université
PARIS-SACLAY

Probabilistic Approaches to the Analysis of Voting Outcomes' Variability

Approches probabilistes pour l'analyse de la variabilité des résultats électoraux

Thèse de doctorat de l'université Paris-Saclay

École doctorale n° 573, Interfaces : matériaux, systèmes, usages (INTERFACES)
Spécialité de doctorat: Informatique
Graduate School : Sciences de l'ingénierie et des systèmes. Référent : CentraleSupélec

Thèse préparée dans l'unité de recherche **Mathématiques et Informatique pour la Complexité des Systèmes** (Université Paris-Saclay, CentraleSupélec), sous la direction de **Vincent MOUSSEAU**, Professeur, et le co-encadrement de **Anaëlle WILCZYNSKI**, Maitre de conférence

Thèse soutenue à Paris-Saclay, le 12 décembre 2025, par

Henri SURUGUE

Composition du jury

Membres du jury avec voix délibérative

Vincent MERLIN Directeur de recherche CNRS, Université de Caen Normandie	Rapporteur
Lirong XIA Professeur, Rutgers University	Rapporteur
Umberto GRANDI Professeur, Université Toulouse Capitole	Examinateur
Jérôme LANG Directeur de recherche CNRS, Université Paris Dauphine - PSL	Examinateur
Maria POLUKAROV Professeure, King's College London	Examinatrice

THESE DE DOCTORAT

NNT : 2025UPAST181

ÉCOLE DOCTORALE
Interfaces:
matériaux, systèmes, usages

Titre: Approches probabilistes pour l'analyse de la variabilité des résultats électoraux

Mots clés: Choix social computationnel, Théorie du vote, Probabilités, Préférences, Vote stratégique.

Résumé: La théorie du choix social s'intéresse au problème de l'agrégation des préférences individuelles en une décision collective. Les règles de vote constituent l'un des principaux outils utilisés, mais la plupart d'entre elles peuvent poser des problèmes, et des résultats fondateurs dans ce domaine montrent qu'il n'existe pas de règle de vote parfaite. Plutôt que de chercher à identifier une règle de vote idéale, cette thèse adopte une perspective différente en se concentrant sur les résultats électoraux eux-mêmes.

A l'aide d'une approche probabiliste, elle étudie plus précisément comment ces résultats peuvent varier selon différents facteurs, fournissant ainsi des éléments concrets et simples sur leurs implications en pratique.

Une première source de variabilité réside dans le choix de la règle de vote. Puisqu'il en existe plusieurs, une question naturelle est de savoir si elles ont tendance à produire le même vainqueur. Pour y répondre, nous étudions la probabilité d'accord entre différentes règles. Les travaux précédents montrent que cette probabilité est faible sous des distributions uniformes de préférences. Cependant, nous nous concentrons sur des distributions plus structurées, couramment utilisées dans les études empiriques, et montrons que l'accord est en réalité nettement plus élevé dans de tels contextes.

Une deuxième source de variabilité est liée à la possibilité de comportements stratégiques. Étant donné que les règles de vote peuvent être manipulées par des électeurs qui ne voteraient pas sincèrement, nous analysons comment le vote stratégique affecte la diversité des résultats dans les élections avec la règle de pluralité. En particulier, nous examinons s'il est possible de prédire le vainqueur, à la fois d'un point de vue computationnel et qualitatif. Par exemple, nous étudions si le comportement stratégique augmente la probabilité que le vainqueur de Condorcet, celui qui bat tous les autres candidats en duel, soit élu.

Une troisième source de variabilité vient de la possibilité de manipulation de l'information par un acteur externe. L'existence de comportements stratégiques pose la question du pouvoir donné à ceux qui diffusent l'information. Nous nous demandons notamment si un institut de sondage peut influencer le résultat d'une élection avec la règle de la pluralité en manipulant les informations qu'il diffuse. Nous abordons cette question sous un angle à la fois computationnel et probabiliste, en étudiant si une telle manipulation est réalisable en pratique, et à quelle fréquence elle est réussie.

Enfin, une dernière source de variabilité étudiée provient de l'incertitude. En pratique, les informations recueillies par les sondages sont souvent incomplètes ou incertaines, ce qui pose la question de leur impact sur les décisions stratégiques des électeurs. Pour traiter ce problème, nous introduisons un nouveau modèle conçu pour prendre en compte ce type d'incertitude, donnant ainsi des outils pour étudier la variabilité qui en découle.

Chacune de ces sources de variabilité des résultats électoraux est donc analysée à travers une approche probabiliste, qui permet d'avoir un sens des proportions sur les résultats et d'apporter des éclairages concrets sur ce qui peut effectivement se produire en pratique.

Title: Probabilistic Approaches to the Analysis of Voting Outcomes' Variability
Keywords: Computational social choice, Voting theory, Probability, Preferences, Strategic voting.

Abstract: Social choice theory deals with the problem of aggregating individual preferences into a collective decision. Voting rules are one of the tools for this purpose, but most of them come with issues, and foundational results in the field of social choice show that no perfect voting rule exists. Rather than striving to identify an ideal rule, this thesis adopts a different perspective by focusing on election outcomes themselves.

Specifically, it investigates through a probabilistic approach, how these outcomes may vary depending on several factors, thus providing concrete and simple insights into their practical implications.

A first source of variability is the choice of voting rule. Since different voting rules exist, a natural question is whether they tend to produce the same outcome. To address this, we study the probability of agreement among various rules. Previous results in the literature show that this probability is low under uniform distributions of preferences. However, we focus on more structured distributions, commonly used in empirical studies, and show that agreement is in fact significantly higher in such settings.

A second source of variability is the possibility of strategic behavior. Since voting rules are susceptible to manipulation by insincere voters, we investigate how strategic voting impacts the diversity of outcomes in plurality elections. In particular, we examine whether it is possible to predict the winner, both from a computational point of view and from a qualitative perspective. For instance, whether strategic behavior increases the likelihood that the Condorcet winner, that is the candidate who beats every other candidate in pairwise comparisons, is elected.

A third source of variability is the potential manipulation of information by an external agent. The existence of strategic voting raises the question of how much power is given to those who disseminate information. We ask whether a polling institute can influence the outcome of plurality elections by manipulating the information it broadcasts. We approach this question from both a computational and a probabilistic perspective, examining whether such manipulation is computationally feasible and how frequently it succeeds.

A final source of variability considered comes from uncertainty. In practice, the information collected through polls is often neither complete nor certain, raising the question of how this affects voters' strategic decisions. To address this, we introduce a new model designed to handle such challenges, thereby providing tools to capture the variability arising from uncertainty.

Thus, each source of variability in election outcomes is studied using a probabilistic approach, which helps gain a sense of proportion regarding the results and provides concrete insights into what can happen in practice.

Acknowledgments

First of all, I would like to sincerely thank Lirong Xia and Vincent Merlin for accepting to review my thesis. I am very grateful for the time you took to read it and for your thoughtful and encouraging reports. I would also like to thank Maria Polukarov, Jérôme Lang, and Umberto Grandi for being part of my thesis committee. I have worked on many of your papers throughout my thesis, and it is a great pleasure and honor to receive your feedback and to have you all present at my PhD defense.

Then, I would like to thank the two main figures of my thesis, my PhD advisors: Anaëlle Wilczynski and Vincent Mousseau. Everything began with my first exchange with Vincent, who not only believed in this project from day one but has also been the most encouraging presence throughout. Like every great teacher, Vincent has the rare ability to always be supportive, not only with scientific guidance, but also with thoughtful and personal advice, which has been invaluable over these three years. Anaëlle has been closely involved at every stage of my thesis, providing careful and insightful feedback on my work. I am especially grateful for the sense of rigor and precision I have developed thanks to her work.

Moreover, I would like to thank my two co-authors outside my advisors' team: Magdaléna and Sébastien. Before thanking Sébastien, I especially want to thank Khaled, who had the great idea of suggesting that I visit him. It turned out to be a wonderful encounter, both scientifically and on a human level. Thank you, Sébastien, for your warm welcome in Compiègne and for the great moments we shared working together. I would also like to thank Magdaléna, not only for our research collaboration, but also for being a valuable source of relevant and encouraging feedback during the end of my thesis. I am truly grateful and happy to have the chance to work with you.

This thesis would not have been such a pleasant journey without the many colleagues who have been both a source of scientific inspiration and a constant motivation through the good times we shared. When I entered the lab on my very first day, I had the chance to be warmly welcomed by Manuel. I want to thank you here for the care and seriousness with which you take the task of welcoming newcomers, it has this wonderful effect of making people feel almost immediately comfortable and like they truly belong. You clearly didn't end up doing your PhD with Vincent by chance! However, he was not the only one who contributed to the welcoming environment I was lucky to be part of. I truly enjoyed spending many lunchtimes with colleagues like Pascale, Wassila, Marc, Romain, Francesco, Karl, Mathieu, Céline, Alexandre, Paolo, Myriam, Véronique, Sarah, Pierre, Inès, Lamia, Thomas, Gabriel, Magdaléna, Khaled, and many others... I also want to thank Fabienne, Suzanne and Leila for making all our research organization a lot easier.

Outside the lab, I would first like to thank Nicolas, who took his role in my doctoral follow-up committee more than seriously. He provided valuable insights at every stage of my PhD, including during the search for a plan after the thesis! I would also like to thank Anne for accepting to be part of this committee. Many other researchers outside the lab also deserve my thanks for the interesting discussions, including Jordan, Farès, Caren, Olivier, Théo, Richard, Ariana, François, Pierre, Ariane, Srdjan, Arthur, and many others...

Je continue en français pour les remerciements plus personnels, peut-être que cela donnera à certains davantage envie de les lire ! Je voudrais d'abord remercier les amis que j'ai la chance d'avoir à mes côtés. Ce n'est pas tous les jours qu'on a l'occasion de le faire, alors merci pour tous les bons moments passés ensemble et pour votre curiosité, que ce soit pour mon travail ou d'autres aspects de ma vie : Benjamin, Juliette, Estelle, Hugo, Elisabeth, Jocelyn, Eliot, Richard, Manuel ...

Dans cette dernière partie, je souhaiterais remercier ma famille, en commençant par mon frère Gauthier. Merci pour les discussions partagées autour de certains sujets de recherche et de beaucoup de questions d'enseignement. Je veux aussi remercier énormément ma sœur, qui, avec son intérêt tout modéré pour mon travail, me rappelle qu'il y a bien d'autres choses que la recherche dans la vie. J'ai beaucoup de chance de vous avoir tous les deux. Merci à mon père, qui est fier de moi chaque jour, pour à peu près tout, et qui exprime son affection parfois de manière discrète... mais que je sais maintenant mieux identifier. Enfin, merci à ma mère, dont le soutien est infaillible, et qui je le sais cherche toujours à faire le mieux pour moi.

Ensuite, je voudrais remercier toute ma famille au sens large : mes oncles et tantes, mes grands-parents, ainsi que mes cousins et cousines, qui ont contribué de près ou de loin au succès de ma thèse. Je tiens à remercier tout particulièrement Odile et Alain, qui m'ont accueilli de nombreuses fois lorsque je devais être tôt au labo, et qui portent toujours un vif intérêt à ce que je fais.

Enfin, je veux remercier celle qui m'accompagne au quotidien, elle se reconnaîtra. Tu es à la fois une source de joie immense et la personne que je souhaite le plus avoir à mes côtés dans tous les moments de ma vie, qu'ils soient importants ou anodins, heureux ou plus difficiles. Depuis le début de ma thèse, tu as été une présence joyeuse, curieuse, rassurante, et je mesure chaque jour la chance que j'ai de t'avoir à mes côtés.

Contents

1 - Introduction

Nineteen years ago, in a small French town, a group of schoolchildren was asked to elect their representatives for the local youth city council. Without hesitation, one of them raised his hand to run. Back home, he started reflecting seriously on how to win the election, given that four candidates were competing and each had to outline a political agenda. Should he propose installing football goals in the schoolyard to sway the electorate? His campaign worked and he was elected. Following two intense terms, he retired from a promising political career ... at the age of thirteen. Years later, he studied mathematics, computer science, and economics, fields that, perhaps surprisingly, intersect around a common theme: the study of collective decisions. Then, he wrote his first research paper on poll manipulation in political elections. Coincidence? Maybe not. Social choice theory lies at the heart of this intersection. It draws on tools from all three fields to formally study how individual preferences can be combined into a collective decision.

Social choice theory is interested in the problem of aggregating individual preferences. One well-known type of aggregation is *voting rules*, which are both a central topic in social choice theory and the focus of this work. Imagine an election with two candidates, A and B, and nine voters, where five voters prefer A over B, and the remaining four prefer B over A. In this case, it is natural to elect candidate A as the winner, as she receives the majority of the votes. Nevertheless, let us make this example slightly more complex by introducing a third candidate, C. Suppose that two voters who prefer A to B and two voters who prefer B to A now rank C first, while all remaining voters rank C in third position. The situation has already become less straightforward. Indeed, if we think of the classical *plurality* rule, in which each voter votes for her favorite candidate, and the candidate with the most votes wins, then candidate C would be elected. However, a majority of voters ranked C last. This example highlights a key limitation of the plurality rule: it ignores voters' full preference rankings and relies only on their first choices.

Coming back to our initial example with two candidates, A and B, one might also expect that the addition of a less popular third candidate, C, should not affect the relative ranking between the original two candidates. For instance, if two voters who preferred A to B now rank C first, while all other voters rank C last, then C is an irrelevant candidate. Nevertheless, under the plurality rule, candidate B ends up winning the election, receiving a higher plurality score than A. We say that plurality is violating the property called *independence of irrelevant alternatives*.

Imagine another election with three candidates, A, B, and C, and five voters. Two voters prefer candidate A over candidate B, and candidate B

over candidate C. Two other voters prefer candidate B over A, and A over C. Finally, one voter prefers candidate C over B, and B over A. In a plurality election, where each voter gives her vote to her most preferred candidate, both A and B receive two votes, while C receives only one. Ties are broken lexicographically; since A comes first in the tie-breaking order, A is declared the winner. However, the voter who prefers C the most has an incentive to change her vote and support candidate B instead because her first candidate C has no chance. By doing so, B would receive three votes and win the election, which is a more favorable outcome for that voter who ranks A last. Such a situation, where a voter can misreport her true preferences in order to obtain a more favorable outcome, is known as *strategic voting*. In this case, the voter is said to be *pivotal*. A voting rule that allows such behavior is said to be not *strategyproof*. Properties such as independence of irrelevant alternatives, strategyproofness, and others are referred to as *axioms*.

We also distinguish between two types of solutions to our problem of aggregating preferences: a *social welfare function*, which outputs a complete ranking of the candidates, and a *voting rule*, which selects one or more winners. For instance, the plurality rule assigns scores to candidates based on how many voters rank them first. These scores can be used to generate a complete ranking of the candidates, which is an example of a social welfare function. However, if the goal is simply to select one or more winners, one can apply a voting rule that picks the candidates ranked highest in the ordering. These examples show that the simplest rule, plurality, is far from perfect, as it violates the property of independence of irrelevant alternatives and allows for strategic voting. One might consider designing a more complex rule that would resolve these issues.

Nevertheless, social choice researchers have brought even additional bad news: for more than three candidates, the search of a social welfare function satisfying very reasonable axioms is hopeless. Indeed, combining the axiom of *independence of irrelevant alternatives* with weak Pareto condition, which essentially states that if every voter prefers A to B, then the social ranking must place A above B, leads to the conclusion that only a dictatorship can satisfy both axioms, this is Arrow's impossibility theorem [Arrow, 1950]. *Dictatorship* is defined as a social welfare function that always outputs the ranking of a single voter chosen in advance, and is of course undesirable. In addition, another fundamental impossibility result is the Gibbard and Satterthwaite theorem [Gibbard, 1973; Satterthwaite, 1975], which essentially states that, with more than three candidates, strategic voting is unavoidable.

These two results are central in social choice theory and make the problem of selecting a voting procedure far more complex than one might initially expect. From a broader perspective, a closely related question to that of selecting a voting rule is to understand how much the outcome can vary when the voting rule is changed. Indeed, the relevance of selecting a particular voting rule arises only if the outcome frequently depends on the chosen rule. In other words, having a voting rule that satisfies desirable properties may be of little importance if the outcome remains the same regardless of the rule used, since the final decision would not be affected. There are two natural approaches to addressing the question of agreement between voting rules, each grounded in the nature of the data used. First, one can analyze real-world data from various elections, such as American elections [Regenwetter et al., 2007], Romanian elections [Roescu, 2014], or the parliamentary elections in the Austrian federal state of Styria [Darmann et al., 2019]. Being grounded on actual data, it gives you answers to real cases but is susceptible to providing insights that apply only to the specific election under study. In contrast, a second approach relies on synthetic data and is complementary to the first, as it addresses the question of agreement between voting rules with greater generality and flexibility regarding election size. A drawback of this approach is its potential lack of empirical relevance. This approach began with the seminal work of Gehrlein and Fishburn [1980], which was limited to three candidates and a small set of voting rules. It was later extended to a broader range of voting rules [Merlin et al., 2000; Lepelley et al., 2000b; Gehrlein, 2003]. These papers, among others, have been surveyed in the work of Gehrlein and Lepelley [2010]. Thus, this topic is important to the field of social choice, as the choice of different voting rules can be viewed as a key source of variability in election outcomes.

The Gibbard and Satterthwaite impossibility theorem, previously discussed, also opens up a whole research question: if strategic behavior is unavoidable with more than three candidates, then we might be interested in modeling such strategic behavior. To address this question, a key issue lies in the information available to voters, as strategic decisions depend on what voters know about the current state of the election. Two main approaches have been proposed in the literature. The first, due to Myerson and Weber [1993], adopts a Bayesian game-theoretic framework in which all voters simultaneously assess the situation and decide whether to change their vote based on their utilities and probabilistic beliefs about other voters. In contrast, the second approach, introduced by Meir et al. [2010], models strategic behavior as an iterative process: at each step, a single voter updates her vote if she is pivotal, based on the current state of the election. In this framework, the information available evolves dynamically as voters react sequentially. Of course, both models are simplifications. The first relies on strong assumptions regarding utilities, probabilistic models, and the one shot nature of the decision-making process. The second assumes sequential information updates at each round, in a way similar to polls, which may be unrealistic in contexts where polls are sparse, and it assumes arbitrary strategic behavior from voters. Nevertheless, these two models have met consensus in the literature, as highlighted in the survey by Meir [2018].

Therefore, these models raise many questions, both about the interpretation of their results and about their underlying modeling framework. First, we may want to assess whether these strategic voting models often predict different winners, either through empirical [Thompson et al., 2013] or computational analysis [Rabinovich et al., 2015]. In other words, we aim to determine whether the variability in outcomes induced by strategic voting is substantial. Indeed, it would allow us to study the impact of strategic voting on the outcome of the election. To address this question, one may also turn to a more theoretical framework that studies the problem through probabilistic average-case analyses [Brânzei et al., 2013; Kavner and Xia, 2024]. Another relevant aspect is whether, despite the inevitability of strategic behavior, the final winner obtained after strategic voting may still satisfy other desirable properties. To this purpose, we present one other solution concept introduced by de Caritat marquis de Condorcet [1785], namely the *Condorcet winner*. Imagine an election with three candidates, A, B, and C, and seven voters. Three voters prefer candidate A over B, and B over C. Two voters prefer B to C, and C to A. Finally, two voters prefer C to B, and B to A. Then, candidate B defeats both A and C in pairwise comparisons, even though A is the plurality winner. Somehow, the underlying intuition behind this solution concept is that if we elect A, then after the election, a majority of the voters would have preferred B instead. In precise terms, a Condorcet winner is a candidate which beats every other candidate in pairwise comparisons. Such a candidate does not always exist because pairwise comparisons can cycle, referred to as a *Condorcet cycle*. However, when it does, one might expect it to be elected, but this is not always the case. For example, if we use the previous example and replace one voter who preferred C to B, and B to A with a voter who prefers C to A, and A to B, we obtain a Condorcet cycle: A is preferred to B, B to C, and C to A. A voting rule that always elects the Condorcet winner when it exists is called *Condorcet-consistent*. It is then natural to ask whether strategic voting increases or decreases the probability of electing such a winner [Grandi et al., 2013]. This question addresses the qualitative impact of strategic behavior, since being a Condorcet winner is a qualitative property that can apply to any candidate.

Moreover, as previously mentioned, the winner of strategic voting is strongly connected with the information available to voters, and may influence strategic voters depending on the nature of the information [Endriss et al., 2016; Reijngoud and Endriss, 2012]. This raises important concerns about the power held by those who disseminate such information. One may question whether polling institutes can influence the outcome of an election [Wilczynski, 2019; Baumeister et al., 2020]. Therefore, one may question the extent to which an external agent can influence the variability of election outcomes. For instance, one might ask whether a polling institute can affect the result of an election by manipulating the information it disseminates, thereby making the outcome itself susceptible to external influence.

