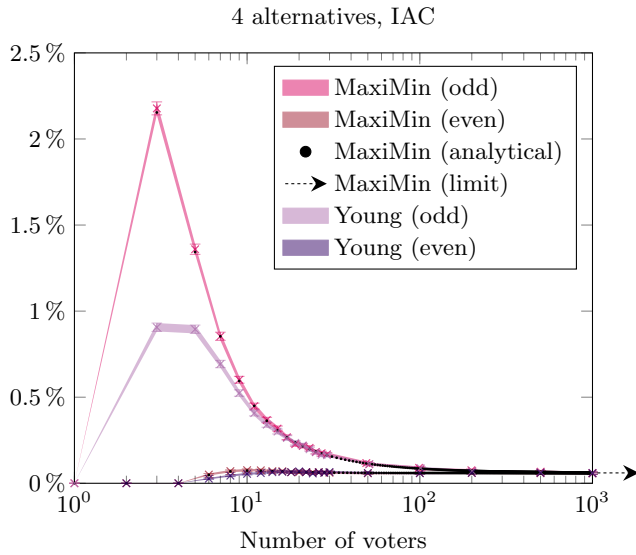


Figure 2: Comparison between CLP probabilities for MaxiMin and Young’s rule under IAC in four-alternative elections

5.3 Results under Other Preference Models

Figure 4, as one would expect, shows that under more realistic assumptions the probability of the CLP decreases further in four-alternative elections with 50/51 voters, with the highest probability occurring under the unrealistic assumption of IC and the lowest probability under what may be the most realistic model in many settings, the spatial

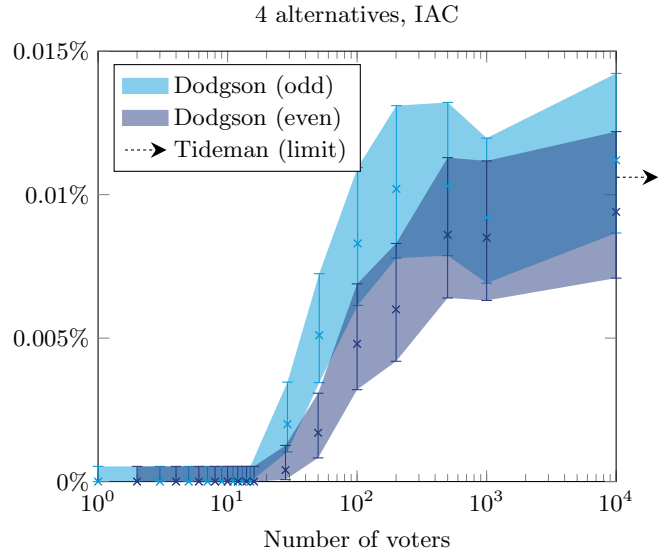


Figure 3: CLP probabilities for Dodgson’s rule under IAC in four-alternative elections; the seemingly large confidence intervals are due to the small scale of the graph.

model. In our experiments, Dodgson’s rule *never* selected a CL in the spatial model.

Similarly, we could not find any occurrence of the CLP in real-world data, which may be considered the strongest evidence that the CLP virtually never materializes in practice.⁸

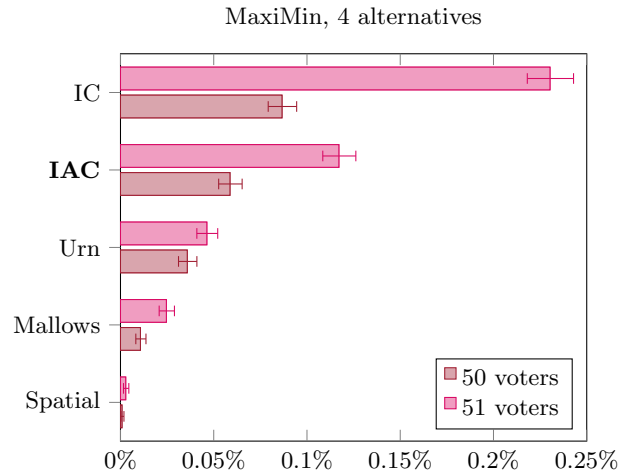


Figure 4: CLP probabilities in four-alternative elections for varying preference models and MaxiMin

⁸We tested 314 preference profiles with strict orders from the PREFLIB library as well as the roughly 11 million four-alternative elections which Mattei et al. [26] derived from the Netflix Prize data. While about 54,000 of those elections were susceptible to the CLP, it never occurred under the rules we considered in this paper. In contrast, under plurality it already occurred in twelve out of the 314 PREFLIB-instances.

6. AGENDA CONTRACTION PARADOX

Recall that the *agenda contraction paradox* (ACP) occurs when a reduced set of alternatives (created by the unavailability of losing alternatives) influences the outcome of an election. For many cases, it may be considered a generalization of the CLP as the following argument shows. Suppose the CL x is uniquely selected by a voting rule which implements majority rule on two-alternative choice sets. Then restricting A to $\{x, y\}$ for some alternative $y \neq x$ yields the new winner y (since $g_{xy} > 0$).

As we will see, the ACP is much more of a practical problem than the CLP. The picture, however, is not black and white. Whether or not it is a serious threat depends on the voting rule, the underlying preference distribution, and on the parity of the number of voters.

6.1 Varying Voting Rules

The ACP probability strongly varies for different voting rules (see Figure 5). Borda's rule generally exhibits the worst behavior of the rules studied, with probabilities of up to 56%, and with 34% for large electorates with 1,000 voters. In contrast, Copeland's rule is quite robust to the ACP for large electorates (with only about 8% occurrence probability for 1,000 voters).⁹

The reason for this gap between Borda's and Copeland's rule appears to be two-fold: First, Condorcet extensions are safe from this paradox as long as a CW exists; Borda's rule, by contrast, is not. Second, the discriminatory power of voting rules (i.e., their ability to select small winning sets) strongly supports the paradox. As soon as a single majority-dominated alternative is selected, the ACP has to occur. For large numbers of voters, this is in line with Copeland's rule being least discriminating among those evaluated. The essential set is among the most discriminating known voting rules immune to the ACP, but presumably less discriminating than Copeland's rule.

The behavior of MaxiMin is almost identical to that of Young's and Dodgson's rule. Confirming our approximate "limit" results of 1,000 voters, we were able to analytically compute the limit probability for MaxiMin as $\frac{331}{2048} \approx 16\%$. This is in perfect congruence with the (rounded) values for MaxiMin, Young's rule, and Dodgson's rule.

It should also be noted that with fewer than 100 voters, the parity of the number of voters plays a major role. For even numbers, significantly higher probabilities arise (which is particularly true for Copeland's rule, see above). At least part of this can be explained by a reduced probability for CWs in these cases.

For more alternatives (see the right-hand side of Figure 5), the relative behavior remains vastly unchanged with probabilities further increasing to values larger than 40% to 80% (mostly since the likelihood of a CW decreases roughly at the same rate).

6.2 Varying Preference Models

Figure 6 extends the analysis of the previous section by additionally considering preference models beyond IAC. The overall picture regarding the different rules remains the same. For large electorates Copeland's rule outperforms the other rules, whereas Borda's rule performs worst.

⁹For small *even* numbers of voters, Copeland's rule also frequently fails agenda contraction, which is also visible in Figure 6 and explains the seemingly high values in Table 2.

Regarding the different preference models, three classes emerge from Figure 6.

First, for Mallows- ϕ we observe probabilities that are vanishing with increased numbers of voters. Under the spatial model this is true as well, with the surprising exception of Borda's rule, for which the picture looks completely different and the probability does not go below 20% in the spatial model. Presumably, this can be explained by Borda's inability to select the CW in this setting, a hypothesis that deserves further study, however. On the contrary, the other rules appear to be benefitting from the fact that the existence of a CW becomes very likely under models with high voter interdependence.

Second, as expected, the assumption of IC serves as an upper bound for all other preference models. The results for IAC are not much lower, fostering the impression that IAC could also be an unrealistic upper bound.