Another concern regarding the information available to voters is its intrinsic uncertainty. Indeed, constructing polls necessarily involves sampling a subset of the electorate, as it is not feasible to query every voter. This sampling process introduces uncertainty into the information provided [Meir, 2017]. This observation raises the question of how voters make decisions under uncertainty, whether they follow a pessimistic criterion [Gilboa and Schmeidler, 1989], a combination of optimism and pessimism [Hurwicz, 1951], or some other behavior. This naturally leads to the question of how such decision-making behaviors may impact outcomes' variability. In addition, one may also raise the issue of uncertain voters [Kreiss and Augustin, 2020], as it is not always clear whether a voter is certain of her preference or able to provide a complete ranking over the candidates. Note that the majority of existing models rely on the strong assumption that each voter provides a complete and certain order over the candidates. Therefore, considering more general assumptions may lead to greater variability in election outcomes.

The goal of this thesis is to study the various sources of variability in election outcomes. This variability can stem from multiple sources: it may arise from the voting rule itself, from the susceptibility of all rules to strategic behaviors, from manipulation through polling information, or from incomplete or uncertain preference profiles. These problems share a common perspective: the aim is not to propose a novel "perfect" procedure for aggregating preferences, nor to characterize the normative properties of voting systems, but rather to focus on existing voting rules and to analyze their effects in concrete terms, by analyzing who is elected, and how this outcome may vary. Therefore, a relevant way to tackle this problem is through a probabilistic approach, as our concern is to measure the variability of the outcome, which is, by nature, a quantitative question. This approach contrasts with much of the literature in computational social choice, where the focus is often on addressing problems from an axiomatic or computational perspective [Brandt et al., 2016]. Nevertheless, we note that some literature on our questions does exist, and a non-exhaustive list of works using probabilistic approaches to our questions has already been mentioned, including the survey by Gehrlein and Lepelley [2010] on agreement between voting rules, the studies by Brânzei et al. [2013] or Kavner and Xia [2024] on iterative voting, and the work of Kreiss and Augustin [2020] on uncertain voters. When a theoretical issue arises in an election context, it is natural to ask whether it occurs frequently or only in rare cases, thereby allowing us to assess the magnitude and practical relevance of the problem. For instance, when considering the possibility of poll manipulation in plurality elections, we may not only ask whether such manipulation is possible or computationally easy to perform, but also whether it occurs frequently in practice, as this provides concrete insights of the phenomenon.

In Chapter 2, we give a presentation of the preliminary concepts, notations, and tools necessary for understanding this thesis. Indeed, we present the main concepts needed for the thesis, including the modeling of voters' preferences, an overview of voting rules, the basic strategic voting model, and the probabilistic tools for preference modeling.

In Chapter 3, we study how different voting rules may produce the same outcome. Although no voting rule is perfect from an axiomatic perspective, we show that the probability of agreement between them can be very high in certain contexts. In particular, we demonstrate that when the distribution of voters' preferences in the population satisfies some natural structural assumptions, this probability can exceed the levels previously computed in the literature.

In Chapter 4, we analyze how strategic voting can impact the outcomes of plurality elections. We adapt the concepts of *possible* and *necessary* winners in this context. A candidate is a possible winner if there exists a sequence of strategic moves that leads to her victory, and a necessary winner if every possible sequence of strategic moves results in her winning. We study these notions from both computational and probabilistic perspectives. We then assess the quality of the winner under strategic voting by examining whether strategic behavior increases or decreases the probability that the elected candidate is a Condorcet winner.

In Chapter 5, we explore the susceptibility of political elections to poll manipulation through strategic voting. Indeed, the possibility of strategic voting using poll information opens the question of the power granted to polling institutes. We will tackle that problem from a computational and probabilistic perspectives. Indeed, we are interested in two questions which are the ability to compute a manipulation and whether or not this strategy is often successful.

In Chapter 6, we extend the model of uncertainty in iterative voting introduced by Meir [2017] to introduce quantitative aspects. We use tools from imprecise probability to bridge the strategic voting models of Myerson and Weber [1993] and Meir et al. [2010], while also extending the framework to model new situations. In particular, we relax the assumption that voters' preferences are complete and certain.

2 - Preliminaries and Notations

This chapter is dedicated to the introduction of all basic notations that will be used in the whole thesis. Let us describe step by step all the objects that will help us describe an election. We obviously start with voters and candidates. Let $N = \{1, \ldots, n\}$ be a set of voters and $M = \{x_1, \ldots, x_m\}$ a set of m candidates. For notational convenience, we will use $[k]$ instead of $\{1, \ldots, k\}$ for any positive integer k, making therefore $N = [n]$.

2.1 . Preferences

We want to model voters' preferences regarding the candidates. Two main approaches are commonly used in the literature to represent agents' preferences: the ordinal approach, which models preferences as rankings over alternatives, and the cardinal approach, which assigns utility values to capture preference intensity. Following the former, we adopt the ordinal approach and represent preferences using linear orders. Furthermore, we will assume a linear order until Chapter 6, which stipulates that each voter can strictly compare any two candidates and rank all of them. Precisely, each voter $i \in N$ has preferences over candidates represented by a linear order $\succ_i$ over M. We may drop the i when the linear order is not associated to a specific voter. Let $top(\succ_i)$ be the preferred candidate of i, i.e.,

$$top(\succ_i) \succeq_i x, \quad \text{for every } x \in M.$$

We define analogously the worst preferred candidate of i, i.e.,

$$x \succeq_i worst(\succ_i), \quad \text{for every } x \in M.$$

Let $N^x := \{i \in N : top(\succ_i) = x\}$ be the set of voters who prefer x to any other candidate, and, for a given subset of voters $A \subseteq N$, let us define the set of voters who prefer x to y.

$$A^{x \succ y} := \{i \in A : x \succ_i y\}$$

However, one can argue that the hypothesis of complete linear order is not always verified since voters could have incomplete or uncertain preferences. Indeed, a voter may not be able to rank a candidate or may not be sure of her preferences. We will try to relax this hypothesis in Chapter 6.

Nevertheless, we assume it from Chapter 3 to Chapter 5 and summarize all preferences in one object called *preference profile* of the election that we denote $\mathscr{P} = (\succ_i)_{i \in N}$.

Here is a simple preference profile that illustrates the definition.

Example 1. *Consider an election with four voters and three candidates, with voters' preferences as follows:*

$$\begin{array}{ccccc} x_1 & \succ_1 & x_2 & \succ_1 & x_3 \\ x_1 & \succ_2 & x_3 & \succ_2 & x_2 \\ x_2 & \succ_3 & x_3 & \succ_3 & x_1 \\ x_3 & \succ_4 & x_2 & \succ_4 & x_1 \end{array}$$

In this example, voter 1 prefers x_1 to x_2 and x_2 to x_3.

We denote the set of all possible preference orders for m candidates by Π^m. Then, for a given preference order $\succ \in \Pi^m$, we define the rank of candidate x in $\succ$, denoted by $r_\succ(x)$, as

$$r_\succ(x) := |\{y \in M : y \succ x\}| + 1.$$

We also present a distance between two preferences called *Kendall tau distance*, called $dist_{KT}$, to evaluate the similarity between two preference orders $\succ$ and $\succ'$ in Π^m, by counting the number of pairwise comparisons on which the two orders disagree, i.e., $dist_{KT}(\succ, \succ') = |\{(x, y) \in M^2 : x \succ y$ and $y \succ x\}|$.

Continuing with the case $m = 3$ as in Example 1, we can explicitly enumerate all possible preference orders.

Example 2. *Consider an election with three candidates, we can enumerate all possible voters' preferences as follows:*

$$\begin{array}{ccccc} x_1 & \succ & x_2 & \succ & x_3 \\ x_1 & \succ & x_3 & \succ & x_2 \\ x_2 & \succ & x_1 & \succ & x_3 \\ x_2 & \succ & x_3 & \succ & x_1 \\ x_3 & \succ & x_1 & \succ & x_2 \\ x_3 & \succ & x_2 & \succ & x_1 \end{array}$$

In general, we can easily deduce that $|\Pi^m| = m!$ since each preference order corresponds to a permutation of the set of candidates M of size m.

2.2 . A Special Case of Preferences: Single-Peaked Preferences

In many situations, certain preferences may be excluded for various reasons, as they correspond to scenarios that are unlikely or even impossible. For example, consider an individual ranking temperatures ranging from $10°$C to $25°$C in 5-degree increments. A preference such as $25°\mathrm{C} \succ 20°\mathrm{C} \succ 15°\mathrm{C} \succ 10°\mathrm{C}$ indicates a consistent liking for higher temperatures. However, a preference like $25°\mathrm{C} \succ 10°\mathrm{C} \succ 15°\mathrm{C} \succ 20°\mathrm{C}$ appears incoherent, as it violates the natural ordering of temperature. In such cases, it is reasonable to restrict attention to *single-peaked* preferences, where each individual has a most-preferred option (a "peak") and preferences decrease monotonically as one moves away from this peak in either direction along a meaningful axis (e.g., temperature from low to high). The same reasoning applies in political settings, where preferences often align along a one-dimensional left-right ideological spectrum.

Building on the idea we roughly discussed, we now consider a common preference restriction, namely single-peakedness [Black, 1948].

Definition 1 (Single-peakedness [Black, 1948])**.** *A preference profile $\mathscr{P} \in (\Pi^m)^n$ is single-peaked if there exists an axis $>$ on M such that, for every voter $i \in N$, and each triple of candidates $x > y > z$, we have $y \succ_i x$ or $y \succ_i z$.*

All along the thesis, we consider, w.l.o.g., an axis $>$ on M such that $x_1 > \cdots > x_m$. Let $\Pi^m_>$ be the set of all possible single-peaked preference orders with respect to an axis $>$ on M.

Continuing with the case $m = 3$, we can enumerate all possible single-peaked preferences orders with respect to an axis.

Example 3. *Consider an election with three candidates and an axis $x_1 > x_2 > x_3$, we can enumerate all possible single-peaked voters' preferences with respect to this axis as follows:*

$$\begin{array}{ccccc} x_1 & \succ & x_2 & \succ & x_3 \\ x_2 & \succ & x_1 & \succ & x_3 \\ x_2 & \succ & x_3 & \succ & x_1 \\ x_3 & \succ & x_2 & \succ & x_1 \end{array}$$

In general, we can easily deduce that $|\Pi^m_>| = 2^{m-1}$, independently of the axis $>$.

Following the previous example, we give in Figure 2.1 an illustration of how single-peaked preferences can be represented with respect to a specific axis. We observe the presence of a single peak for each preference , which visually illustrates the concept of being single-peaked.

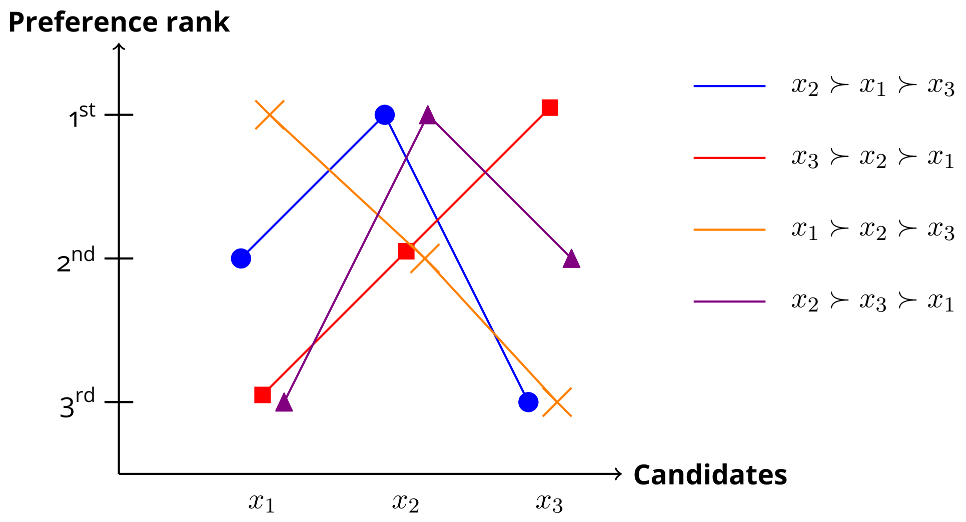

Figure 2.1: All possible single-peaked preference orders for three candidates with respect to the axis $x_1 > x_2 > x_3$

Since we have modeled voters' preferences, we are now ready to address one of the main concern of voting theory, namely the choice of a voting rule to aggregate voters' opinions.

2.3 . Voting Rules

A voting rule $\mathcal{F} : (\Pi^m)^n \to 2^M \setminus \{\emptyset\}$ selects a non-empty subset of candidates for each preference profile $\mathscr{P} \in (\Pi^m)^n$. For the rest of this thesis, most of the time, we reduce our attention to voting rules that choose a unique winner. When we have ties, the solution is to use a tie-breaking rule. We arbitrarily choose the *lexicographic tie-breaking*, denoted by $\rhd$, which is an order over the candidates.

We first present a class of voting rules, namely *scoring rules*.

Definition 2 (Scoring rules)**.** *A scoring rule $\mathcal{F}$ is associated with a score function $\alpha = (\alpha_1, \ldots, \alpha_m)$ and selects the candidates maximizing this score, i.e., $\mathcal{F}(\mathscr{P}) \in \arg\max_{x \in M} \sum_{i \in N} \alpha_{r_{\succ_i}(x)}$ for every preference profile $\mathscr{P} \in (\Pi^m)^n$.*

We denote by $\mathcal{W}_{\mathcal{F}}$ the set of winners under a given voting rule $\mathcal{F}$. This set is often reduced to a singleton, as we primarily focus on rules that elect a single winner. We then make a distinction between the ballot and the underlying preference. Let $b_i \in 2^M \setminus \{\emptyset\}$ denote the ballot submitted by voter i and $b \in M^n$ denote the *ballot profile*, i.e., $b := (b_1, \ldots, b_n)$. One of the simplest scoring rules we could think of is *plurality*, namely the rule where each voter votes for their preferred candidate, and the candidate with the highest number of votes is elected, we denote the winner set by $\mathcal{W}_P$.

Formally, the winner under Plurality of the ballot profile b is

$$\mathcal{W}_P(b) \in \arg\max_{x \in M} |\{i \in N : b_i = x\}|$$

For the particular case of plurality rule we add other notations: the ballot profile b from which b_i is excluded denoted by b_{-i} and $b^\top$ the *truthful* ballot profile where all voters submit their sincere preferences, i.e., $b_i^\top = top(\succ_i)$ for every voter $i \in N$.

Let us build an example to illustrate the plurality rule.

Example 4. *Consider an election with seven voters and three candidates, with voters' preferences as follows:*

$$\begin{array}{ccccc}
x_1 & \succ_1 & x_2 & \succ_1 & x_3 \\
x_1 & \succ_2 & x_2 & \succ_2 & x_3 \\
x_1 & \succ_3 & x_3 & \succ_3 & x_2 \\
x_2 & \succ_4 & x_3 & \succ_4 & x_1 \\
x_2 & \succ_5 & x_3 & \succ_5 & x_1 \\
x_3 & \succ_6 & x_2 & \succ_6 & x_1 \\
x_3 & \succ_7 & x_2 & \succ_7 & x_1
\end{array}$$

Therefore, in this example, x_1 receives 3 points, x_2 receives 2 points, and x_3 also receives 2 points. Hence, x_1 is elected under the plurality rule.

We then consider a specific type of scoring rules, namely *positional scoring rule* (PSR). A PSR $\mathcal{F}$ is defined by a *score vector* $\alpha = (\alpha_1, \ldots, \alpha_m)$ such that

$$\alpha_1 \geq \cdots \geq \alpha_m \quad \text{and} \quad \alpha_1 > \alpha_m.$$

Each candidate receives points from each voter based on the candidate's position in the voter's preference ranking: a candidate ranked j-th receives α_j points.

The k-*approval* voting rule, for $k \in [m-1]$, is a particular case of PSR where $\alpha_j = 1$ for all $j \in [k]$, and $\alpha_j = 0$ for all $k < j \leq m$. The plurality rule is the 1-approval rule and the *veto* rule is the $(m-1)$-approval rule. The Borda rule is the PSR characterized by an evenly spaced positional score vector, e.g., $\alpha = (m-1, m-2, \ldots, 1, 0)$.

Let us consider the same preference profile as in Example 4 and compare with the Borda rule.

Example 5. *Consider an election with seven voters and three candidates, with voters' preferences as follows:*

$$\begin{array}{ccccc}
x_1 & \succ_1 & x_2 & \succ_1 & x_3 \\
x_1 & \succ_2 & x_2 & \succ_2 & x_3 \\
x_1 & \succ_3 & x_3 & \succ_3 & x_2 \\
x_2 & \succ_4 & x_3 & \succ_4 & x_1 \\
x_2 & \succ_5 & x_3 & \succ_5 & x_1 \\
x_3 & \succ_6 & x_2 & \succ_6 & x_1 \\
x_3 & \succ_7 & x_2 & \succ_7 & x_1
\end{array}$$

In this example x_1 is elected using plurality. However, under the Borda rule, the candidate ranked first in a preference order receives 2 points, the second receives 1 point, and the third receives 0 points. Therefore, x_1 obtains 6 points, x_2 obtains 8 points, and x_3 obtains 7 points. Hence, x_2 is elected under the Borda rule.

All these voting rules appear natural, as they are based on intuitive ideas of assigning points to candidates. The simplest rule one can think of is plurality, where each voter assigns one point to a single candidate. Building on this idea, other scoring rules assign points to candidates with respect to their position in the voter's preference ranking, namely *positional scoring rule*. These rules have the particularity of evaluating each candidate individually, based solely on their position in the voters' preference orders, which is not always the case for other voting rules.

Another type of rule commonly used in practice is scoring rules with runoff. A typical example is *plurality with runoff*. In this rule, only the two candidates with the highest plurality scores in the first round proceed to the second round. In case of a tie, a predefined tie-breaking rule must be applied. Then, the preference profile is updated by removing all other candidates from each voter's ranking. The candidate who wins under plurality in this reduced profile is declared the overall winner. Building on that rule, we also have *instant-runoff*, in which one candidate with the lowest plurality score is eliminated in each round (up to $m-1$ rounds).

However, other solution concepts have been proposed, instead of evaluating the candidates on their absolute position in the voters' preferences, other voting rules take into account pairwise comparisons of candidates. For this purpose, we present the notion of Condorcet winner.

Definition 3 (Condorcet winner)**.** *A candidate x is the* Condorcet winner *if she beats all the other candidates in pairwise comparisons, i.e., $|N^{x \succ y}| > |N^{y \succ x}|$, for every candidate $y \in M \setminus \{x\}$. A weak Condorcet winner x is such that $|N^{x \succ y}| \geq |N^{y \succ x}|$, for every candidate $y \in M \setminus \{x\}$.*

We define symmetrically the (weak) Condorcet loser. In general, a (weak) Condorcet winner does not always exist. Otherwise, we obtain a Condorcet cycle, as illustrated in the following example.

Example 6. *Let us consider the following profile with 3 candidates and 3 voters:*

$$\begin{array}{ccccc} x_1 & \succ_1 & x_2 & \succ_1 & x_3 \\ x_2 & \succ_2 & x_3 & \succ_2 & x_1 \\ x_3 & \succ_3 & x_1 & \succ_3 & x_2 \end{array}$$

If we compute pairwise comparisons we have:

- *Comparing x_1 and x_2, we see that voters 1 and 3 prefer x_1, so $|N^{x_1 \succ x_2}| \geq |N^{x_2 \succ x_1}|$.*

- *Comparing x_2 and x_3, we see that voters 1 and 2 prefer x_2, so $|N^{x_2 \succ x_3}| \geq |N^{x_3 \succ x_2}|$.*

- *Comparing x_3 and x_1, we see that voters 2 and 3 prefer x_3, so $|N^{x_3 \succ x_1}| \geq |N^{x_1 \succ x_3}|$.*

Let us illustrate these computations:

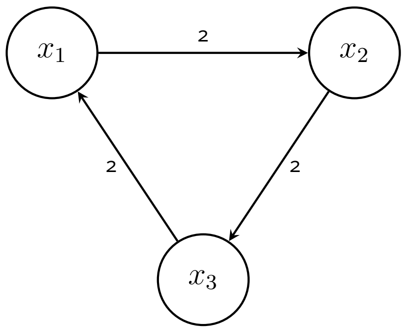

Figure 2.2: Majority graph of a Condorcet cycle

Therefore, we have a Condorcet cycle, since no candidate wins all pairwise comparisons.

However, a (weak) Condorcet winner always exists when the preferences are single-peaked and m is odd (when m is even) [Black, 1958]. A voting rule which always elects the Condorcet winner, when it exists, is called *Condorcet-consistent*. Note that PSRs are not Condorcet-consistent [de Caritat marquis de Condorcet, 1785].

Let us illustrate all these concepts.