Third, the urn model yields much lower values compared to IAC and IC. The absolute numbers, however, are still beyond acceptable levels (between 4% and 23% for 1,000 voters).

The findings in the empirical data corroborate our experimental findings. In PREFLIB the ACP occurs 17 times for Borda, three times for Copeland and exactly once for MaxiMin as well as Young's and Dodgson's rule. In the Netflix data set, where the number of voters is at least 350, Copeland performs much better than the other Condorcet extensions (4,400 compared to 18,470 occurrences for the other Condorcet extensions). Borda's rule virtually always suffers from the ACP on this data set: there are 54,620 instances of ACPs already when considering profiles that do not have a CW (there are 54,650 of such).

7. CONCLUSION

We investigated the likelihood of the CLP and the ACP using Ehrhart theory, computer simulations, and empirical data. The CLP is often cited as a major flaw of some Condorcet extensions such as Dodgson's rule, Young's rule, and MaxiMin. For example, Fishburn regards Condorcet extensions that suffer from the CLP (specifically referring to the three rules mentioned above) as "dubious extensions of the basic Condorcet criterion" ([18], p. 480).¹⁰ While this is intelligible from a theoretical point of view, our results have shown that the CLP is of virtually no practical concern. The ACP, on the other hand, frequently occurs under various distributional assumptions about the voters' preferences. The extent to which it is real threat, however, strongly depends on the voting rule, the underlying distribution of preferences, and, surprisingly, the parity of the number of voters. Our main quantitative results for the worst case are summarized in Table 2. Potential future work includes the analysis of other voting paradoxes (such as monotonicity failures or the no-show paradox) and other rules (such as Nanson's rule or Black's rule).

Acknowledgments

Christian Geist was supported by Deutsche Forschungsgemeinschaft under grant BR 2312/9-1. The authors also thank Nicholas Mattei for providing the Netflix Prize data and Christof Söger for his guidance regarding NORMALIZ.

¹⁰Fishburn [18] actually analyzes violations of "Smith's Condorcet principle", which is weaker than the CLP.

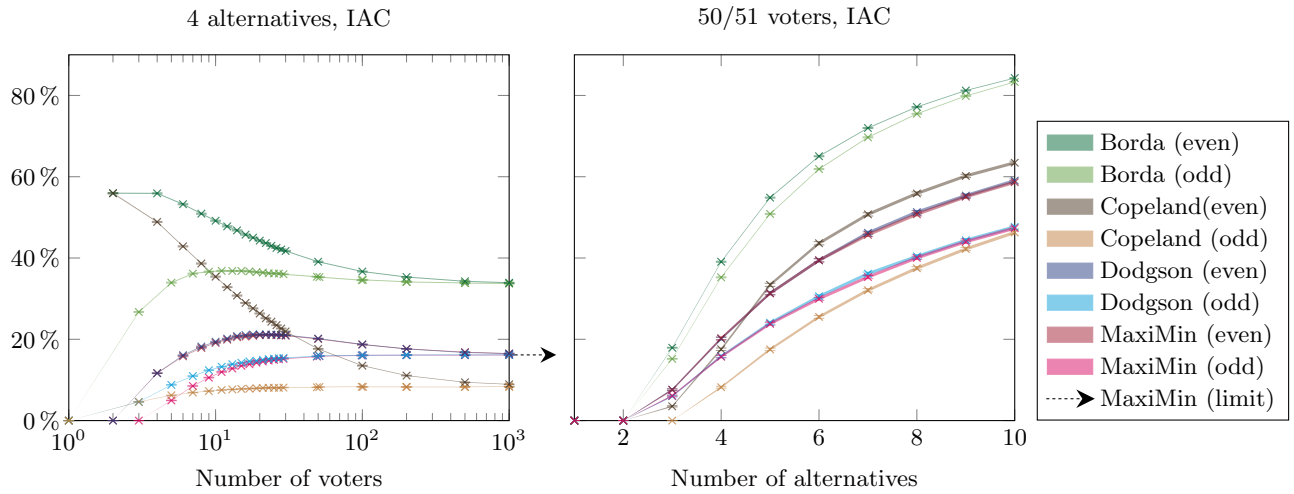


Figure 5: Comparison between ACP probabilities for different voting rules under IAC

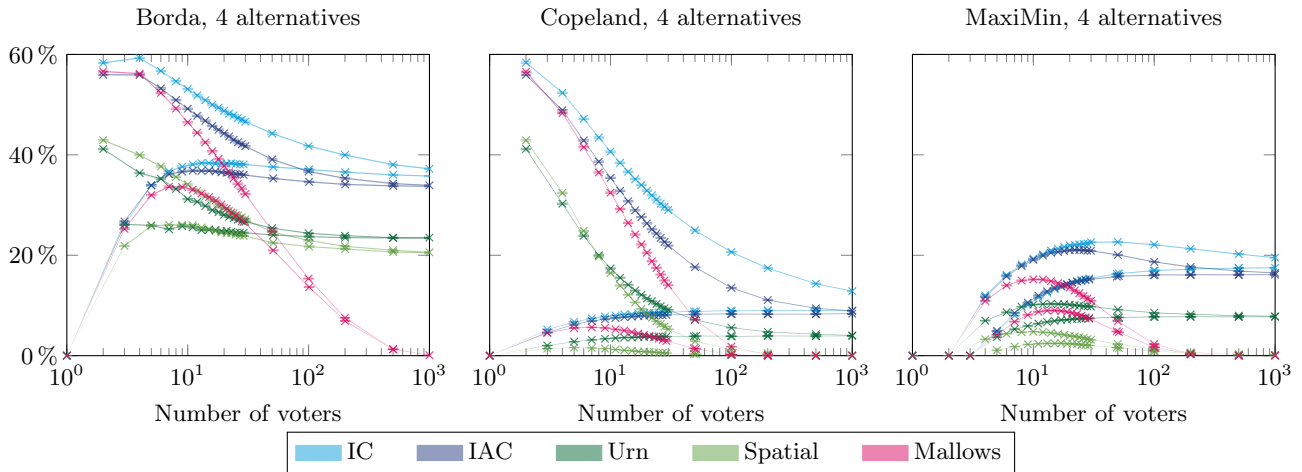


Figure 6: Comparison between ACP_{Borda} , ACP_{MaxiMin} and ACP_{Copeland} for varying preference models in four-alternative elections; the values of ACP_{Young} and ACP_{Dodgson} are omitted since they strongly resemble the ones of ACP_{MaxiMin} .

Paradox	Condorcet loser paradox (CLP)			Agenda contraction paradox (ACP)		
	IAC		IC	IAC		IC
Model	$\{1, \dots, 1000\}$	$\{50, 51\}$	$\{50, 51\}$	$\{1, \dots, 1000\}$	$\{50, 51\}$	$\{1, \dots, 1000\}$
n						
m	4		4	4		4
Essential set	0%	0%	0%	0%	0%	0%
Borda	0%	0%	0%	56%	84%	59%
Copeland	0%	0%	0%	56%	63%	58%
Dodgson	0.01%	0.005%	0.005%	21%	59%	23%
Young	1%	0.15%	0.25%	21%	59%	23%
MaxiMin	2.2%	0.15%	0.25%	21%	59%	23%

Table 2: Rounded maximal CLP and ACP probabilities which occurred during our simulations