Example 7. *Consider an election with seven voters and three candidates, with voters' preferences as follows:*

$$
\begin{array}{ccccc}
x_1 & \succ_1 & x_2 & \succ_1 & x_3 \\
x_1 & \succ_2 & x_2 & \succ_2 & x_3 \\
x_1 & \succ_3 & x_3 & \succ_3 & x_2 \\
x_2 & \succ_4 & x_3 & \succ_4 & x_1 \\
x_2 & \succ_5 & x_3 & \succ_5 & x_1 \\
x_3 & \succ_6 & x_2 & \succ_6 & x_1 \\
x_3 & \succ_7 & x_2 & \succ_7 & x_1
\end{array}
$$

In the pairwise comparisons, candidate x_2 defeats both x_1 and x_3, while x_3 only defeats x_1, and no candidate defeats x_2. Therefore, x_2 is the Condorcet winner. In the veto rule, candidate x_2 receives no last-place rankings and thus obtains the maximum possible score. In contrast, x_1 is ranked last by 4 voters and x_3 by 2 voters, receiving fewer points. Therefore, x_2 is the winner under the veto rule. In the first round, x_1 receives the most first-place votes and qualifies for the runoff *along with x_2. In the second round, x_2 is preferred over x_1 by a majority of voters (4 out of 7). Therefore, x_2 is the winner under the plurality with runoff rule. The results are summarized in the following table.*

Voting Rule	**Scores**			**Winner**
	x_1	x_2	x_3	
Plurality	3	2	2	x_1
Borda	12	14	11	x_2
Condorcet (pairwise wins)	0	2	1	x_2
Veto (2-approval)	3	7	5	x_2
Plurality with runoff (2nd round)	3	4	–	x_2

Table 2.1: Scores and winners under various voting rules for the profile in Example 7.

These different voting rules illustrate the diversity of ways to aggregate individual preferences. This naturally leads to the question of which rule to choose in practice.

2.4 . A Taste of Impossibility Theorems in Voting Theory

Some seminal works on voting theory by Arrow [1950] study the design of good voting rules with an axiomatic approach. We present the first impossibility theorem using the more general concept of *social welfare function* which is slightly more general than the one of voting rules. Indeed, a social welfare function takes a preference profile and returns a ranking of all candidates with respect to a criterion that should describe the best for the society as a whole, namely $\mathcal{SWF} : (\Pi^m)^n \to \Pi^m$, denoted $\succ_{\mathcal{SWF}}$. We can easily recover a voting rule (i.e. social choice function) with a unique winner by taking the first one of that ranking.

Let us now give some examples of axiomatic properties that may be desirable. First of all, we want the social welfare function to avoid *dictatorship*, meaning that the social ranking should not always coincide with the preferences of a single individual i fixed in advance, regardless of the preferences of the other voters. We might also want that a social welfare function prefers one candidate to another if every voter prefers this candidate. Thus, we say that a social welfare function is *weakly Paretian* if, for any two candidates $a, b \in M$, if for every i, $a \succ_i b$, then $a \succ_{\mathcal{SWF}} b$.

Another desirable axiomatic property could be the one of *independence of irrelevant alternatives*, which is the idea that if two candidates are ranked in a certain order, the addition of a new candidate should not affect their relative ranking. Specifically, the relative ranking $a \succ_{\mathcal{SWF}} b$ or $b \succ_{\mathcal{SWF}} a$ should not be affected by another candidate c.

We are now ready to present Arrow's impossibility theorem.

Theorem 1 (Arrow's impossibility theorem, 1950)**.** *When there are three or more candidates, every social welfare function that is weakly Paretian and independent of irrelevant alternatives must be a dictatorship.*

Many ideas have been brought into the field to try to circumvent this impossibility result, including relaxing some axioms, restricting preference domain (single-peaked profile for example) or using probabilistic tools (see Brandt et al. [2016] for a survey).

A second classical result that we will need to contextualize our work is the impossibility theorem of Gibbard and Satterthwaite [Gibbard, 1973; Satterthwaite, 1975], basically saying that no voting rule can avoid strategic voting when there are more than three candidates. To introduce it, we will need to make some adaptations and introduce new concepts.

The adaptation of the notion of dictatorship for a voting rule is that the voting rule should not always coincide with the top preference of a single individual i fixed in advance, regardless of the preferences of the other voters. A quite natural and desirable property for a voting rule is *non-imposition*, which means that any candidate can be elected in at least one possible profile of individual preferences, i.e., for all candidate c, there exists a preference profile $\mathscr{P}$ where the voting rule elects c. Additionally, one might wish for a voting rule to be *strategyproof*, implying that no voter has an incentive to misreport her preference to obtain a better outcome. Indeed, a voting rule $\mathcal{F}$ is strategyproof for a voter i if for all profile $\mathscr{P}$ there is no other profile $\mathscr{P}'$ where i has changed its vote such that $\mathcal{F}(\mathscr{P}') \succ_i \mathcal{F}(\mathscr{P})$. A voting rule is said to be strategyproof if it is strategyproof for all voters. Finally, we say that a voting rule is resolute if it always elects one candidate.

Theorem 2 (Gibbard-Satterthwaite's impossibility theorem, 1973; 1975)**.** *When there are three or more candidates, every voting rule that is resolute, nonimposed and strategyproof must be a dictatorship.*

This theorem has inspired numerous works in computational social choice, either aiming to limit voters' ability to manipulate outcomes through algorithmic complexity (see Chapter 6 of Brandt et al. [2016] for a survey) or allow manipulation and trying to model strategic behavior (see [Meir, 2022] for a survey).

2.5 . Strategic Voting

This line of research attempts to circumvent the impossibility result stated in Theorem 2 by allowing for voter manipulation and analyzing strategic voting through a game-theoretical approach. Indeed, since voters may act strategically, a ballot does not necessarily reflect the voter's true preferences. For instance, under the plurality rule, the chosen candidate might not be the voter's top ranked option, as she may have voted strategically to improve the election outcome from her perspective. We will examine this phenomenon as an iterative model called *iterative voting*.

In iterative voting, the goal is to study strategic behavior from a game-theoretical and discrete perspective. At each step, one voter is allowed to change her ballot. The process reaches a *Nash equilibrium* when no voter wishes to update her ballot given the available information. However, convergence is not guaranteed in general, cycles in strategic moves may prevent the process from terminating. In this thesis, we focus on one specific voting rule, plurality, due to its simplicity, naturalness, and widespread use. Its structural properties also make the analysis of strategic behavior more tractable.

Therefore, we consider plurality with lexicographic tie-breaking $\triangleright$ to be the voting rule.

Recall that the winner under plurality of the ballot profile b is $\mathcal{W}_P(b) \in \arg\max_{x \in M} s_x(b)$, where $s_x(b) := |\{i \in N : b_i = x\}|$ and a lexicographic tie-breaking, denoted by $\triangleright$, is used if necessary. By abuse of notation, we sometimes write s_x instead of $s_x(b)$. Let I_n^m be the set of all possible candidates' scores under plurality, i.e., $I_n^m := \{s \in \mathbb{N}^m \mid \sum_{j=1}^m s_j = n\}$. By abuse of notation, we sometimes write $\mathcal{W}_P(s)$ to refer to the winner of a score vector s. Let $s^\top$ denote the candidates' scores in $b^\top$. We will present the classical iterative voting model introduced by Meir et al. [2010], along with its subsequent developments as reviewed in the survey by Meir [2022], and generalized by Wilczynski [2019].

Initially, all voters vote truthfully, therefore the initial ballot profile b^0 is exactly the truthful ballot profile $b^\top$. Then they change their ballot strategically following a *best response* strategy which consists in supporting their preferred candidate within the set of so-called *potential winners*. A candidate y is a potential winner for voter i, at a given step where the current score vector is s, if i believes that voting for y will make candidate y the new winner, i.e., $s^{-i}_{\mathcal{W}_P(s^{-i})} - s^{-i}_y + \mathbb{1}_{\mathcal{W}_P(s^{-i}) \triangleright y} \leqslant 1$, where s^{-i} denotes the score vector s without counting the current ballot b_i of voter i. Let PW_i^t denote the set of potential winners for voter i at step t, and PW^t the set of all potential winners at step t, i.e., $PW^t := \bigcup_{i \in N} PW_i^t$.

Example 8. *Figure 2.3 shows a situation in which voter i, whose preference order is $x_3 \succ_i x_1 \succ_i x_2$, has two potential winners x_1 and x_2, that is, $PW_i = \{x_1, x_2\}$.*

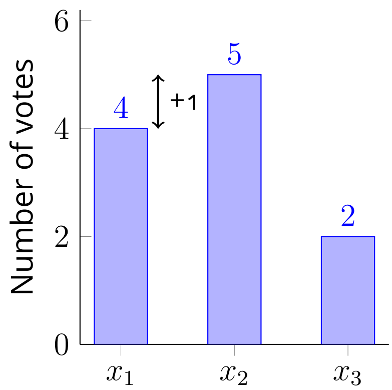

Figure 2.3: Pivotal situation for a voter $x_3 \succ x_1 \succ x_2$

When only a score vector is mentioned without a reference to a specific time step t, we may directly write $PW(s)$ to denote the set of potential winners according to a given score vector s. Note that this definition is independent of the preference profile.

At every step, one voter may change her vote to a best response according to the following behavior. For each voter i at step t, where the current winner is denoted by w^{t-1}: i deviates from her current ballot b_i^{t-1} to another ballot b_i^t supporting candidate $y \in PW_i^{t-1} \setminus \{w^{t-1}\}$ if y is her most preferred candidate within PW_i^{t-1}. Note that this behavior corresponds to "direct" best response [Meir et al., 2010]. For example, in Figure 2.3, a move from x_3 to x_1 is a direct best response.

Example 9. *Consider an election with five voters and four candidates, with voters' preferences as follows:*

$$\begin{array}{ccccccc}
x_1 & \succ_1 & x_3 & \succ_1 & x_4 & \succ_1 & x_2 \\
x_2 & \succ_2 & x_1 & \succ_2 & x_3 & \succ_2 & x_4 \\
x_3 & \succ_3 & x_2 & \succ_3 & x_1 & \succ_3 & x_4 \\
x_4 & \succ_4 & x_2 & \succ_4 & x_1 & \succ_4 & x_3 \\
x_4 & \succ_5 & x_3 & \succ_5 & x_1 & \succ_5 & x_2
\end{array}$$

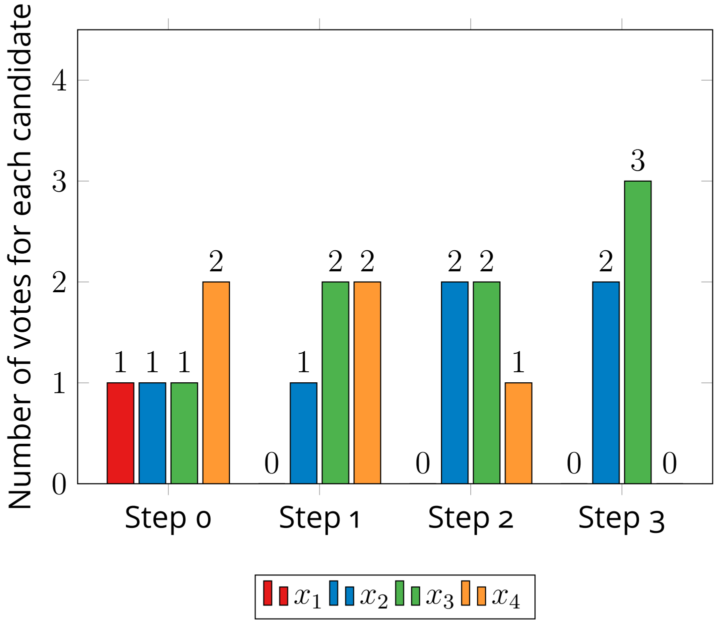

Figure 2.4: Number of votes per candidate over the steps of iterative voting

When needed, a lexicographic tie-breaking rule is used. Initially, in the truthful preference profile, x_4 is the winner. Voter 1 has an incentive to change her vote from x_1 to x_3, as this would make x_3 the new winner, which she prefers to x_4. Then, voter 4 may deviate from x_4 to x_2 in order to elect x_2 instead of x_3, which is a better outcome for her. Finally, voter 5 can change her vote from x_4 to x_3 to ensure that x_3 is elected rather than x_2, which she prefers. At this point, no voter has an incentive to deviate further, and x_3 is thus the final winner.

This example gives us a glimpse of how this model works, but it says nothing about the impact of the order in which strategic moves are made. The question of describing the diversity and qualitative characteristics of the winner will be addressed in Chapter 4. However, one other concern with this model is that polls are not fully reliable. This unreliability arises for several reasons. For instance, polling institutes only survey a subset of voters, and some respondents may be unsure of their views at the time of the poll or may intentionally misrepresent their preferences. These questions will be addressed respectively in Chapter 6 and Chapter 5.

We now turn to a final prerequisite for reading this thesis: the probabilistic tools used to study the variability of voting outcomes. As already motivated in the introduction, we adopt a probabilistic perspective to provide quantitative insights into the voting problems under consideration.

2.6 . Probabilistic Modeling of Preferences

Instead of analyzing voting rules purely from an axiomatic angle, we could take a quantitative perspective. Although the axiomatic approach offers valuable insights into the qualitative properties that voting rules satisfy, it does not provide any information about how frequently these properties hold when they are not always satisfied. More generally, many papers investigate voting problems as decision problems, addressing interesting computational questions and sometimes complementing the analysis with experiments. Empirical studies based on computer simulations can be a valuable complement to theoretical results in various voting problems, such as manipulation, winner determination, bribery, or the analysis of possible and necessary winners [Brandt et al., 2016]. Ideally, one would rely on real-world data to quantify such phenomena, for instance by using the Preflib platform [Mattei and Walsh, 2013]. However, real-world data are often limited in scope and highly context dependent, which may hinder the generalizability of experimental findings. In contrast, synthetic data enable the simulation of elections of arbitrary size and offer full control over the experimental parameters.

However, for experiments to be meaningful, we also need to simulate realistic elections, raising the question of a compromise between realism and flexibility. A large number of statistical distributions exist for generating elections [Boehmer et al., 2024]. We present these objects, which will be key to our analysis, and will study them from both a practical and a theoretical perspective. Let us denote as $C(n, \Pi^m)$ the probability distribution of drawing n preference orders from Π^m to constitute a preference profile $\mathscr{P} \in (\Pi^m)^n$. Such a probability distribution $C(n, \Pi^m)$ is called a *culture*. Let us first consider independent and identical drawings of voters' preferences such that we can either look at the distribution $C(n, \Pi^m)$ as a whole object or n drawings of preferences $\succ_i$. We denote by $\mathbb{P}_C$ the associated probability distribution. We also denote by $C(n, \Pi^m_{sub})$ a probability distribution over a subdomain $\Pi^m_{sub} \subseteq \Pi^m$ of the set of preference profiles. We begin by presenting the simplest possible culture one can consider in the absence of any prior information, namely the *impartial culture*.

Definition 4 (Impartial culture)**.** *The impartial culture, called IC, draws every preference order $\succ$ from Π^m with uniform probability, i.e., $\mathbb{P}_{IC}(\succ) = \frac{1}{m!}$.*

Let us give a short example to illustrate impartial culture.

Example 10. *Consider an election of three candidates, we might consider Π^3 and give equal weight to all orders, i.e., for all $\succ$ from Π^3 , $\mathbb{P}_{IC}(\succ) = \frac{1}{6}$.*

Another point of view to choose a distribution of preferences, without any additional information, is the *impartial anonymous culture*, which assumes a uniform distribution over preference profiles.

Definition 5 (Impartial anonymous culture)**.** *The impartial anonymous culture, called IAC, draws every profile $\mathscr{P}$ from $(\Pi^m)^n$ with uniform probability.*

Let us give a short example to illustrate impartial anonymous culture.

Example 11. *Consider an election of three candidates and four voters, we might consider $(\Pi^3)^4$ and give equal weight to all preference profiles, i.e., for all $\mathscr{P}$ from $(\Pi^3)^4$, $\mathbb{P}_{IAC}(\mathscr{P}) = \frac{1}{\binom{9}{5}} = \frac{1}{126}$.*

Therefore, this provides an idea of the quantitative differences. Mathematically, under IC the voters' preferences are independent, whereas under IAC the entire profile is drawn at once, making the voters' preferences dependent. From a broader perspective, the underlying idea is quite different: IAC draws a single profile from the set of all possible profiles, whereas IC draws each voter's preference independently. In particular, this means that under the IAC model, all plurality score vectors are drawn with equal probability, whereas the IC model tends to assign higher probability to more balanced scores.

These two cultures are very important to study because they do not favor any candidate by construction. However, if we wish to consider a structure in the voters' preferences, such as a single-peaked axis representing a left-right political spectrum, then impartial cultures are far from capturing such a setting. The following remark provides further insight into that idea.

Remark 3. *Consider an election of three candidates, we might suppose that a left-right axis $>$ structure preferences in the sense of Definition 1 (e.g., political context). Without lost of generality, we can suppose that the axis is $x_1 > x_2 > x_3$. Then, as we have already seen in Example 3, $x_1 \succ x_3 \succ x_2$ and $x_3 \succ x_1 \succ x_2$ are not single-peaked with respect to $>$. The question then becomes how to draw such preferences in this context. Two approaches can address this issue: either exclude these preferences altogether (assign them zero probability) or include them but with a lower probability.*

Following the previous remark, we define *single-peaked cultures* as distributions that only generate single-peaked orders with respect to a given axis.

Definition 6 (Single-peaked culture)**.** *For a given axis $>$ over M, a culture $C(n, \Pi^m_{sub})$ is said to be single-peaked if $C(n, \Pi^m_{sub}) = C(n, \Pi^m_{>})$.*

We might sample single-peaked preference orders by drawing them uniformly on the restricted space of single-peaked preference orders $\Pi^m_{>}$. The associated culture then refers to Walsh's model [Walsh, 2015]. Another way to impartially sample single-peaked preference orders is to use the model of Conitzer [Conitzer, 2007], which generates preferences in $\Pi^m_{>}$ by first choosing a peak uniformly at random among the candidates, and then iteratively selecting the next candidate uniformly from those immediately to the left or right of the last chosen one on the axis, thus ensuring that each candidate is equally likely to be ranked first under plurality.

Definition 7 (Walsh's distribution)**.** *The* Walsh's distribution $\pi_W : \Pi^m_> \to [0,1]$ *is such that* $\pi_W(\succ_i) = \frac{1}{2^{m-1}}$*, for every* $\succ_i \in \Pi^m_>$*.*

Here is a small example to illustrate Definition 7.

Example 12. *Consider an election of three candidates, if we suppose an axis* $x_1 > x_2 > x_3$*. Walsh's distribution is given by:*

Order $\succ$	*Probability*
$(x_1 \succ x_2 \succ x_3)$	0.25
$(x_2 \succ x_1 \succ x_3)$	0.25
$(x_2 \succ x_3 \succ x_1)$	0.25
$(x_3 \succ x_2 \succ x_1)$	0.25

Definition 8 (Conitzer's distribution)**.** *The* Conitzer's distribution $\pi_C : \Pi^m_> \to [0,1]$ *is such that* $\pi_C(\succ_i) = \frac{1}{m} \cdot \frac{1}{2^{\min\{r_{\succ_i}(x_1), r_{\succ_i}(x_m)\}-1}}$ *for every* $\succ_i \in \Pi^m_>$*.*

This definition is equivalent to the algorithm proposed by Conitzer [2007]. The peak is selected uniformly at random, corresponding for the $\frac{1}{m}$ term. Once the peak is fixed, the next candidate is chosen uniformly among the two candidates adjacent on the axis $>$, making the process dependent on the relative positions of the two extreme candidates. Specifically, once one of these two extreme candidates is selected, the rest of the ranking is completed by successively adding the remaining candidates on the same side with respect to the axis $>$. Here is a small example to illustrate Definition 8.

Example 13. *Consider an election of three candidates, if we suppose an axis* $x_1 > x_2 > x_3$*. Walsh's distribution is given by:*

Order $\succ_i$	*Probability*
$(x_1 \succ x_2 \succ x_3)$	$1/3$
$(x_2 \succ x_1 \succ x_3)$	$1/6$
$(x_2 \succ x_3 \succ x_1)$	$1/6$
$(x_3 \succ x_2 \succ x_1)$	$1/3$

These two approaches will be expanded and studied in Chapter 3. We can also mention another culture which is also a restriction on a single-peaked axis but with impartial anonymous culture [Saari and Valognes, 1999].

On a different note, one might consider alternative structures, such as *Euclidean distributions* [Boehmer et al., 2024]. This distribution relies on a different structure than the single-peaked domain, namely the Euclidean domain. Candidates and voters are placed as points in a d-dimensional space, where each voter ranks candidates according to their proximity (with respect to a chosen distance), preferring closer candidates over those farther away. For example, in two dimensions, this corresponds to placing the candidates in a two-dimensional plane. In the context of political elections, the two axes could represent cultural liberalism and economic liberalism. From a more technical perspective, an important distinction is that the Euclidean model accounts for distances between candidates in a quantitative way, unlike the single-peaked domain which is purely ordinal.

Next, we detail on the second point in Remark 3 by presenting the *Mallows culture*, where the probability of a preference order decreases with its distance from a fixed reference order.

Definition 9 (Mallows culture, 1957). *For given $\sigma \in \Pi^m$ and $\phi \in [0,1]$, the Mallows culture, called $\mathcal{M}^{\phi,\sigma}$, draws every preference order with a probability related to its distance to the reference ranking σ, more precisely, $\mathbb{P}_{\mathcal{M}^{\phi,\sigma}}(\succ_i) = \frac{1}{Z}\phi^{dist_{KT}(\succ_i,\sigma)}$ where $Z = \sum_{\succ_i \in \Pi^m} \phi^{dist_{KT}(\succ_i,\sigma)}$.*

Note that culture $\mathcal{M}^{1,\sigma}$ corresponds to the impartial culture.

Let us give a toy example for a three candidates election.

Example 14. *Consider an election of three candidates, we might consider $\sigma = x_1 \succ x_2 \succ x_3$. The Mallows distribution with parameter $\phi = 0.5$ and the Kendall-Tau distance is given by:*

$$\mathbb{P}_{\mathcal{M}^{0.5,\sigma}}(\succ_i) = \frac{1}{Z}(0.5)^{dist_{KT}(\succ,\sigma)},$$

Order $\succ$	$dist_{KT}(\succ,\sigma)$	$(0.5)^{dist_{KT}(\succ,\sigma)}$
$(x_1 \succ x_2 \succ x_3)$	0	1
$(x_1 \succ x_3 \succ x_2)$	1	0.5
$(x_2 \succ x_1 \succ x_3)$	1	0.5
$(x_2 \succ x_3 \succ x_1)$	2	0.25
$(x_3 \succ x_1 \succ x_2)$	2	0.25
$(x_3 \succ x_2 \succ x_1)$	3	0.125

For example, $dist_{KT}((x_1 \succ x_2 \succ x_3),(x_1 \succ x_3 \succ x_2)) = 1$, since one can obtain the latter by interchanging x_2 and x_3 in the former. We now compute the re-normalization constant:

$$Z = 1 + 0.5 + 0.5 + 0.25 + 0.25 + 0.125 = 2.625$$

With the following computation

$$\mathbb{P}_{\mathcal{M}^{0.5,\sigma}}(\succ) = \frac{dist_{KT}(\succ,\sigma)}{2.625}$$

we are able to conclude our toy example:

Order $\succ$	*Probability*
$(x_1 \succ x_2 \succ x_3)$	$\frac{1}{2.625} \approx 0.381$
$(x_1 \succ x_3 \succ x_2)$	$\frac{0.5}{2.625} \approx 0.190$
$(x_2 \succ x_1 \succ x_3)$	$\frac{0.5}{2.625} \approx 0.190$
$(x_2 \succ x_3 \succ x_1)$	$\frac{0.25}{2.625} \approx 0.095$
$(x_3 \succ x_1 \succ x_2)$	$\frac{0.25}{2.625} \approx 0.095$
$(x_3 \succ x_2 \succ x_1)$	$\frac{0.125}{2.625} \approx 0.048$

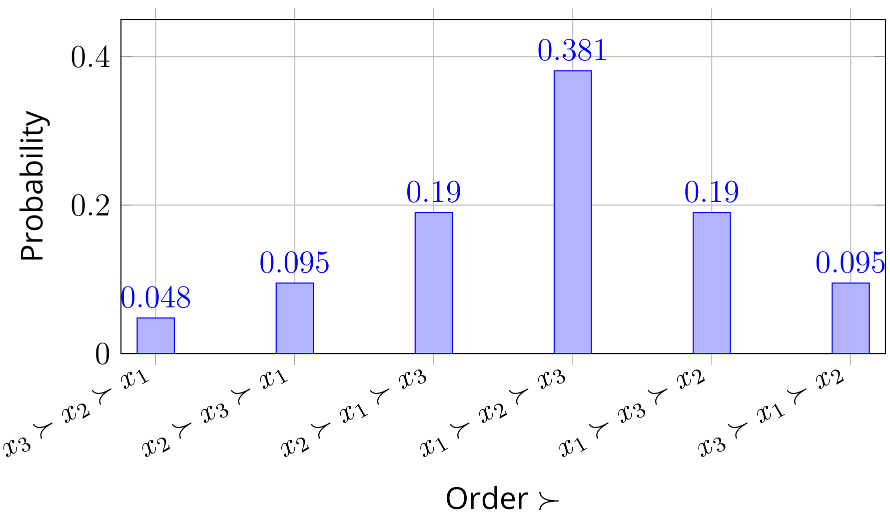

Figure 2.5: Mallows distribution of preferences with three candidates with $\phi = 0.5$ and reference order $\sigma = x_1 \succ x_2 \succ x_3$

This model presented in Figure 2.5 is frequently used in experiments [Boehmer et al., 2024], being the second most popular after impartial culture, as it appears in 28.5% of the papers considered in this survey of numerical experiments. We also want to present a last model which is a bit different since it does not suppose the independence of voters' drawings.

Definition 10 (Pólya-Eggenberger culture, 1923)**.** *The* Pólya-Eggenberger urn model, *P-E*(r) *where* $r \in \mathbb{R}_+^*$ *and* $R = m! \cdot r$, *draws a preference profile* $\mathscr{P}$ *in* $(\Pi^m)^n$ *as follows:*

- *We consider an urn initially containing* $m!$ *balls representing the* $m!$ *different preference orders from* Π^m*, i.e., each* ℓ^{th} *preference order from* Π^m *is initially drawn with probability* $\beta_\ell = \frac{1}{m!}$
- *To generate our preference profile* $\mathscr{P}$ *with* n *voters, we proceed as follows: for each voter, we draw a ball from the urn and assign the corresponding preference order to the voter. After each draw, the ball is returned to the urn along with* R *additional copies, where* $R > 0$*.*

For example, taking $r = \frac{1}{9}$ gives the well-known UM10 model which is 10%-correlation urn. Note that culture P-E($\frac{1}{m!}$) corresponds to the impartial anonymous culture.

Of course, many other ways to simulate preference data exist [Boehmer et al., 2024]. However, we only present those that are both common and useful for the remainder of the thesis.

We end up this section of preliminaries with some probabilistic tools that will be useful.

2.7 . Probabilistic Tools

This section is devoted to presenting some probabilistic results that will be needed for the remainder of the thesis. We begin with two lemmas from Hoeffding [1994], which will be useful to derive bounds on the speed of convergence toward a limiting behavior. Indeed, while some results are asymptotic, typically with respect to the number of voters, it is often desirable to assess whether these results still hold in small elections of reasonable size.

Lemma 4 ([Hoeffding, 1994])**.** *Let X_k be some independent real random variables, and $(a_k)_{k\in[n]}$ and $(b_k)_{k\in[n]}$ two real sequences such that for every $k \in [n]$, we have $a_k < b_k$ and $\mathbb{P}(a_k \leqslant X_k \leqslant b_k) = 1$. Then, for every $t > 0$, $\mathbb{P}(S_n - \mathbb{E}(S_n) \geqslant t) \leqslant e^{\frac{-2t^2}{\sum_{k=1}^n (b_k - a_k)^2}}$, where $S_n = \sum_{k=1}^n X_k$.*

Lemma 5 ([Hoeffding, 1994])**.** *Let X_k be some independent real random variables, and $(a_k)_{k\in[n]}$ and $(b_k)_{k\in[n]}$ two real sequences such that for every $k \in [n]$, we have $a_k < b_k$ and $\mathbb{P}(a_k \leqslant X_k \leqslant b_k) = 1$. Then, for every $t > 0$, $\mathbb{P}(S_n - \mathbb{E}(S_n) \leqslant -t) \leqslant e^{\frac{-2t^2}{\sum_{k=1}^n (b_k - a_k)^2}}$, where $S_n = \sum_{k=1}^n X_k$.*

These lemmas are particularly useful when studying asymptotic phenomena, typically with respect to the number of voters in a voting situation, and when we wish to determine how quickly such phenomena arise as the number of voters increases. They allow us in particular to derive bounds on the probability that a given event occurs.

Another result that will be useful, even though it is quite simple, is the following:

Lemma 6 ([Bonferroni, 1936])**.** $\mathbb{P}(\bigcap_{i=1}^n A_i) \geqslant \sum_{i=1}^n \mathbb{P}(A_i) - (n-1)$.

We end this section by another probabilistic result, namely the Glivenko-Cantelli theorem [Cantelli, 1935]. In our voting situations, this theorem allows us to verify that a phenomenon indeed occurs asymptotically with respect to the number of voters. If a property holds for the limit profile of a probability distribution over preferences, then any sequence of empirical profiles sampled independent and identical from that distribution will asymptotically satisfy the property almost surely.

Theorem 7 (Glivenko–Cantelli, 1935)**.** *Let $(\Omega, \mathcal{A}, \mathbb{P})$ be a probability space. Let $(X_i)_{1\leq i\leq n}$ be independent and identically distributed random variables with common cumulative distribution function F. For $\omega \in \Omega$, we denote by $F_n(\cdot, \omega)$ the empirical distribution function of the sample $(X_i)_{1\leq i\leq n}(\omega)$. Almost surely, the empirical distribution function F_n converges uniformly to the distribution function F, in other words,*

$$\mathbb{P}_\omega \left(\lim_{n\to\infty} \|F_n(\cdot, \omega) - F\|_\infty = 0 \right) = 1.$$

We now have all the necessary tools to begin our study of voting outcomes' variability.

3 - Agreement Among Voting Rules on the Same Outcome

Abstract

Many different voting rules have been proposed in the literature and they can select very different alternatives. This naturally raises the question of whether this diversity in outcomes often occurs. Our goal in this chapter is to offer a probabilistic perspective on this question. Previous works have shown that the probability of agreement between voting rules is generally quite low under impartial culture. We use a similar probabilistic approach on single-peaked cultures. We describe voting rules to agree under standard single-peaked cultures, and show that the probability of agreement between rather large families of voting rules is much higher under such cultures, with fast convergence of this probability with respect to the number of voters. We finally provide some insights on other structured preference distributions, observing that many exhibit similar convergence in agreement, including the Mallows' distribution. Our study reveals a tendency of several well-known voting cultures to bias the outcome of voting rules, which is worth knowing before conducting experiments on synthetic data.

Résumé

De nombreuses règles de vote ont été proposées dans la littérature et elles peuvent sélectionner des alternatives très différentes. La question se pose alors de savoir si cette diversité de gagnants se produit souvent en pratique. Ce chapitre est donc consacré à la question de l'accord entre les règles de vote. Des travaux antérieurs ont montré que la probabilité d'accord entre règles de vote est généralement assez faible sous des cultures de préférences uniformes. Nous utilisons une approche probabiliste similaire sur des cultures "single-peaked". Nous décrivons les règles de vote s'accordent sous ces cultures et nous montrons que la probabilité d'accord d'une grande famille de règles de vote est beaucoup plus élevée dans ce cadre, avec une convergence rapide en fonction du nombre de votants. Enfin, nous donnons un aperçu d'autres distributions de préférences structurées, en observant que nombre d'entre elles présentent une convergence similaire, y compris la distribution de Mallows. Notre étude révèle que plusieurs cultures de vote connues ont tendance à biaiser le résultat des règles de vote, ce qui est important à savoir avant de mener des expériences sur des données synthétiques.

Most of the content of this chapter is based on a paper co-authored with Vincent Mousseau and Anaëlle Wilczynski, which was accepted at the 28^{th} European Conference on Artificial Intelligence (ECAI 2025) [Mousseau et al., 2025b].

3.1 . Introduction

Our investigation into the variability of voting outcomes begins with this first chapter, which focuses on how outcomes change under different voting rules. This is a central issue in voting theory where much attention is given to designing effective voting systems. However, as we recall in Chapter 2 the social choice literature is famous for impossibility theorems, e.g., Arrow theorem [Arrow, 1950] or Gibbard-Satterthwaite theorem [Gibbard, 1973; Satterthwaite, 1975] basically stating that no perfect voting rule exists. Many different voting rules have been designed along the years, and a large body of literature, including these two impossibility theorems, is devoted to their axiomatic characterization [Arrow et al., 2010]. In fact, different voting rules can select very different alternatives. Here is a motivating example of our research question.

Example 15. *Consider an election with five candidates and fourteen voters, with voters' preferences as follows:*

$$\begin{array}{lccccccccc}
\forall i \in \{1, \cdots, 5\}, & x_1 & \succ_i & x_3 & \succ_i & x_2 & \succ_i & x_4 & \succ_i & x_5 \\
\forall i \in \{6, \cdots, 9\}, & x_5 & \succ_i & x_2 & \succ_i & x_3 & \succ_i & x_4 & \succ_i & x_1 \\
\forall i \in \{10, \cdots, 12\}, & x_4 & \succ_i & x_3 & \succ_i & x_2 & \succ_i & x_5 & \succ_i & x_1 \\
\forall i \in \{13, \cdots, 14\}, & x_2 & \succ_i & x_4 & \succ_i & x_5 & \succ_i & x_3 & \succ_i & x_1
\end{array}$$

In this preference profile, candidate x_1 is the winner under plurality, x_2 under Borda, x_3 is the Condorcet winner (and thus the winner under any Condorcet-consistent rule) and also the winner under 2-approval, while x_5 wins under plurality with runoff.

This example has illustrated that the election winner can strongly depend on the choice of voting rule. The purpose of this chapter is to investigate whether such phenomena occur frequently. More specifically, we aim to study cases in which preference profiles exhibit a high degree of structure, such as when preferences are aligned along a single axis or centered around a specific preference.

Note that for the rest of the chapter, we limit our work to positional scoring rules and Condorcet-consistent rules. Also, in the case of $m = 2$ candidates, all PSRs coincide and gaining one point for a candidate in a PSR is equivalent for this candidate to be ranked before the other candidate, breaking the gap between absolute and relative evaluation of candidates. In that case, majority voting can appear as the only reasonable voting rule [May, 1952]. Therefore, given the focus of our chapter, we reasonably assume that $m > 2$.

This second example illustrates that a slight modification of the preference profile can completely alter the previously observed phenomenon regarding the agreement between voting rules. Indeed, by changing the preferences of only two voters, we observe the exact opposite behavior: plurality, Borda, 2-approval and Condorcet-consistent rules all agree on the same winner.

Example 16. *Consider an election with five candidates and fourteen voters, with voters' preferences as follows:*

$$\begin{array}{lccccccccc}
\forall i \in \{1, \cdots, 5\}, & x_1 & \succ_i & x_3 & \succ_i & x_2 & \succ_i & x_4 & \succ_i & x_5 \\
\forall i \in \{6, 7\}, & x_5 & \succ_i & x_2 & \succ_i & x_3 & \succ_i & x_4 & \succ_i & x_1 \\
\forall i \in \{8, 9\}, & x_3 & \succ_i & x_2 & \succ_i & x_5 & \succ_i & x_4 & \succ_i & x_1 \\
\forall i \in \{10, 11, 12\}, & x_4 & \succ_i & x_3 & \succ_i & x_2 & \succ_i & x_5 & \succ_i & x_1 \\
\forall i \in \{13, 14\}, & x_2 & \succ_i & x_4 & \succ_i & x_5 & \succ_i & x_3 & \succ_i & x_1
\end{array}$$

In this profile, candidate x_3 is the winner under plurality, Borda, 2-approval and is the Condorcet winner (and thus the winner under any Condorcet-consistent rule).

These two examples are particularly illustrative, as they show two different behaviors can arise: a disagreement of all voting rules under consideration in Example 15 or an alignment of all voting rules as Example 16. This motivates the current chapter, which aims to quantify how frequently one phenomenon or its complement arises.

This issue has been raised by many articles [Gehrlein and Lepelley, 2010] which study the probability that different voting rules disagree on their outcome. Indeed, exploring the agreement among voting rules can help understand the similarity between voting rules, in an orthogonal perspective than the axiomatic study.

Most of the works on voting rules' agreement focus on the impartial (anonymous) culture, where each preference order (or score), is uniformly drawn from the whole set of linear orders over candidates. Such study is necessary because the impartial culture can arguably be seen as the most neutral. However, it does not capture real voters' preferences which are usually far from being uniformly distributed as survey in [Boehmer et al., 2024]. Moreover, most results on impartial culture highlight that voting rules not often agree [Merlin et al., 2000]. Therefore, exploring more structured and realistic cultures may provide new insights on differences between voting rules. In this chapter, we will focus on cultures generating single-peaked preferences [Black, 1948], which make sense in several contexts such as, e.g., political elections where a left-right axis can structure most voters' preferences. Even though single-peaked cultures are still far from being a perfect match to real data [esc, 2021], they are much more realistic than impartial culture, so these models can be seen as a better approximation of the reality in some contexts (e.g., political elections).

In another point of view, studying agreement between voting rules under single-peaked cultures can also improve the understanding of such cultures. A key question in computational social choice, and in particular in voting theory, is how to generate relevant synthetic data for experiments on elections [Boehmer et al., 2024], as already explained in Section 2.6.

Among all statistical distributions to generate synthetic data, single-peaked distributions are quite often used, as reported by Boehmer et al. [2024]. Therefore, exploring voting rules' agreement under single-peaked cultures is relevant to better understand these commonly used cultures and better interpret experimental studies. Let us illustrate possible issues in the interpretation of experiments. For instance, if one would like to compare how often different rules violate the majority criterion (i.e., a candidate ranked first by half of the voters should be elected), then experiments could be used. However, the conclusions may be very different depending on the voting culture used to generate synthetic data. In particular, using single-peaked cultures may lead to different conclusions compared to impartial culture, especially if the results on voting rules' agreement are very different. In particular, if two voting rules frequently agree under a given culture then the results will be similar because the voting rules are close under that culture, not because of the problem itself. In any case, knowing how the statistical tool works is a prerequisite for a good empirical study.

In this chapter, we study the probability of agreement of different voting rules under single-peaked cultures. Up to our best knowledge, this question has been surprisingly neglected for cultures more structured than impartial ones. One notable exception is the work of Chatterjee and Storcken [2020] on unimodal profiles. We focus our study on two well-known models to generate single-peaked elections: Walsh's [Walsh, 2015] and Conitzer's [Conitzer, 2007] models. They consider different ways of uniformly drawing single-peaked preference orders: either uniformly within the whole single-peaked domain [Walsh, 2015], or uniformly with respect to the peak candidate in the order [Conitzer, 2007].

We particularly examine positional scoring rules (PSRs), which compute scores for the candidates based on their position in the voters' preferences. This family covers many famous voting rules, such as k-approval rules like plurality or veto, and the Borda rule. We show that for both Walsh's and Conitzer's distributions, many PSRs tend to elect the median candidate(s) in the single-peaked axis, which turns out to be the asymptotic Condorcet winner, implying that these rules also agree with Condorcet-consistent rules. We also provide a lower bound on the speed of convergence to such a winner, meaning that this result holds for reasonable election sizes. We characterize these rules for both cultures and observe that this set is larger for Walsh's distribution, which is coherent with its definition. Conitzer's distribution seems to be more neutral toward the candidates, in the sense of probability to be elected. We further study this aspect by examining when single-peaked distributions are unbiased, i.e., when they do not favor any candidate with respect to a given voting rule.

We also provide some insights on the agreement among voting rules under two other structured preference distributions: unimodal distributions, which include Mallows' cultures [Mallows, 1957], where we complete Theorem 4.1 of Chatterjee and Storcken [2020] to prove a rapid convergence to a large probability of agreement; and Pólya-Eggenberger urns [Eggenberger and Pólya, 1923], where we show that even if the probability of agreement remains high, the convergence toward one is not guaranteed. Finally, we examine real-world elections to understand how these phenomena appear in practice.

Related Work:

The question of agreement among voting rules was initiated by Gehrlein and Fishburn [1980, 1983] who give an explicit probability of agreement between two positional scoring rules in the case of three candidates under impartial culture. They prove that the probability of all scoring rules to agree in large elections is 0.5346. Numerical experiments have also been conducted to determine [Gehrlein, 1986; Nurmi, 2012] the value of the probability of agreement on synthetic data. Many necessary conditions have then been derived to characterize the agreement of all positional scoring rules [Merlin et al., 2000; Moulin, 1989; Saari, 2012], or the violation of the majority principle [Lepelley and Merlin, 1998]. In particular, Merlin et al. [2000] give the probability (i.e., 0.50116) under impartial culture that many rules (including positional scoring rules, elimination rules and Condorcet-consistent rules) agree on the same winner in the case of three candidates. This work was complemented via Monte-Carlo simulations by Lepelley et al. [2000a] for more than three candidates. Similar results with explicit formulas have been found under anonymous impartial culture for three candidates [Gehrlein, 2002]. Most of these works focus on three candidates, sometimes four [Kamwa and Merlin, 2019], under the impartial (sometimes anonymous) culture and try to provide explicit formulas. One closely related line of work studies agreement among elimination voting rules with three candidates under IAC restricted to the single-peaked domain [Lepelley and Vidu, 2000]. In contrast, we focus on single-peaked distributions with an arbitrary number of candidates and analyze the conditions of convergence toward the same outcome.