REFERENCES

- [1] 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.
- [2] A. I. Barvinok. A polynomial time algorithm for counting integral points in polyhedra when the dimension is fixed. *Mathematics of Operations Research*, 19(4):769–779, 1994.
- [3] J. Bennett and S. Lanning. The Netflix prize. In *Proceedings of KDD Cup and Workshop*, 2007.
- [4] F. Brandt and H. G. Seedig. A tournament of order 24 with two disjoint TEQ-retentive sets. Technical report, <http://arxiv.org/abs/1302.5592>, 2013.
- [5] 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.
- [6] F. Brandt, F. Fischer, P. Harrenstein, and M. Mair. A computational analysis of the tournament equilibrium set. In *Proc. of 23rd AAAI Conference*, pages 38–43. AAAI Press, 2008.
- [7] F. Brandt, M. Chudnovsky, I. Kim, G. Liu, S. Norin, A. Scott, P. Seymour, and S. Thomassé. A counterexample to a conjecture of Schwartz. *Social Choice and Welfare*, 40:739–743, 2013.
- [8] F. Brandt, V. Conitzer, and U. Endriss. Computational social choice. In G. Weiß, editor, *Multiagent Systems*, chapter 6, pages 213–283. MIT Press, 2nd edition, 2013.
- [9] F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
- [10] 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.
- [11] 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>.
- [12] Y. Chevaleyre, U. Endriss, J. Lang, and N. Maudet. A short introduction to computational social choice. In *Proc. of 33rd Conference on Current Trends in Theory and Practice of Computer Science (SOFSEM)*, volume 4362 of *LNCS*, pages 51–69. Springer, 2007.
- [13] C. J. Clopper and E. S. Pearson. The use of confidence or fiducial limits illustrated in the case of the binomial. *Biometrika*, 26(4):404–413, 1934.
- [14] D. E. Critchlow, M. A. Fligner, and J. S. Verducci. Probability models on rankings. *Journal of Mathematical Psychology*, 35:294–318, 1991.
- [15] 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.
- [16] E. Ehrhart. Sur les polyèdres rationnels homothétiques à n dimensions. *Compte-Rendus de l'académie des Sciences*, 254(4):616, 1962.
- [17] P. C. Fishburn. Paradoxes of voting. *The American Political Science Review*, 68(2):537–546, 1974.
- [18] P. C. Fishburn. Condorcet social choice functions. *SIAM Journal on Applied Mathematics*, 33(3):469–489, 1977.
- [19] W. V. Gehrlein. *Condorcet's Paradox*. Springer, 2006.
- [20] W. V. Gehrlein and D. Lepelley. *Voting paradoxes and group coherence: the Condorcet efficiency of voting rules*. Springer-Verlag, 2010.
- [21] 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.
- [22] H. C. Huang and V. C. H. Chua. Analytical representation of probabilities under the IAC condition. *Social Choice and Welfare*, 17(1):143–155, 2000.
- [23] 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.
- [24] J. I. Marden. *Analyzing and Modeling Rank Data*. Number 64 in Monographs on Statistics and Applied Probability. Chapman & Hall, 1995.
- [25] 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.
- [26] N. Mattei, J. Forshee, and J. Goldsmith. An empirical study of voting rules and manipulation with large datasets. In *Proc. of 4th COMSOC*, 2012.
- [27] J. Oren, Y. Filmus, and C. Boutilier. Efficient vote elicitation under candidate uncertainty. In *Proc. of 23rd IJCAI*, pages 309–316. AAAI Press, 2015.
- [28] 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.
- [29] M. Regenwetter, B. Grofman, A. A. J. Marley, and I. M. Tsetlin. *Behavioral Social Choice: Probabilistic Models, Statistical Inference, and Applications*. Cambridge University Press, 2006.
- [30] J. Rothe, editor. *Economics and Computation: An Introduction to Algorithmic Game Theory, Computational Social Choice, and Fair Division*. Springer, 2015.
- [31] A. Schürmann. Exploiting polyhedral symmetries in social choice. *Social Choice and Welfare*, 40(4):1097–1110, 2013.
- [32] N. Tideman and F. Plassmann. Modeling the outcomes of vote-casting in actual elections. In D. S. Felsenthal and M. Machover, editors, *Electoral Systems: Paradoxes, Assumptions, and Procedures*, Studies in Choice and Welfare. Springer, 2012.
- [33] T. N. Tideman. Independence of clones as a criterion for voting rules. *Social Choice and Welfare*, 4(3):185–206, 1987.
- [34] M. C. Wilson and G. Pritchard. Probability calculations under the IAC hypothesis. *Mathematical Social Sciences*, 54(3):244–256, 2007.