In another perspective, many works have studied the Condorcet efficiency of voting rules (see Gehrlein and Lepelley [2010] for a survey), i.e., their probability to elect a Condorcet winner, which can be seen as exploring how much these rules agree with Condorcet-consistent rules. This question has also been investigated for structured cultures, such as impartial (anonymous) culture over the single-peaked domain [Gehrlein, 2003; Lepelley, 1994; Lepelley et al., 2000b], and Pólya-Eggenberger urns [Gehrlein and Lepelley, 2009] but, as far as we know, only for three candidates.

Another close question is the notion of consensus [Elkind et al., 2010; Hadjibeyli and Wilson, 2019], which is essentially setting a distance to find the closest election that satisfies consensus, i.e., the one where the minimum number of voters would disagree. Beyond voting rule agreement, the likelihood of the occurrence of voting paradoxes has been widely investigated [Gehrlein and Lepelley, 2010; Xia, 2020]. In addition, following the idea of asymptotic results, many studies have been conducted in machine learning, making the link between a voting rule and a maximum likelihood estimator [Azari Soufiani et al., 2014; Caragiannis et al., 2014; Xia, 2014]. In the same perspective, a work on the asymptotic probability of ties in elections was proposed [Xia, 2021]. While these directions may sometimes be outside of voting theory, it highlights the importance of our research question.

3.2 . Probabilistic Frameworks Adapted for Voting Rules

This section is devoted to presenting the probabilistic foundations of our problem, specifically the notion of expected winners within a culture, its convergence bounds, and how sampling can be performed in the single-peaked domain.

When considering independent and identical voter preference drawings, the culture can be defined as drawing n preference orders $\succ_i$ from a given preference distribution $\pi^m : \Pi^m \to [0,1]$ with $\sum_{\succ_i \in \Pi^m} \pi^m(\succ_i) = 1$. The probability for a candidate x_j to be ranked at position $k \in [m]$ under preference distribution π^m is given by $\mathbb{P}^m_\pi(j,k) = \sum_{\succ_i \in \Pi^m : r_{\succ_i}(x_j)=k} \pi^m(\succ_i)$. Moreover, the probability for a candidate x to be ranked before a candidate y under preference distribution π^m is given by $\mathbb{P}^m_\pi(x \succ_i y) = \sum_{\succ_i \in \Pi^m : x \succ_i y} \pi^m(\succ_i)$. When the context is clear, the superscript m may be omitted.

Let $S^{\mathcal{F}}(x)$ denote the random variable giving the score of a candidate $x \in M$ for a voting rule $\mathcal{F}$ when all preference orders are drawn identically. Let $\mathbb{E}_\pi[S^{\mathcal{F}}(x)]$ denote the expected score of candidate x for voting rule $\mathcal{F}$ under distribution π. For a PSR $\mathcal{F}$ characterized by a positional score vector α and a preference distribution π, the expected score of each candidate x is given by $\mathbb{E}_\pi[S^{\mathcal{F}}(x)] = \sum_{\succ_i \in \Pi^m} \pi(\succ_i) \cdot \alpha_{r_{\succ_i}(x)}$.

A special attention should be paid to the last subsection, Section 3.7.2, where voter preferences are not drawn independently and identically.

3.2.1 . Convergence to the Expected Winners

When voters' preferences are identically and independently drawn with respect to distribution π and $\mathbb{E}_\pi[S^{\mathcal{F}}(x)]$ is finite for any $x \in M$, by the law of large numbers, the expected winners $\mathcal{W}_\pi(\mathcal{F})$ of $\mathcal{F}$ under π are $\mathcal{W}_\pi(\mathcal{F}) := \arg\max_{x\in M} \mathbb{E}_\pi[S^{\mathcal{F}}(x)]$. A candidate x is an asymptotic (weak) Condorcet winner under distribution π if $\mathbb{P}_\pi(x \succ_i y) > \frac{1}{2}$ (resp., $\mathbb{P}_\pi(x \succ_i y) \geq \frac{1}{2}$), for every $y \in M \setminus \{x\}$.

In addition to the guarantee of convergence to the election of expected winners, we provide below a lower bound on the probability that an expected winner actually wins, when we draw voters' preferences independently and identically with respect to a distribution π. This allows to get a more concrete idea on whether this can happen in practice.

Theorem 8. *For a positional scoring rule $\mathcal{F}$ defined by a score vector α, and a preference distribution π over the set of candidates M, the set of expected winners is defined as*

$$\mathcal{W}_\pi(\mathcal{F}) = \arg\max_{x\in M} \mathbb{E}_\pi[S^{\mathcal{F}}(x)].$$

When this set is a singleton, i.e., $\mathcal{W}_\pi(\mathcal{F}) = \{x\}$, the probability that $\mathcal{F}$ elects x satisfies

$$\mathbb{P}_\pi(x \in \mathcal{F}(\mathscr{P})) \geq L_\pi(\mathcal{F}),$$

where:

$$L_\pi(\mathcal{F}) := 1 - 2 \cdot \max_{y\in M\setminus\mathcal{W}_\pi(\mathcal{F})} \exp\left(\frac{-2n \cdot \left(\mu_\pi^{\mathcal{F}}(y) - \mathbb{E}_\pi[S^{\mathcal{F}}(y)]\right)^2}{(\max_j \alpha_j - \min_j \alpha_j)^2}\right)$$

and $$\mu_\pi^{\mathcal{F}}(y) := \frac{\max_{x\in M} \mathbb{E}_\pi[S^{\mathcal{F}}(x)] + \mathbb{E}_\pi[S^{\mathcal{F}}(y)]}{2}$$

Proof. Let $\mathbb{E}_\pi[S^{\mathcal{F}}(y)]_i$ be the expected score of candidate y with rule $\mathcal{F}$ for voter i. Let $\mathcal{W}_\pi(\mathcal{F}) = \{x\}$, we have:

$$\mathbb{P}_\pi(x \in \mathcal{F}(\mathscr{P})) = \mathbb{P}_\pi[\forall y \neq x, \sum_{i=1}^n S^{\mathcal{F}}(x)_i > \sum_{i=1}^n S^{\mathcal{F}}(y)_i]$$

Using Bonferroni's inequality, which states that for any events $A_1, \ldots, A_m$ in a probability space $(\Omega, \mathcal{F}, \mathbb{P})$,

$$\mathbb{P}(A_1 \cap \cdots \cap A_m) \geq \sum_{i=1}^m \mathbb{P}(A_i) - (m-1),$$

we obtain:

$$\mathbb{P}_\pi[\forall y \neq x, \sum_{i=1}^n S^{\mathcal{F}}(x)_i > \sum_{i=1}^n S^{\mathcal{F}}(y)_i]$$

$$\geqslant \sum_{x\neq y} \mathbb{P}_\pi[\sum_{i=1}^n S^{\mathcal{F}}(x)_i > \sum_{i=1}^n S^{\mathcal{F}}(y)_i] - (m-2)$$

$$\geqslant (m-1) \cdot \min_{y \neq x} \mathbb{P}_\pi[\sum_{i=1}^{n} S^{\mathcal{F}}(x)_i > \sum_{i=1}^{n} S^{\mathcal{F}}(y)_i] - (m-2)$$

Let us now compute a lower bound for $\mathbb{P}_\pi[\sum_{i=1}^{n} S^{\mathcal{F}}(x)_i > \sum_{i=1}^{n} S^{\mathcal{F}}(y)_i]$. Using again Bonferroni's inequality we have:

$$\mathbb{P}_\pi[\sum_{i=1}^{n} S^{\mathcal{F}}(x)_i > \sum_{i=1}^{n} S^{\mathcal{F}}(y)_i]$$

$$\geqslant \mathbb{P}_\pi[\sum_{i=1}^{n} S^{\mathcal{F}}(x)_i < n \cdot \mu_\pi^{\mathcal{F}}(y)] + \mathbb{P}_\pi[\sum_{i=1}^{n} S^{\mathcal{F}}(x)_i > n \cdot \mu_\pi^{\mathcal{F}}(y)] - 1$$

Now, we work on each term separately,

$$\mathbb{P}_\pi[\sum_{i=1}^{n} S^{\mathcal{F}}(x)_i < n \cdot \mu_\pi^{\mathcal{F}}(y)]$$

$$= 1 - \mathbb{P}_\pi[\sum_{i=1}^{n} S^{\mathcal{F}}(x)_i \geqslant n \cdot \mu_\pi^{\mathcal{F}}(y)]$$

Using the first Hoeffding's inequality (Lemma 4) with $a_i = \min_y \alpha_y$ and $b_i = \max_y \alpha_y$,

$$\mathbb{P}_\pi[\sum_{i=1}^{n} S^{\mathcal{F}}(x)_i \geqslant n \cdot \mu_\pi^{\mathcal{F}}(y)]$$

$$= \mathbb{P}_\pi[\sum_{i=1}^{n} S^{\mathcal{F}}(x)_i - n \cdot \mathbb{E}_\pi[S^{\mathcal{F}}(y)]_i \geqslant n \cdot \mu_\pi^{\mathcal{F}}(y) - n \cdot \mathbb{E}_\pi[S^{\mathcal{F}}(y)]_i]$$

$$\leqslant e^{\frac{-2n(\mu_\pi^{\mathcal{F}}(y) - \mathbb{E}_\pi[S^{\mathcal{F}}(y)]_i)^2}{(\max_y \alpha_y - \min_y \alpha_y)^2}}$$

We reproduce the exact same reasoning for the second term $\mathbb{P}_\pi[\sum_{i=1}^{n} S^{\mathcal{F}}(x)_i > n \cdot \mu_\pi^{\mathcal{F}}(y)]$ but we use the second Hoeffding's inequality (Lemma 5). We summarize and find:

$$\mathbb{P}_\pi[\sum_{i=1}^{n} S^{\mathcal{F}}(x)_i > \sum_{i=1}^{n} S^{\mathcal{F}}(y)_i]$$

$$\leqslant 1 - e^{\frac{-2n(\mu_\pi^{\mathcal{F}}(y) - \mathbb{E}_\pi[S^{\mathcal{F}}(y)]_i)^2}{(\max_y \alpha_y - \min_y \alpha_y)^2}} - e^{\frac{-2n(\mu_\pi^{\mathcal{F}}(y) - \mathbb{E}_\pi[S^{\mathcal{F}}(x)]_i)^2}{(\max_y \alpha_y - \min_y \alpha_y)^2}}$$

Finally, we get:

$$\mathbb{P}_\pi(x \in \mathcal{F}(\mathscr{P})) \geqslant 1 - 2 \cdot \max_{y \neq x} e^{\frac{-2n \cdot (\mu_\pi^{\mathcal{F}}(y) - \mathbb{E}_\pi[S^{\mathcal{F}}(y)])^2}{(\max_y \alpha_y - \min_y \alpha_y)^2}}$$

□

We can thus deduce a lower bound for the speed of convergence for the agreement of several voting rules.

Corollary 9. *For two positional scoring rules $\mathcal{F}_1$ and $\mathcal{F}_2$ whose expected winner under a preference distribution π are the same, i.e., $C := \mathcal{W}_\pi(\mathcal{F}_1) = \mathcal{W}_\pi(\mathcal{F}_2)$, the probability of their agreement for electing the same unique candidate from C is such that: $\mathbb{P}_\pi(\mathcal{F}_1(\mathscr{P}) = \mathcal{F}_2(\mathscr{P})) \geqslant \min\{L_\pi(\mathcal{F}_1), L_\pi(\mathcal{F}_2)\}$.*

Proof. We apply twice Theorem 8 and deduce that both rules $\mathcal{F}_1$ and $\mathcal{F}_2$ have to agree on the same outcome with a probability higher than the minimum of both lower bounds. □

3.2.2 . Single-Peaked Distributions

We recall that single-peaked cultures are distributions $C(n, \Pi^m)$ such that for a given axis $>$ over M, $C(n, \Pi^m) = C(n, \Pi^m_>)$. This subsection addresses the question of how to sample single-peaked elections without relying on any additional information. Specifically, we assume that all candidates occupying a similar position relative to the axis $>$ should be treated identically.

We define a bijection on ranking $\tau : [m] \to [m]$ as a function that associates to a ranking another ranking.

Let us define the symmetry with respect to the single-peaked axis via the bijection $\tau : [m] \to [m]$ which associates with each candidate x_j its symmetric candidate $x_{\tau(j)}$ where $\tau(j) = m - j + 1$. We denote by $[\succ]^\tau =\succ'$ the preference order obtained by replacing each candidate in $\succ$ with its image under the function τ. We give a small example to illustrate this concept.

Example 17. *If $m = 4$, then $[x_2 \succ_i x_3 \succ_i x_4 \succ_i x_1]^\tau = x_3 \succ_i x_2 \succ_i x_1 \succ_i x_4$.*

A single-peaked preference distribution $\pi : \Pi^m_> \to [0, 1]$ is said to be *symmetric* if $\mathbb{P}^m_\pi(j, 1) = \mathbb{P}^m_\pi(\tau(j), 1)$, for every candidate $x_j \in M$. Symmetric single-peaked distributions form a rather large family of single-peaked distributions which include, e.g., the distributions π such that $\mathbb{P}^m_\pi(x_j \succ_i x_{j+1}) = \mathbb{P}^m_\pi(x_{\tau(j)} \succ_i x_{\tau(j+1)})$ for every $j \in [\lfloor \frac{m}{2} \rfloor]$, but not only. Using symmetric single-peaked distributions turns out to be very natural, in order to derive experiments on the single-peaked domain, without any additional information than the single-peaked axis. Indeed, these distributions satisfy the property of treating similar candidates similarly, which we believe to be the least biased assumption one can make in the absence of prior information.

In particular, two distributions have met consensus in the literature to sample single-peaked elections, and we already presented them in Chapter 2: Walsh's distribution (Definition 7) and Conitzer's distribution (Definition 8). They happen to be symmetric and capture different types of impartial culture on the single-peaked domain. Roughly, the idea is either to uniformly draw every single-peaked preference order in Walsh's model, or to uniformly draw every peak candidate and then construct the rest of the preference order by uniformly choosing the next candidate to rank between the closest available candidates on the single-peaked axis in Conitzer's model. One point worth mentioning is that under Conitzer's distribution, each candidate is elected with the same probability under the plurality rule. In a sense, this culture can be considered unbiased toward all candidates with respect to plurality. This idea will later be explored and developed in much greater depth in Section 3.6.

In this chapter, we aim at understanding the behavior of voting rules under single-peaked distributions. In particular, we analyze the conditions under which PSRs agree, the location of the expected winners with respect to the single-peaked axis and whether they are asymptotic (weak) Condorcet winners.

3.3 . Asymptotic Condorcet Winners in The Single-Peaked Domain

We begin by introducing some preliminaries on the single-peaked domain, including the enumeration of single-peaked preference orders where a specific candidate is ranked at a given position, as well as a first result concerning the asymptotic election of the Condorcet winner. Let us start with structural properties of the single-peaked domain. We first recall that $|\Pi^m_{\triangleright}| = 2^{m-1}$.

Observation 10. *Candidate x_j can never be ranked at a position $k > \max\{j, m-j+1\}$ in a single-peaked order.*

Proof. If $k > m-j+1$ (resp., $k > j$), then it means that there are not enough positions between position k and position m to place at least all candidates $x_1, \ldots, x_{j-1}$ (resp., $x_{j+1}, \ldots, x_m$), which is necessary in order to rank x_j at position k, by single-peakedness. It follows that, under such a condition, no single-peaked preference order can rank x_j at position k. □

Lemma 11 (Boehmer et al. [2022]). *The number of single-peaked preference orders in $\Pi^m_{\triangleright}$ in which candidate x_j is ranked at position k is given by the following formula, for each $j, k \in [m]$:*

$$\mathscr{D}_m(j,k) = 2^{k-2}\left(\binom{m-k}{j-1} + \binom{m-k}{j-k}\right).$$

Let C^* denote the set of median candidates in the single-peaked axis, this set is a singleton in case m is odd and is a pair of candidates in case m is even, i.e.,

$$C^* := \begin{cases} \{x_{\lceil \frac{m}{2} \rceil}\} & \text{if } m \text{ is odd} \\ \{x_{\frac{m}{2}}, x_{\frac{m}{2}+1}\} & \text{if } m \text{ is even} \end{cases}$$

By convention, $\binom{n}{k} = 0$ when $k > n$ or $k < 0$.

These candidates play an important role in the single-peaked domain. We first show below that more preference orders rank them at good positions compared to the other candidates.

Lemma 12. *For every median candidate $x_c \in C^*$ and any other candidate $x_j \in M \setminus C^*$, there exists an index $\gamma_m(j) \in [\max\{j, m-j+1\}]$ such that $\mathscr{D}_m(c,k) \geq \mathscr{D}_m(j,k)$ for every $1 \leq k \leq \gamma_m(j)$ and $\mathscr{D}_m(j,k) > \mathscr{D}_m(c,k)$ for every $\gamma_m(j) < k \leq \max\{j, m-j+1\}$.*

Proof. Let us compare a median candidate $x_c \in C^*$ and another candidate $x_j \in M \setminus C^*$ where, w.l.o.g., $c := \lceil \frac{m}{2} \rceil$ and $j < c$. Our goal is to compare $\mathscr{D}_m(c,k)$ and $\mathscr{D}_m(j,k)$ for a given position $k \in [m-j+1]$, and thus, by Lemma 11, to compare $\binom{m-k}{c-1} + \binom{m-k}{c-k}$ and $\binom{m-k}{j-1} + \binom{m-k}{j-k}$. Observe that $\binom{m-k}{c-k} = \binom{m-k}{m-c} = \binom{m-k}{\lfloor \frac{m}{2} \rfloor}$, and thus $\binom{m-k}{c-k} = \binom{m-k}{c-\mathbb{1}_{\{m \text{ odd}\}}}$, implying that $\binom{m-k}{c-k} = \binom{m-k}{c-1}$ when m is odd.

Let us recall that, when n is fixed, the binomial coefficient $\binom{n}{\ell}$ is strictly increasing from $\ell = 0$ to $\ell = \frac{n}{2}$ and then strictly decreasing from $\ell = \frac{n}{2}$ to $\ell = n$ (in case n is odd, the two maximal values are taken for $\ell = \lfloor \frac{n}{2} \rfloor$ and $\ell = \lceil \frac{n}{2} \rceil$, so it is fine to simply consider that the closest ℓ is to $\frac{n}{2}$, the biggest the value $\binom{n}{\ell}$). In our case, we have $n = m-k$, therefore the maximal value of $\binom{m-k}{\ell}$ is taken for ℓ the closest to $\frac{m-k}{2}$. Observe that $\frac{m-k}{2} \geq c-k$. It follows that $\binom{m-k}{c-k} \geq \binom{m-k}{j-k}$, since we are in the increasing part. Moreover, the maximal value $\ell = \frac{m-k}{2}$ is always closer to $c-1$ than to $j-k$: if $\frac{m-k}{2} \geq c-1$ it is obvious and if $\frac{m-k}{2} < c-1$, then supposing $j-k$ is closer would imply $\frac{m-k}{2} - j + k < c - 1 - \frac{m-k}{2}$ and thus, because $j < c = \lceil \frac{m}{2} \rceil$, we would have $c > m - j + 1 \geq \lfloor \frac{m}{2} \rfloor + 2$, a contradiction. It follows that $\binom{m-k}{c-1} \geq \binom{m-k}{j-k}$.

First observe that if $k \leq c-j+1$, then we have $j-1 \leq c-k$ and thus $j-k \leq j-1 \leq c-k \leq c-1$. Since $c-k \leq \frac{m-k}{2}$, it follows that $\binom{m-k}{c-k} \geq \binom{m-k}{j-1}$, and thus, since $\binom{m-k}{c-1} \geq \binom{m-k}{j-k}$, we have that $\mathscr{D}_m(c,k) \geq \mathscr{D}_m(j,k)$.

By Observation 10, we know that $\mathscr{D}_m(j,k) = 0$ iff $k > \max\{j, m-j+1\}$, therefore $\mathscr{D}_m(c,k) = 0$ when $k > \lfloor \frac{m}{2} \rfloor + 1$ and $\mathscr{D}_m(j,k) = 0$ when $j > m-j+1$. It follows that $\mathscr{D}_m(j,k) > \mathscr{D}_m(c,k)$ for all $\lfloor \frac{m}{2} \rfloor + 1 < k \leq m-j+1$.

To summarize, now we know that when $k \leq c-j+1$, we have $\mathscr{D}_m(c,k) \geq \mathscr{D}_m(j,k)$ and when $k > \lfloor \frac{m}{2} \rfloor + 1$, we have $\mathscr{D}_m(j,k) > \mathscr{D}_m(c,k)$. It means that there exists an index k such that $\mathscr{D}_m(c,k) \geq \mathscr{D}_m(j,k)$ and $\mathscr{D}_m(j,k+1) > \mathscr{D}_m(c,k+1)$. Let us consider the greatest such index k_0 as our base case and suppose, by induction, that $\mathscr{D}_m(c,k') \geq \mathscr{D}_m(j,k')$ for all indices k' such that $k \leq k' \leq k_0$ for a given index $k \leq k_0$. We will prove that if $\mathscr{D}_m(c,k) \geq \mathscr{D}_m(j,k)$ holds, then $\mathscr{D}_m(c,k-1) \geq \mathscr{D}_m(j,k-1)$ also holds, which will be sufficient to prove our statement about the existence of a unique threshold $\gamma_m(j)$ to distinguish when $\mathscr{D}_m(c,k) \geq \mathscr{D}_m(j,k)$ and when $\mathscr{D}_m(j,k) > \mathscr{D}_m(c,k)$.

Suppose that $\mathscr{D}_m(c,k) \geq \mathscr{D}_m(j,k)$ holds for a given position k. By Lemma 11, it means that $\binom{m-k}{c-1} + \binom{m-k}{c-\mathbb{1}_{\{m \text{ odd}\}}} \geq \binom{m-k}{j-1} + \binom{m-k}{j-k}$. By Pascal's identity, we thus have $\binom{m-k+1}{c-1} - \binom{m-k}{c-2} + \binom{m-k+1}{c-\mathbb{1}_{\{m \text{ odd}\}}} - \binom{m-k}{c-1-\mathbb{1}_{\{m \text{ odd}\}}} \geq \binom{m-k+1}{j-1} - \binom{m-k}{j-2} + \binom{m-k+1}{j-k+1} - \binom{m-k}{j-k+1}$, which is equivalent to $\binom{m-k+1}{c-1} + \binom{m-k+1}{c-\mathbb{1}_{\{m \text{ odd}\}}} \geq \binom{m-k+1}{j-1} + \binom{m-k+1}{j-k+1} + \binom{m-k}{c-2} + \binom{m-k}{c-1-\mathbb{1}_{\{m \text{ odd}\}}} - \binom{m-k}{j-2} - \binom{m-k}{j-k+1}$. If $\binom{m-k}{c-2} + \binom{m-k}{c-1-\mathbb{1}_{\{m \text{ odd}\}}} - \binom{m-k}{j-2} - \binom{m-k}{j-k+1} \geq 0$ holds, then we have $\binom{m-k+1}{c-1} + \binom{m-k+1}{c-\mathbb{1}_{\{m \text{ odd}\}}} \geq \binom{m-k+1}{j-1} + \binom{m-k+1}{j-k+1}$ and our claim follows, i.e., we have $\mathscr{D}_m(c, k-1) \geq \mathscr{D}_m(j, k-1)$. Let us thus assume, for the sake of contradiction, that $\binom{m-k}{c-2} + \binom{m-k}{c-1-\mathbb{1}_{\{m \text{ odd}\}}} - \binom{m-k}{j-2} - \binom{m-k}{j-k+1} < 0$.

$$
\begin{aligned}
&\binom{m-k}{c-2} + \binom{m-k}{c-1-\mathbb{1}_{\{m \text{ odd}\}}} < \binom{m-k}{j-2} + \binom{m-k}{j-k+1} \\
\Leftrightarrow &\binom{m-k}{c-1} \cdot \frac{c-1}{m-k-c+2} + \binom{m-k}{c-\mathbb{1}_{\{m \text{ odd}\}}} \cdot \frac{c-\mathbb{1}_{\{m \text{ odd}\}}}{m-k-c+1+\mathbb{1}_{\{m \text{ odd}\}}} \\
&< \binom{m-k}{j-1} \cdot \frac{j-1}{m-k-j+2} + \binom{m-k}{j-k} \cdot \frac{m-j}{j-k+1} \\
\Leftrightarrow &\frac{c-1}{m-k-c+2} \cdot \left(\binom{m-k}{c-1} + \binom{m-k}{c-\mathbb{1}_{\{m \text{ odd}\}}}\right) + \\
&\binom{m-k}{c} \cdot \frac{(m-k+1) \cdot \mathbb{1}_{\{m \text{ even}\}}}{(m-k-c+2)(m-k-c+1)} < \\
&\binom{m-k}{j-1} \cdot \frac{j-1}{m-k-j+2} + \binom{m-k}{j-k} \cdot \frac{m-j}{j-k+1}
\end{aligned}
$$

Since we have assumed $\mathscr{D}_m(c, k) \geq \mathscr{D}_m(j, k)$, it follows that:

$$
\begin{aligned}
&\frac{c-1}{m-k-c+2} \cdot \left(\binom{m-k}{j-1} + \binom{m-k}{j-k}\right) + \\
&\binom{m-k}{c} \cdot \frac{(m-k+1) \cdot \mathbb{1}_{\{m \text{ even}\}}}{(m-k-c+2)(m-k-c+1)} < \\
&\binom{m-k}{j-1} \cdot \frac{j-1}{m-k-j+2} + \binom{m-k}{j-k} \cdot \frac{m-j}{j-k+1} \\
\Leftrightarrow &\binom{m-k}{j-1} \cdot \left(\frac{c-1}{m-k-c+2} - \frac{j-1}{m-k-j+2}\right) + \\
&\binom{m-k}{c} \cdot \frac{(m-k+1) \cdot \mathbb{1}_{\{m \text{ even}\}}}{(m-k-c+2)(m-k-c+1)} < \\
&\binom{m-k}{j-k} \cdot \left(\frac{m-j}{j-k+1} - \frac{c-1}{m-k-c+2}\right) \\
\Leftrightarrow &\binom{m-k}{j-1} \cdot \left(\frac{c-1}{m-k-c+2} - \frac{j-1}{m-k-j+2}\right) + \\
&\binom{m-k}{j-1} \cdot \frac{\prod_{p=1}^{c-j+1}(m-k-c+p)}{\prod_{p=1}^{c-j+1}(j-1+p)} \cdot \frac{(m-k+1) \cdot \mathbb{1}_{\{m \text{ even}\}}}{(m-k-c+2)(m-k-c+1)} <
\end{aligned}
$$

$$
\begin{aligned}
&\binom{m-k}{j-1}\cdot\frac{\prod_{p=1}^{k-1}(j-k+p)}{\prod_{p=1}^{k-1}(m-k-j+1+p)}\cdot\left(\frac{m-j}{j-k+1}-\frac{c-1}{m-k-c+2}\right)\\
\Leftrightarrow&\left(\frac{c-1}{m-k-c+2}-\frac{j-1}{m-k-j+2}\right)+\\
&\frac{\prod_{p=1}^{c-j}(m-k-c+1+p)}{\prod_{p=1}^{c-j+1}(j-1+p)}\cdot\frac{(m-k+1)\cdot\mathbb{1}_{\{m\text{ even}\}}}{(m-k-c+2)}<\\
&\frac{\prod_{p=1}^{k-1}(j-k+p)}{\prod_{p=1}^{k-1}(m-k-j+1+p)}\cdot\left(\frac{m-j}{j-k+1}-\frac{c-1}{m-k-c+2}\right)\\
\Leftrightarrow&\frac{(c-j)(m-k+1)}{(m-k-c+2)(m-k-j+2)}+\\
&\frac{\prod_{p=1}^{c-j}(m-k-c+1+p)}{\prod_{p=1}^{c-j+1}(j-1+p)}\cdot\frac{(m-k+1)\cdot\mathbb{1}_{\{m\text{ even}\}}}{(m-k-c+2)}<\\
&\frac{\prod_{p=1}^{k-1}(j-k+p)}{\prod_{p=1}^{k-1}(m-k-j+1+p)}\cdot\frac{(m-k+1)(m-c-j+1)}{(j-k+1)(m-k-c+2)}\\
\Leftrightarrow&\frac{(c-j)}{(m-k-j+2)}+\frac{\prod_{p=1}^{c-j}(m-k-c+1+p)}{\prod_{p=1}^{c-j+1}(j-1+p)}\cdot\mathbb{1}_{\{m\text{ even}\}}<\\
&\frac{\prod_{p=1}^{k-1}(j-k+p)}{\prod_{p=1}^{k-1}(m-k-j+1+p)}\cdot\frac{(m-c-j+1)}{(j-k+1)}\\
\Leftrightarrow&(c-j)+\frac{\prod_{p=1}^{c-j+1}(m-k-c+1+p)}{\prod_{p=1}^{c-j+1}(j-1+p)}\cdot\mathbb{1}_{\{m\text{ even}\}}<\\
&\frac{\prod_{p=2}^{k-1}(j-k+p)}{\prod_{p=2}^{k-1}(m-k-j+1+p)}\cdot(m-c-j+1)\\
\Leftrightarrow&(c-j)\cdot\frac{\prod_{p=2}^{k-1}(m-k-j+1+p)}{\prod_{p=2}^{k-1}(j-k+p)}+\\
&\frac{\prod_{p=1}^{c-j+1}(m-k-c+1+p)}{\prod_{p=1}^{c-j+1}(j-1+p)}\cdot\frac{\prod_{p=2}^{k-1}(m-k-j+1+p)}{\prod_{p=2}^{k-1}(j-k+p)}\cdot\mathbb{1}_{\{m\text{ even}\}}\\
&<(m-c-j+1)\\
\Leftrightarrow&(c-j)\cdot\frac{\prod_{p=2}^{k-1}(m-k-j+1+p)}{\prod_{p=2}^{k-1}(j-k+p)}+\\
&\frac{\prod_{p=1}^{k+c-j-1}(m-k-c+1+p)}{\prod_{p=1}^{k+c-j-1}(j-k+1+p)}\cdot\mathbb{1}_{\{m\text{ even}\}}<(m-c-j+1)
\end{aligned}
$$

Since $j < c = \lceil \frac{m}{2} \rceil$, we have $m-j+1 > j$ and $m-c \geq j$, therefore

$\frac{\prod_{p=2}^{k-1}(m-k-j+1+p)}{\prod_{p=2}^{k-1}(j-k+p)} > 1$ and $\frac{\prod_{p=1}^{k+c-j-1}(m-k-c+1+p)}{\prod_{p=1}^{k+c-j-1}(j-k+1+p)} \geq 1$. It follows that $(c-j) + \mathbb{1}_{\{m \text{ even}\}} < (m-c-j+1)$, whereas $m-c+1 = \lfloor \frac{m}{2} \rfloor + 1 = c + \mathbb{1}_{\{m \text{ even}\}}$, a contradiction. □

Moreover, we show below that many natural single-peaked distributions favor the median candidates by tending to make them (weak) Condorcet winners. This proposition can be put in perspective with the theorem of Black [1958], which essentially states that every single-peaked profile admits a weak Condorcet winner, namely the median candidate(s).

Proposition 13. *Every symmetric single-peaked preference distribution makes the median candidate(s) asymptotic weak Condorcet winner(s). When m is odd, the unique median candidate is the asymptotic Condorcet winner under any symmetric single-peaked distribution π which assigns a positive probability to rank the median candidate first, i.e., $\mathbb{P}_\pi(c,1) > 0$ for $x_c \in C^*$.*

Proof. Let us consider a symmetric single-peaked distribution π. Let us compare a median candidate $x_c \in C^*$ and any other candidate $x_j \in M \setminus C^*$ where, w.l.o.g., $c = \lceil \frac{m}{2} \rceil$ and $j < c$. By single-peakedness, a preference order with a candidate x_ℓ as a peak candidate must rank x_c before x_j if $\ell \geq c$. It follows that $\mathbb{P}_\pi(x_c \succ_i x_j) \geq \sum_{\ell=c}^{m} \mathbb{P}_\pi(\ell,1)$. Recall that $\sum_{\ell=1}^{m} \mathbb{P}_\pi(\ell,1) = 1$.

If m is odd then, by symmetry, we have $\sum_{\ell=1}^{c-1} \mathbb{P}_\pi(\ell,1) = \sum_{\ell=c+1}^{m} \mathbb{P}_\pi(\ell,1)$, and thus $\mathbb{P}_\pi(x_c \succ_i x_j) \geq \sum_{\ell=c}^{m} \mathbb{P}_\pi(\ell,1) \geq \frac{1}{2}$. This inequality is strict if $\mathbb{P}_\pi(c,1) > 0$.

If m is even then, by symmetry, we have $\sum_{\ell=1}^{c} \mathbb{P}_\pi(\ell,1) = \sum_{\ell=c+1}^{m} \mathbb{P}_\pi(\ell,1)$, and thus $\mathbb{P}_\pi(x_c \succ_i x_j) \geq \sum_{\ell=c}^{m} \mathbb{P}_\pi(\ell,1) \geq \frac{1}{2}$. It remains to compare x_c with the other median candidate x_{c+1}. The arguments are similar: a preference order with a candidate x_ℓ as a peak candidate must rank x_c before x_{c+1} if $\ell \leq c$. Therefore, $\mathbb{P}_\pi(x_c \succ_i x_{c+1}) \geq \sum_{\ell=1}^{c} \mathbb{P}_\pi(\ell,1) = \frac{1}{2}$. □

Beyond the question of electing the Condorcet winner, we will explore whether certain voting rules tend to elect the median candidate within a specific single-peaked culture, namely we will investigate the expected winner under various voting rules given one single-peaked culture assumption.

3.4 . Agreement under Walsh's Distribution

We first study Walsh's distribution (Definition 7), which can be seen as impartial culture on the single-peaked domain. As such, the probability that a candidate is ranked at a given rank follows from Lemma 11.

Observation 14. *The probability $\mathbb{P}_{\pi_W}(j,k)$ that candidate x_j is ranked at position k under Walsh's distribution, for each $j,k \in [m]$, is equal to $\mathbb{P}_{\pi_W}(j,k) = \frac{\mathscr{D}_m(j,k)}{2^{m-1}}$.*

In Example 12, we observe that under plurality rule, candidate x_2 is strongly favored, as 0.5 of the total probability weight ranks x_2 in first place. The following result aims to show that, in the general case, the median candidate is always included in the set of expected winners, regardless of the election size and the chosen PSR. Specifically, we establish that this distribution favors the median candidates since their expected score under every PSR is at least as large as the one of any other candidate.

Proposition 15. *For every PSR $\mathcal{F}$, the median candidates always belong to the expected winners of $\mathcal{F}$ under Walsh's distribution, i.e., $C^* \subseteq \mathcal{W}_{\pi_W}(\mathcal{F})$.*

Proof. When comparing the expected score of a candidate $c \in C^*$ with the one of any other candidate $x_j \in M \setminus C^*$, we can restrict our attention, w.l.o.g., to the median candidate $x_c := x_{\lceil \frac{m}{2} \rceil} \in C^*$ and to any candidate x_j such that $j < \lceil \frac{m}{2} \rceil$ (by symmetry with respect to the single-peaked axis). The expected score of a candidate x_j, for Walsh's distribution and a PSR $\mathcal{F}$ characterized by the positional score vector α, is given by $\mathbb{E}_{\pi_W}[S^{\mathcal{F}}(x_j)] = \sum_{k=1}^{m} \mathbb{P}_{\pi_W}(j,k) \cdot \alpha_k = \sum_{k=1}^{m} \frac{\mathscr{D}_m(j,k)}{2^{m-1}} \cdot \alpha_k$. By the fact that $j < c = \lceil \frac{m}{2} \rceil$, we have $\max\{j, m-j+1\} = m-j+1$ and $\max\{c, m-c+1\} = \lfloor \frac{m}{2} \rfloor + 1$. And thus, by Observation 10, $\mathscr{D}_m(c,k) = 0$ for every $k > \lfloor \frac{m}{2} \rfloor + 1$ and $\mathscr{D}_m(j,k) = 0$ for every $k > m-j+1 > \lfloor \frac{m}{2} \rfloor + 1$. Therefore, we have $\mathbb{E}_{\pi_W}[S^{\mathcal{F}}(x_j)] = \sum_{k=1}^{m-j+1} \frac{\mathscr{D}_m(j,k)}{2^{m-1}} \cdot \alpha_k$ and $\mathbb{E}_{\pi_W}[S^{\mathcal{F}}(x_c)] = \sum_{k=1}^{\lfloor \frac{m}{2} \rfloor + 1} \frac{\mathscr{D}_m(j,k)}{2^{m-1}} \cdot \alpha_k$ Let us compare the expected scores of both candidates:

$$
\begin{aligned}
&\mathbb{E}_{\pi_W}[S^{\mathcal{F}}(x_c)] - \mathbb{E}_{\pi_W}[S^{\mathcal{F}}(x_j)] \\
=& \sum_{k=1}^{\lfloor \frac{m}{2} \rfloor + 1} \frac{\mathscr{D}_m(c,k)}{2^{m-1}} \cdot \alpha_k - \sum_{k=1}^{m-j+1} \frac{\mathscr{D}_m(j,k)}{2^{m-1}} \cdot \alpha_k \\
=& \frac{1}{2^{m-1}} \left(\sum_{k=1}^{\gamma_m(j)} (\mathscr{D}_m(c,k) - \mathscr{D}_m(j,k)) \cdot \alpha_k + \right. \\
& \left. \sum_{k=\gamma_m(j)+1}^{m-j+1} (\mathscr{D}_m(c,k) - \mathscr{D}_m(j,k)) \cdot \alpha_k \right) \\
\geq& \frac{\alpha_{\gamma_m(j)}}{2^{m-1}} \left(\sum_{k=1}^{\lfloor \frac{m}{2} \rfloor + 1} \mathscr{D}_m(c,k) - \sum_{k=1}^{m-j+1} \mathscr{D}_m(j,k) \right) \\
=& 0
\end{aligned}
$$

The inequality comes from the fact that, by Lemma 12, $\mathscr{D}_m(c,k) \geq \mathscr{D}_m(j,k)$, for every $k \in [\gamma_m(j)]$ and $\mathscr{D}_m(c,k) < \mathscr{D}_m(j,k)$ for every $\gamma_m(j) < k \leq m-j+1$, and that $\alpha_1 \geq \cdots \geq \alpha_\gamma \geq \cdots \geq \alpha_m$. The last equality to 0 is because $\sum_{k=1}^{\lfloor \frac{m}{2} \rfloor + 1} \mathscr{D}_m(c,k) = \sum_{k=1}^{m-j+1} \mathscr{D}_m(j,k) = 2^{m-1}$.

Hence, the expected score of a median candidate x_c is always at least as good as the expected score of any other candidate, which completes the proof. $\square$

The next result aims to go a step further by providing a characterization of all PSRs for which the median candidates are the only expected winners. Returning to Example 12, we are not only interested in knowing that b is favored under plurality, but also in identifying all rules under which b emerges as an expected winner.

We identify them as the PSRs whose associated positional score vector α is such that there exists an index $\ell \in [\lfloor \frac{m}{2} \rfloor + 1]$ with $\alpha_\ell > \alpha_{\ell+1}$. We call them *first-prioritizing* PSRs. Note that all k-approval rules for $k \leq \lfloor \frac{m}{2} \rfloor + 1$ are first-prioritizing, as well as the Borda rule.

Theorem 16. *The median candidates are the unique expected winners of a PSR $\mathcal{F}$ under Walsh's distribution, i.e., $\mathcal{W}_{\pi_W}(\mathcal{F}) = C^*$, iff $\mathcal{F}$ is first-prioritizing.*

Proof. Consider first a PSR $\mathcal{F}$ characterized by a positional score vector α such that there exists an index $\ell \in [\lfloor \frac{m}{2} \rfloor + 1]$ for which $\alpha_\ell > \alpha_{\ell+1}$. Let us compare, w.l.o.g., the median candidate x_c with $c := \lceil \frac{m}{2} \rceil$ and a candidate x_j such that $1 \leq j < c$ where, by definition, $x_j \in M \setminus C^*$. By the fact that $j < c = \lceil \frac{m}{2} \rceil$, we have $\max\{j, m-j+1\} = m-j+1$ and $\max\{c, m-c+1\} = \lfloor \frac{m}{2} \rfloor + 1$. And thus, by Observation 10, $\mathscr{D}_m(c,k) = 0$ for every $k > \lfloor \frac{m}{2} \rfloor + 1$ and $\mathscr{D}_m(j,k) = 0$ for every $k > m-j+1 > \lfloor \frac{m}{2} \rfloor + 1$. Let us compare the expected scores of both candidates:

$$
\begin{aligned}
&\mathbb{E}_{\pi_W}[S^{\mathcal{F}}(x_c)] - \mathbb{E}_{\pi_W}[S^{\mathcal{F}}(x_j)] \\
=&\sum_{k=1}^{\lfloor \frac{m}{2} \rfloor + 1} \frac{\mathscr{D}_m(c,k)}{2^{m-1}} \cdot \alpha_k - \sum_{k=1}^{m-j+1} \frac{\mathscr{D}_m(j,k)}{2^{m-1}} \cdot \alpha_k \\
=&\frac{1}{2^{m-1}} \left(\sum_{k=1}^{\gamma_m(j)} (\mathscr{D}_m(c,k) - \mathscr{D}_m(j,k)) \cdot \alpha_k + \right. \\
&\left. \sum_{k=\gamma_m(j)+1}^{m-j+1} (\mathscr{D}_m(c,k) - \mathscr{D}_m(j,k)) \cdot \alpha_k \right) \\
>&\frac{\alpha_{\gamma_m(j)}}{2^{m-1}} \left(\sum_{k=1}^{\lfloor \frac{m}{2} \rfloor + 1} \mathscr{D}_m(c,k) - \sum_{k=1}^{m-j+1} \mathscr{D}_m(j,k) \right) \\
=&0
\end{aligned}
$$

The inequality comes from the fact that, by Lemma 12, $\mathscr{D}_m(c,k) \geq \mathscr{D}_m(j,k)$, for every $k \in [\gamma_m(j)]$ and $\mathscr{D}_m(c,k) < \mathscr{D}_m(j,k)$ for every $\gamma_m(j) < k \leq m-j+1$, and that $\alpha_1 \geq \cdots \geq \alpha_{\gamma_m(j)} \geq \cdots \geq \alpha_m$. This inequality is strict because there exists an index ℓ such that $1 \leq \ell \leq \lfloor \frac{m}{2} \rfloor + 1 < m - j + 1$ for which $\alpha_\ell > \alpha_{\ell+1}$. Hence, the expected score of a median candidate is always greater than the expected score of any other candidate x_j, and thus the median candidates are the only expected winners.

Consider now a PSR $\mathcal{F}$ characterized by a positional score vector α where $\alpha_1 = \cdots = \alpha_\ell$ for a given $\ell > \lfloor \frac{m}{2} \rfloor + 1$. Let us compare, w.l.o.g., the median candidate x_c with $c := \lceil \frac{m}{2} \rceil$ and the candidate x_{c-1} (which must exist since $m > 2$) where, by definition, $x_{c-1} \in M \setminus C^*$. By the fact that $c - 1 < c = \lceil \frac{m}{2} \rceil$, we have $\max\{c, m - c + 1\} = \lfloor \frac{m}{2} \rfloor + 1$ and $\max\{c - 1, m - (c - 1) + 1\} = \lfloor \frac{m}{2} \rfloor + 2$. And thus, by Observation 10, $\mathscr{D}_m(c, k) = 0$ for every $k > \lfloor \frac{m}{2} \rfloor + 1$ and $\mathscr{D}_m(c - 1, k) = 0$ for every $k > \lfloor \frac{m}{2} \rfloor + 2$. Note that, by assumption, we have $\ell > \lfloor \frac{m}{2} \rfloor + 1$, and thus $\alpha_1 = \cdots = \alpha_{\lfloor \frac{m}{2} \rfloor + 1} = \alpha_{\lfloor \frac{m}{2} \rfloor + 2}$. Let us compare the expected scores of both candidates:

$$
\begin{aligned}
&\mathbb{E}_{\pi_W}[S^{\mathcal{F}}(x_c)] - \mathbb{E}_{\pi_W}[S^{\mathcal{F}}(x_{c-1}))] \\
=& \sum_{k=1}^{\lfloor \frac{m}{2} \rfloor + 1} \frac{\mathscr{D}_m(c, k)}{2^{m-1}} \cdot \alpha_k - \sum_{k=1}^{\lfloor \frac{m}{2} \rfloor + 2} \frac{\mathscr{D}_m(c-1, k)}{2^{m-1}} \cdot \alpha_k \\
=& \frac{\alpha_1}{2^{m-1}} \left(\sum_{k=1}^{\lfloor \frac{m}{2} \rfloor + 1} \mathscr{D}_m(c, k) - \sum_{k=1}^{\lfloor \frac{m}{2} \rfloor + 2} \mathscr{D}_m(c-1, k) \right) \\
=& 0
\end{aligned}
$$

Hence, the expected scores of x_c and x_{c-1} are equal, whereas x_{c-1} is not a median candidate. Therefore, the median candidates are not the only expected winners. □

By Proposition 13 and Theorem 16, the first-prioritizing PSRs tend to elect the (weak) Condorcet winner(s) under Walsh's distribution. It follows that asymptotically, under Walsh's distribution, all first-prioritizing PSRs and all Condorcet-consistent rules tend to agree on the same outcome, namely the median candidates.

Corollary 17. *Under Walsh's distribution, all first-prioritizing PSRs and Condorcet-consistent rules asymptotically agree to elect the median candidates.*

We show a good lower bound for the convergence to the same outcome for a subset of first-prioritizing PSRs, which include the rules k-approval and Borda rule.

Proposition 18. *For all k-approval voting rules that are first-prioritizing and the Borda rule, under Walsh's distribution, the probability of their agreement for electing one candidate from C^* is lower bounded by $L_\pi(\mathcal{F}_1)$ where $\mathcal{F}_1$ refers to the* *plurality rule. Thus, the speed of convergence for the plurality rule constitutes a lower bound.*

Proof. Let $\mathcal{A}$ be the set of all k-approval rules that are first-prioritizing. Thanks to Corollary 9, it is enough to look at $\min_{\mathcal{F}\in\mathcal{A}}\{L_\pi(\mathcal{F})\}$. Using the expression of L in Theorem 8, we can greatly simplify our question to the finding of $\mathcal{F}$ such that $\max_{x\in M}\mathbb{E}_{\pi_W}[S^{\mathcal{F}}(x)] - \mathbb{E}_{\pi_W}[S^{\mathcal{F}}(y)]$ is minimal, where $\max_{x\in M}\mathbb{E}_{\pi_W}[S^{\mathcal{F}}(x)] = \mathbb{E}_{\pi_W}[S^{\mathcal{F}}(c)]$ if $c \in C^*$ and $y \in M \setminus C^*$. However, $\max_{x\in M}\mathbb{E}_{\pi_W}[S^{\mathcal{F}}(x)] - \mathbb{E}_{\pi_W}[S^{\mathcal{F}}(y)] = \sum_{k=1}^{\lfloor\frac{m}{2}\rfloor+1}\frac{\mathscr{D}_m(c,k)}{2^{m-1}}\cdot\alpha_k - \sum_{k=1}^{m-j+1}\frac{\mathscr{D}_m(y,k)}{2^{m-1}}\cdot\alpha_k$. This can again be reduced to the following minimization: $\sum_{k=1}^{\lfloor\frac{m}{2}\rfloor+1}\mathscr{D}_m(c,k) - \sum_{k=1}^{m-j+1}\mathscr{D}_m(y,k)$. Thanks to Lemma 12 and the fact that $\sum_{k=1}^{\lfloor\frac{m}{2}\rfloor+1}\mathscr{D}_m(c,k) = 2^{m-1}$, we can conclude that the plurality rule minimizes this quantity. It remains to show that the Borda rule always reaches a higher bound L. In fact, it is enough to show that $\max_{x\in M}\mathbb{E}_{\pi_W}[S^{\mathcal{F}}(x)] - \mathbb{E}_{\pi_W}[S^{\mathcal{F}}(y)]$ is at least $m-1$ times bigger than for plurality since we will divide by $(\max_j \alpha_j - \min_j \alpha_j) = m-1$. Nevertheless, the Borda score for the candidate ranked first is $m-1$, so that this quantity has to be larger for Borda. It is even strictly larger thanks to the next Borda scores. □

As an illustration, we apply Proposition 18 for $m = 5$.

Example 18. *Let us apply Theorem 18 for $m = 5$. We have $C^* = \{x_3\}$ and take $\mathcal{F} = \mathcal{F}_1$, $\mathbb{E}_\pi[S^{\mathcal{F}}(x_3)] = \frac{3}{8}$, $\mathbb{E}_\pi[S^{\mathcal{F}}(x_2)] = \frac{1}{4}$ and $\mu_\pi^{\mathcal{F}}(x_2) = \frac{5}{16}$. We notice that these values maximize the maximum in the right hand side in Theorem 8. Finally, we have $\mathbb{P}(\mathcal{F}_1(\succ) = C^*) \geqslant 1 - 2e^{-\frac{n}{128}}$. Thus, this constitutes a lower bound for all k-approval voting rules that are first-prioritizing and Borda. For instance, when $n = 200$, the lower bound is already at 0.58, for $n = 400$, it is 0.91, and for $n = 600$, it reaches 0.98. This shows that these first-prioritizing voting rules quickly converge in agreement as the number of voters increases under Walsh's distribution.*

3.5 . Agreement under Conitzer's Distribution

We now analyze Conitzer's distribution (Definition 8), which considers a uniform distribution not on the whole single-peaked domain, as Walsh's distribution, but on the peak candidates of the single-peaked orders. In Conitzer's model, one first selects a peak candidate uniformly at random, then completes the preference by iteratively ranking one of the closest remaining candidates along the single-peaked axis. It follows that the probability for a given candidate to be ranked at a given rank is a bit less direct, as already stated by Boehmer et al. [2022].

Lemma 19 (Boehmer et al. [2022]). *The probability $\mathbb{P}_{\pi_C}(j,k)$ that candidate x_j is ranked at position k under Conitzer's distribution, for each $j,k \in [m]$, is equal to $\mathbb{P}_{\pi_C}(j,k) = Q(j,k) + Q(m-j+1,k)$ where*

$$Q(j,k) = \begin{cases} \frac{1}{2m} & \text{if } k < j \\ \frac{k}{2m} & \text{if } k = j \\ 0 & \text{otherwise} \end{cases}.$$

In Example 13, we observe that, under the plurality rule, the probability of electing any candidate is equal. However, depending on the voting rule applied, some candidates may still be favored. The following result aims to establish the expected winner according to certain specific rules. We first characterize the expected winners of all k-approval rules.

Proposition 20. *The expected winners of the k-approval rule $\mathcal{F}$ under Conitzer's distribution are:*

$$\mathcal{W}_{\pi_C}(\mathcal{F}) = \begin{cases} M \text{ if } k = 1 \\ \{x_k, x_{m-k+1}\} \text{ if } 1 < k \leq \lfloor \frac{m}{2} \rfloor + 1 \\ \{x_j \in M : \max\{j, m-j+1\} \leq k\} \text{ otherwise} \end{cases}.$$

Proof. Since Conitzer's distribution is symmetric, we restrict our analysis, w.l.o.g., to the case of a candidate x_j where $j \in [\lceil \frac{m}{2} \rceil]$. By Lemma 19, we have the following expected score for x_j:

$$\mathbb{E}_{\pi_C}[S^{\mathcal{F}}(x_j)] = \sum_{\ell=1}^{k} \mathbb{P}_{\pi_C}(j,k)$$

$$= \sum_{\ell=1}^{\min\{k,j\}} Q(j,\ell) + \sum_{\ell=1}^{\min\{k,m-j+1\}} Q(m-j+1,\ell)$$

$$= \begin{cases} \frac{2k}{2m} & \text{if } k < j \\ \frac{3k-1}{2m} + \frac{k}{2m} \cdot \mathbb{1}_{\{j = \lceil \frac{m}{2} \rceil\}} & \text{if } k = j \\ \frac{2j-1+k}{2m} & \text{if } j < k < m-j+1 \\ 1 & \text{if } k \geq m-j+1 \end{cases}$$

If $k \geq \lfloor \frac{m}{2} \rfloor + 1$, then there exist candidates x_j such that $k \geq m-j+1$, and all of them get the maximal expected score of 1, thus they are all expected winners. If $k = 1$, then $\mathbb{E}_{\pi_C}[S^{\mathcal{F}}(x_1)] = \frac{3k-1}{2m} = \frac{1}{m}$ and $\mathbb{E}_{\pi_C}[S^{\mathcal{F}}(x_j)] = \frac{2k}{2m} = \frac{1}{m}$ for all other candidates x_j. It follows that all candidates are expected winners. Finally, if $1 < k \leq \lfloor \frac{m}{2} \rfloor + 1$, then the expected winners are those corresponding to the case where $k = j$ because $3k-1 > 2k$ when $k > 1$ and $2j-1+k < 3k-1$ when $j < k$. $\square$

The next result aims to go a step further by providing a characterization of all PSRs for which the median candidates are the only expected winners.

Theorem 21. *The median candidates are the unique expected winners of a PSR $\mathcal{F}$ under Conitzer's distribution iff the positional score vector α associated with $\mathcal{F}$ satisfies the following inequality, for every $1 \leq j < \lceil \frac{m}{2} \rceil$:*

$$\sum_{\ell=j+1}^{\lceil \frac{m}{2} \rceil - 1} \alpha_\ell + \beta(m) + \alpha_{\frac{m}{2}} \mathbb{1}_{\{m \text{ even}\}} > \sum_{\ell=\lceil \frac{m}{2} \rceil + 1}^{m-j} \alpha_\ell + \delta(j,m)$$

where $\beta(m) := (\lceil \frac{m}{2} \rceil - 1)\alpha_{\lceil \frac{m}{2} \rceil} + (\lfloor \frac{m}{2} \rfloor + 1)\alpha_{\lfloor \frac{m}{2} \rfloor + 1}$ and $\delta(j,m) := (j-1)\alpha_j + (m-j+1)\alpha_{m-j+1}$.
A sufficient condition is $\beta(m) > \delta(j,m)$, for every $j < \lceil \frac{m}{2} \rceil$.

Proof. Consider a PSR $\mathcal{F}$ characterized by a positional score vector α. Let us compare a median candidate $x_c \in C^*$ and another candidate $x_j \in M \setminus C^*$ where, w.l.o.g., $j < c := \lceil \frac{m}{2} \rceil$. By Lemma 19, the expected score of candidate x_c is given by: $\mathbb{E}_{\pi_C}[S^{\mathcal{F}}(x_c)] = \frac{1}{m} \sum_{\ell=1}^{\lceil \frac{m}{2} \rceil - 1} \alpha_\ell + \frac{\lceil \frac{m}{2} \rceil}{2m} \cdot \alpha_{\lceil \frac{m}{2} \rceil} + \frac{\lfloor \frac{m}{2} \rfloor + 1}{2m} \cdot \alpha_{\lfloor \frac{m}{2} \rfloor + 1} + \frac{1}{2m} \cdot \alpha_{\frac{m}{2}} \cdot \mathbb{1}_{\{m \text{ even}\}}$.

Moreover, the expected score of candidate x_j is given by: $\mathbb{E}_{\pi_C}[S^{\mathcal{F}}(x_j)] = \frac{1}{m} \sum_{\ell=1}^{j-1} \alpha_\ell + \frac{j+1}{2m}\alpha_j + \frac{1}{2m} \sum_{\ell=j+1}^{m-j} \alpha_\ell + \frac{m-j+1}{2m}\alpha_{m-j+1}$.

It follows that the median candidates are unique expected winners iff, for every $j < c$, we have:

$$\mathbb{E}_{\pi_C}[S^{\mathcal{F}}(x_c)] - \mathbb{E}_{\pi_C}[S^{\mathcal{F}}(x_j)] > 0$$

$$\Leftrightarrow$$

$$\frac{1}{m} \sum_{\ell=1}^{\lceil \frac{m}{2} \rceil - 1} \alpha_\ell + \frac{\lceil \frac{m}{2} \rceil}{2m} \alpha_{\lceil \frac{m}{2} \rceil} + \frac{\lfloor \frac{m}{2} \rfloor + 1}{2m} \alpha_{\lfloor \frac{m}{2} \rfloor + 1} + \frac{1}{2m} \alpha_{\frac{m}{2}} \mathbb{1}_{\{m \text{ even}\}}$$

$$-\frac{1}{m} \sum_{\ell=1}^{j-1} \alpha_\ell - \frac{j+1}{2m} \alpha_j - \frac{1}{2m} \sum_{\ell=j+1}^{m-j} \alpha_\ell - \frac{m-j+1}{2m} \alpha_{m-j+1} > 0$$

$$\Leftrightarrow$$

$$\frac{1}{2m} \sum_{\ell=j+1}^{\lceil \frac{m}{2} \rceil - 1} \alpha_\ell + \frac{\lceil \frac{m}{2} \rceil - 1}{2m} \alpha_{\lceil \frac{m}{2} \rceil} + \frac{\lfloor \frac{m}{2} \rfloor + 1}{2m} \alpha_{\lfloor \frac{m}{2} \rfloor + 1} + \frac{1}{2m} \alpha_{\frac{m}{2}} \mathbb{1}_{\{m \text{ even}\}}$$

$$> \frac{j-1}{2m} \alpha_j + \frac{1}{2m} \sum_{\ell=\lceil \frac{m}{2} \rceil + 1}^{m-j} \alpha_\ell + \frac{m-j+1}{2m} \alpha_{m-j+1}$$

$$\Leftrightarrow$$

$$\sum_{\ell=j+1}^{\lceil \frac{m}{2} \rceil - 1} \alpha_\ell + (\lceil \frac{m}{2} \rceil - 1) \alpha_{\lceil \frac{m}{2} \rceil} + (\lfloor \frac{m}{2} \rfloor + 1) \alpha_{\lfloor \frac{m}{2} \rfloor + 1} + \alpha_{\frac{m}{2}} \mathbb{1}_{\{m \text{ even}\}}$$

$$> (j-1)\alpha_j + \sum_{\ell=\lceil \frac{m}{2} \rceil+1}^{m-j} \alpha_\ell + (m-j+1)\alpha_{m-j+1}$$

We always have $\sum_{\ell=j+1}^{\lceil \frac{m}{2} \rceil-1} \alpha_\ell + \alpha_{\frac{m}{2}} \mathbb{1}_{\{m \text{ even}\}} \geq \sum_{\ell=\lceil \frac{m}{2} \rceil+1}^{m-j} \alpha_\ell$. It follows that a sufficient condition to get $\mathbb{E}_{\pi_C}[S^{\mathcal{F}}(x_c)] - \mathbb{E}_{\pi_C}[S^{\mathcal{F}}(x_j)] > 0$ is $(\lceil \frac{m}{2} \rceil - 1)\alpha_{\lceil \frac{m}{2} \rceil} + (\lfloor \frac{m}{2} \rfloor + 1)\alpha_{\lfloor \frac{m}{2} \rfloor+1} > (j-1)\alpha_j + (m-j+1)\alpha_{m-j+1}$, for every $1 \leq j < \lceil \frac{m}{2} \rceil$. $\square$

We observe that the Borda rule satisfies the sufficient condition of Theorem 21, as well as $\lceil \frac{m}{2} \rceil$-approval (and $(\frac{m}{2}+1)$-approval if m is even), proving that these rules eventually elect the median candidates (as already observed in Proposition 20 for the approval rules). While Theorem 21 is not immediately interpretable, the following provides some intuition to better understand its implications. The underlying intuition is that the characterization corresponds to PSRs associated with a score vector $(\alpha_1, \cdots, \alpha_m)$ such that the first half of the scores is strictly greater than the second half but, for more than 4 candidates, not with too big a gap. More precisely, for m=3 and m=4, we must have $\alpha_2 > \alpha_3$ and $\alpha_2 > \alpha_3$ or $\alpha_3 > \alpha_4$, respectively, and for m=5, we must have $\alpha_2 > \alpha_3$ or $\alpha_3 > \alpha_4$ and $a_2 < 5 \cdot \alpha_3 - 4 \cdot \alpha_4$. Note that, in addition to Borda and $\lceil m/2 \rceil$-approval, this also includes, e.g., all PSRs such that $\alpha_i = 0$ if $i > \lfloor m/2 \rfloor + 1$ and $\alpha_2 < 2 \cdot \alpha_{\lfloor m/2 \rfloor+1}$.

A visual interpretation of Theorem 21: We now propose a visual interpretation of our result through the following example.

Example 19. *Consider the case $m = 4$. The space of positional scoring rules, renormalized so that the total score sums to 1, forms a 3-dimensional simplex. We apply Theorem 21 to identify the scoring rules that satisfy the necessary and sufficient conditions for agreement under Conitzer's distribution. The scoring rules satisfying these conditions form a region where α_1 is implicitly determined by the simplex constraint, but nothing else. This region, shown in blue in our figure, includes all rules that ensure agreement under Conitzer's distribution. For instance, Borda's rule, which corresponds to the normalized point $(\frac{1}{2}, \frac{1}{3}, \frac{1}{6}, 0)$, lies in this region. Another example is the rule $(2, 1, 1, 0)$, which normalizes to $(\frac{1}{2}, \frac{1}{4}, \frac{1}{4}, 0)$, also within the blue region. This illustrates that, with four candidates, many scoring rules satisfy the agreement condition. As the number of candidates increases, however, additional constraints emerge, and the region becomes a polytope of higher dimension.*

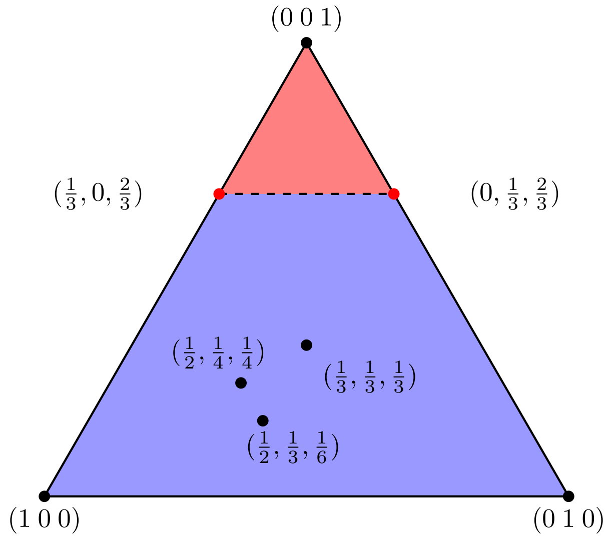

Figure 3.1: Agreement region under Conitzer's distribution ($m = 4$)

Here is a corollary to formalize the previous intuition about Theorem 21. By Proposition 13 and Corollary 22, the Borda rule, $\lceil \frac{m}{2} \rceil$-approval, as well as all rules identified in Theorem 21 tend to elect the (weak) Condorcet winner(s).

Corollary 22. *The median candidates are the unique expected winners of the Borda rule and the $\lceil \frac{m}{2} \rceil$-approval rule (as well as $(\frac{m}{2} + 1)$-approval if m is even) under Conitzer's distribution.*

Proof. We simply show that these rules satisfy the sufficient condition of Theorem 21.

The Borda rule is characterized by the positional score vector $\alpha = (m - 1, \ldots, 0)$, therefore we have $\alpha_j = m - j$, for every $j \in [m]$. Thus, for every $1 \leq j < \lceil \frac{m}{2} \rceil$, we have:

$$
\begin{aligned}
&(\lceil \frac{m}{2} \rceil - 1)\alpha_{\lceil \frac{m}{2} \rceil} + (\lfloor \frac{m}{2} \rfloor + 1)\alpha_{\lfloor \frac{m}{2} \rfloor + 1} \\
&\quad - (j-1)\alpha_j - (m-j+1)\alpha_{m-j+1} \\
=&(\lceil \frac{m}{2} \rceil - 1)\lfloor \frac{m}{2} \rfloor + (\lfloor \frac{m}{2} \rfloor + 1)(\lceil \frac{m}{2} \rceil - 1) \\
&\quad - (j-1)(m-j) - (m-j+1)(j-1) \\
=&(\lceil \frac{m}{2} \rceil - 1)(2\lfloor \frac{m}{2} \rfloor + 1) - (j-1)(2m - 2j + 1)
\end{aligned}
$$

The previous quantity is decreasing with respect to j, therefore it takes its minimum value for $j = \lceil \frac{m}{2} \rceil - 1$, where this quantity is equal to:

$$
\begin{aligned}
&(\lceil\frac{m}{2}\rceil-1)(2\lfloor\frac{m}{2}\rfloor+1)-(\lceil\frac{m}{2}\rceil-2)(2m-2\lceil\frac{m}{2}\rceil+3)\\
=&(\lceil\frac{m}{2}\rceil-2)(2\lfloor\frac{m}{2}\rfloor+1-2m+2\lceil\frac{m}{2}\rceil-3)+(2\lfloor\frac{m}{2}\rfloor+1)\\
=&(\lceil\frac{m}{2}\rceil-2)(-2)+(2\lfloor\frac{m}{2}\rfloor+1)\\
=&-2\lceil\frac{m}{2}\rceil+4+2\lfloor\frac{m}{2}\rfloor+1\\
=&-\mathbb{1}_{\{m\text{ odd}\}}+4-\mathbb{1}_{\{m\text{ odd}\}}+1\\
=&-2\cdot\mathbb{1}_{\{m\text{ odd}\}}+5\\
>&0
\end{aligned}
$$

Hence, the Borda rule satisfies the sufficient condition of Theorem 21.

Under the $\lceil\frac{m}{2}\rceil$-approval rule, $\alpha_j = 1$ for all $1 \le j \le \lceil\frac{m}{2}\rceil$ and $\alpha_j = 0$ for all $j > \lceil\frac{m}{2}\rceil$. Therefore, for every $1 \le j < \lceil\frac{m}{2}\rceil$, we have $(\lceil\frac{m}{2}\rceil - 1)\alpha_{\lceil\frac{m}{2}\rceil} + (\lfloor\frac{m}{2}\rfloor + 1)\alpha_{\lfloor\frac{m}{2}\rfloor+1} - (j-1)\alpha_j - (m-j+1)\alpha_{m-j+1} = \lceil\frac{m}{2}\rceil - 1 + (\lfloor\frac{m}{2}\rfloor + 1)\mathbb{1}_{\{m\text{ odd}\}} - (j-1) > 0$, because $j < \lceil\frac{m}{2}\rceil$. Hence, $\lceil\frac{m}{2}\rceil$-approval satisfies the sufficient condition of Theorem 21. If m is even, then $(\frac{m}{2}+1)$-approval also satisfies the sufficient condition of Theorem 21 because $\alpha_{\lfloor\frac{m}{2}\rfloor+1} = 1$ and thus we have $\lceil\frac{m}{2}\rceil - 1 + \lfloor\frac{m}{2}\rfloor + 1 - (j-1) > 0$. □

Corollary 23. *Under Conitzer's distribution, the Borda rule, $\lceil\frac{m}{2}\rceil$-approval, and Condorcet-consistent rules asymptotically agree to elect the median candidates.*

As an illustration, when we apply Theorem 8 with Borda for $m = 5$.

Example 20. *Let us apply Theorem 18 for $m = 5$. We follow the same type of computation as in Example 18 and get: $\mathbb{P}_{\pi_C}(\mathcal{F}(\mathscr{P}) = C^*) \geqslant 1 - 2e^{-\frac{9n}{3200}}$. For instance, when $n = 500$, the lower bound to elect the median candidate with Borda is already at 0.50, for $n = 1000$, it is 0.87, and for $n = 2000$, it reaches 0.99.*

We have observed a higher level of agreement among voting rules under Conitzer's distribution. This distribution has the notable property of assigning equal probability to each candidate being ranked first. In this sense, it can be interpreted as a way to restore balance among candidates with respect to the plurality rule. Building on this intuition, we now seek to extend this idea to other voting rules by introducing the notion of *unbiased distribution*.

3.6 . Unbiased Distributions

In this section, we aim to identify single-peaked distributions which do not favor any candidate by design, with respect to a given PSR. This led us to the following definition.

Definition 11 (Unbiased distribution)**.** *A preference distribution $\pi : \Pi^m \to [0,1]$ is said to be unbiased with respect to a given PSR $\mathcal{F}$ if all candidates are expected winners of $\mathcal{F}$ under π, i.e., $\mathbb{E}_\pi[S^{\mathcal{F}}(x)] = \mathbb{E}_\pi[S^{\mathcal{F}}(y)]$, for every $x, y \in M$.*

Note that the existence of an unbiased distribution with respect to a given PSR can be decided in polynomial time by solving a system of linear equations with real variables. We first characterize the single-peaked distributions which are unbiased with respect to k-approval rules.

Theorem 24. *There exists an unbiased single-peaked distribution with respect to the k-approval rule iff k divides m.*

Proof. Let us assume that k divides m, i.e., there exists an integer q such that $m = k \cdot q$. Let us partition the set of candidates M in q groups of size k as follows: $\{x_1, x_2, \dots, x_k\}, \{x_{k+1}, \dots, x_{2k}\}, \dots, \{x_{(q-1)k+1}, \dots, x_{qk}\}$ where X_j denotes the group $\{x_{(j-1)k+1}, \dots, x_{jk}\}$ for each $j \in [q]$ and $M = \bigsqcup_{j\in[q]} X_j$. For each group X_j, let us denote by P_j the set of single-peaked preference orders where the k candidates in X_j are ranked among the first k candidates, i.e., $P_j : \{\succ_i \in \Pi^m_> : r_{\succ_i}(x) \leq k, \forall x \in X_j\}$. Observe that P_j is necessarily non-empty for every $j \in [q]$ because, e.g., the following single-peaked order $\succ_i$ belongs to P_j: $x_{(j-1)k+1} \succ_i \cdots \succ_i x_{jk} \succ_i x_{(j-1)k} \succ_i \cdots \succ_i x_1 \succ_i x_{jk+1} \succ_i \cdots \succ_i x_m$. We consider the single-peaked preference distribution $\pi : \Pi^m_> \to [0,1]$ such that $\sum_{\succ_i \in P_j} \pi(\succ_i) = \frac{k}{m} = \frac{1}{q}$ for each $j \in [q]$, and $\pi(\succ_i) = 0$ for all $\succ_i \in \Pi^m_> \setminus \bigcup_{j\in[q]} P_j$. We can check that π is a valid distribution because $\sum_{\succ_i \in \Pi^m_>} \pi(\succ_i) = \sum_{j\in[q]} \sum_{\succ_i \in P_j} \pi(\succ_i) = q \cdot \frac{k}{m} = 1$.

In the k-approval rule, each candidate gains one point per preference order where it is ranked among the first k candidates. Under the described preference distribution π, it occurs for candidate x_ℓ with a positive probability only in preference orders in P_j with the unique j such that $x_\ell \in X_j$. It follows that the expected score of each candidate x_ℓ is equal to $\sum_{\succ_i \in P_j : x_\ell \in X_j} \pi(\succ_i) \cdot 1 = \frac{k}{m}$.

Let us now assume that k does not divide m. Let us denote by q and r the unique integers such that $m = k \cdot q + r$ with $0 < r < k$. Suppose, for the sake of contradiction, that there exists a single-peaked distribution π unbiased with respect to the k-approval rule. We will prove by induction that a preference order ranking candidate $x_{(j-1)k+\ell}$ among the first k candidates, for $\ell \in [k]$, can be assigned a positive probability in π only if all the k candidates $x_{(j-1)k+1}$, …, x_{jk} are ranked among the first k candidates in this preference order, for every $j \in [q]$. For the base case, candidate x_1 gets one point under the k-approval rule iff it is ranked among the first k candidates. However, if x_1 is ranked among the first k candidates then, by single-peakedness, it must also be the case of all candidates x_j for $1 < j \leq k$. Since the expected score of x_1 must be the same as the one of all candidates x_j for $1 < j \leq k$, then no positive probability can be assigned to other preference orders where some candidate x_j, for $1 < j \leq k$, is ranked among the first k candidates. We now assume that a preference order ranking candidate $x_{(j'-1)k+\ell}$ among the first k candidates, for $\ell \in [k]$, can be assigned a positive probability in π only if all the k candidates $x_{(j'-1)k+1}, \dots, x_{j'k}$ are ranked among the first k candidates in this preference order, for every $1 \leq j' < j$, for a given $j \in [q]$. It follows that candidate $x_{(j-1)q+1}$ cannot be ranked within the top k of a preference order with positive probability where some candidate $x_{\ell'}$, for $\ell' < (j-1)q+1$, is also ranked within the top k. Therefore, if $x_{(j-1)q+1}$ is ranked within the top k of a preference order with positive probability, then it must also be the case of all the candidates $x_{(j-1)k+2}, \dots, x_{jk}$. Since the expected score of $x_{(j-1)q+1}$ must be the same as the one of all candidates $x_{(j-1)k+2}, \dots, x_{jk}$, then no positive probability can be assigned to other preference orders where some candidate among $x_{(j-1)k+2}, \dots, x_{jk}$, is ranked among the first k candidates, proving the claim.

Now, let us analyze the case of candidate x_m. If x_m is ranked among the first k candidates then, by single-peakedness, it must also be the case of all the $k-1$ candidates x_j, for $m-k+1 \leq j < m$. Since k does not divide m, there exist integers $j \in [q]$ and $\ell \in [k]$ such that $m-k+1 = (j-1)k+\ell$ and thus candidate x_{m-k+1} is approved in single-peaked orders approving candidates $(j-1)k+\ell'$, for $\ell' \in [k]$, and in the disjoint ones approving candidate x_m, therefore its expected score would be equal to the sum of the expected score of x_m and the expected score of x_{m-k}, contradicting the fact that π is unbiased. □

From Theorem 24, no single-peaked distribution can be unbiased with respect to k-approval, for any $k > m/2$ when $m > 2$, which includes the veto rule (i.e., $(m-1)$-approval). Alternatively, there exists a family of single-peaked distributions which are unbiased with respect to the plurality rule (i.e., 1-approval), including Conitzer's distribution. The description of the family of distributions satisfying this property is given in the proof of the theorem: for each candidate x, the distribution must assign a global sum of probabilities of $\frac{1}{m}$ for all single-peaked preference orders ranking x first. In addition, Conitzer's distribution is unbiased only with respect to plurality, leading to the following statement.

Proposition 25. *Conitzer's distribution is unbiased with respect to a positional scoring rule $\mathcal{F}$ iff $\mathcal{F}$ is the plurality rule.*

Proof. By Proposition 20, all candidates are expected winners of the 1-approval rule (i.e., plurality) under Conitzer's distribution. Therefore, Conitzer's distribution is unbiased with respect to plurality.

Suppose that Conitzer's distribution is unbiased with respect to some positional scoring rule $\mathcal{F}$ defined by the positional score vector $\alpha = (\alpha_1, \dots, \alpha_m)$ such that, by definition, $\alpha_1 \geq \alpha_2 \geq \cdots \geq \alpha_m$ and $\alpha_1 > \alpha_m$. It follows that all candidates are expected winners of $\mathcal{F}$, i.e., $\mathbb{E}_{\pi_C}[S^{\mathcal{F}}(x_i)] = \mathbb{E}_{\pi_C}[S^{\mathcal{F}}(x_j)]$, for every $i, j \in [m]$. By Lemma 19, the expected score of a candidate x_j is the following:

$$\begin{aligned}\mathbb{E}_{\pi_C}[S^{\mathcal{F}}(x_j)] &= \sum_{k=1}^{m} \mathbb{P}_{\pi_C}(j,k) \cdot \alpha_k \\ &= \sum_{k=1}^{m} (Q(j,k) + Q(m-j+1,k)) \cdot \alpha_k \\ &= \sum_{k=1}^{j-1} \frac{1}{2m} \cdot \alpha_k + \frac{j}{2m} \cdot \alpha_j + \\ &\sum_{k=1}^{m-j} \frac{1}{2m} \cdot \alpha_k + \frac{m-j+1}{2m} \cdot \alpha_{m-j+1}\end{aligned}$$

By considering, in particular, candidates x_1 and x_2, we have $\mathbb{E}_{\pi_C}[S^{\mathcal{F}}(x_1)] = \frac{1}{2m}\alpha_1 + \frac{1}{2m}\sum_{k=1}^{m-1}\alpha_k + \frac{1}{2}\alpha_m$ and $\mathbb{E}_{\pi_C}[S^{\mathcal{F}}(x_2)] = \frac{1}{2m}\alpha_1 + \frac{2}{2m}\cdot\alpha_2 + \frac{1}{2m}\sum_{k=1}^{m-2}\alpha_k + \frac{m-1}{2m}\cdot\alpha_{m-1}$. For candidates x_1 and x_2 to be both expected winners, they need to have the same expected score. It follows that:

$$\begin{aligned}&\mathbb{E}_{\pi_C}[S^{\mathcal{F}}(x_1)] = \mathbb{E}_{\pi_C}[S^{\mathcal{F}}(x_2)] \Leftrightarrow \\ &\frac{1}{2m}\alpha_1 + \frac{1}{2m}\sum_{k=1}^{m-1}\alpha_k + \frac{1}{2}\alpha_m = \frac{1}{2m}\alpha_1 + \frac{2}{2m}\cdot\alpha_2 + \\ &\frac{1}{2m}\sum_{k=1}^{m-2}\alpha_k + \frac{m-1}{2m}\cdot\alpha_{m-1} \Leftrightarrow \\ &\frac{1}{2m}\sum_{k=1}^{m-1}\alpha_k + \frac{1}{2}\alpha_m = \frac{2}{2m}\cdot\alpha_2 + \frac{1}{2m}\sum_{k=1}^{m-2}\alpha_k + \frac{m-1}{2m}\cdot\alpha_{m-1} \Leftrightarrow \\ &\frac{1}{2m}\alpha_{m-1} + \frac{1}{2}\alpha_m = \frac{2}{2m}\cdot\alpha_2 + \frac{m-1}{2m}\cdot\alpha_{m-1} \Leftrightarrow \\ &\frac{1}{2}\alpha_m = \frac{2}{2m}\cdot\alpha_2 + \frac{m-2}{2m}\cdot\alpha_{m-1} \Leftrightarrow \\ &\alpha_m = \frac{2}{m}\cdot\alpha_2 + \frac{m-2}{m}\cdot\alpha_{m-1}\end{aligned}$$

Because $\alpha_2 \geq \cdots \geq \alpha_{m-1} \geq \alpha_m$, the fact that $\alpha_m = \frac{2}{m}\cdot\alpha_2 + \frac{m-2}{m}\cdot\alpha_{m-1}$ implies $\alpha_2 = \cdots = \alpha_{m-1} = \alpha_m$. It follows that $\alpha_1 > \alpha_2 = \cdots = \alpha_{m-1} = \alpha_m$, and thus $\mathcal{F}$ corresponds to the plurality rule. □

In contrast, we prove that Walsh's distribution can never be unbiased because, no matter the chosen positional score vector, the expected score of a median candidate will always be strictly greater than the one of an extreme candidate in the single-peaked axis.

Proposition 26. *No PSR can make Walsh's distribution unbiased.*

Proof. Suppose, for the sake of contradiction, that Walsh's distribution π_W is unbiased with respect to a given PSR $\mathcal{F}$ characterized by the positional score vector $\alpha = (\alpha_1, \ldots, \alpha_m)$. We can assume, w.l.o.g., that $\alpha_1 = 1$, $\alpha_m = 0$, and $\alpha_y \in [0,1]$ for every $1 < j < m$. By definition, for every candidates x and y, we have $\mathbb{E}_{\pi_W}[S^{\mathcal{F}}(x)] = \mathbb{E}_{\pi_W}[S^{\mathcal{F}}(y)]$, i.e., $\sum_{\succ_i \in \Pi^m} \pi_W(\succ_i) \cdot \alpha_{r_{\succ_i}(x)} = \sum_{\succ_i \in \Pi^m} \pi_W(\succ_i) \cdot \alpha_{r_{\succ_i}(y)}$, and thus $\sum_{\succ_i \in \Pi^m} \frac{1}{2^{m-1}} \cdot \alpha_{r_{\succ_i}(x)} = \sum_{\succ_i \in \Pi^m} \frac{1}{2^{m-1}} \cdot \alpha_{r_{\succ_i}(y)}$ which implies $\sum_{\succ_i \in \Pi^m} \alpha_{r_{\succ_i}(x)} = \sum_{\succ_i \in \Pi^m} \alpha_{r_{\succ_i}(y)}$.

Consider the extreme candidate x_1 and the median candidate $x_c := x_{\lceil \frac{m}{2} \rceil}$. We must have $\sum_{\succ_i \in \Pi^m} \alpha_{r_{\succ_i}(x_1)} = \sum_{\succ_i \in \Pi^m} \alpha_{r_{\succ_i}(c)}$. By Observation 10, candidate c can never be ranked at a position worse than $\gamma := \lfloor \frac{m}{2} \rfloor + 1$, and thus we have $\sum_{\succ_i \in \Pi^m} \alpha_{r_{\succ_i}(c)} = \sum_{k=1}^{m} \mathscr{D}_m(x_c, k) \cdot \alpha_k = \sum_{k=1}^{\gamma} \mathscr{D}_m(x_c, k) \cdot \alpha_k$ where $\sum_{k=1}^{\gamma} \mathscr{D}_m(x_c, k) = 2^{m-1}$. Since x_1 is an extreme candidate, it is ranked last in half of the single-peaked orders. Therefore, by the fact that $\alpha_m = 0$, we have $\sum_{\succ_i \in \Pi^m} \alpha_{r_{\succ_i}(x_1)} = \sum_{k=1}^{m} \mathscr{D}_m(x_1, k) \cdot \alpha_k = \sum_{k=1}^{m-1} \mathscr{D}_m(x_1, k) \cdot \alpha_k$ where $\sum_{k=1}^{m-1} \mathscr{D}_m(x_1, k) \cdot \alpha_k = 2^{m-2}$. Let us now analyze the difference between $\sum_{\succ_i \in \Pi^m} \alpha_{r_{\succ_i}(x_c)}$ and $\sum_{\succ_i \in \Pi^m} \alpha_{r_{\succ_i}(x_1)}$:

$$
\begin{aligned}
& \sum_{\succ_i \in \Pi^m} \alpha_{r_{\succ_i}(x_c)} - \sum_{\succ_i \in \Pi^m} \alpha_{r_{\succ_i}(x_1)} \\
=& \sum_{k=1}^{\gamma} \mathscr{D}_m(x_c, k) \cdot \alpha_k - \sum_{k=1}^{m-1} \mathscr{D}_m(x_1, k) \cdot \alpha_k \\
=& \sum_{k=1}^{\gamma} (\mathscr{D}_m(x_c, k) - \mathscr{D}_m(x_1, k)) \cdot \alpha_k - \sum_{k=\gamma+1}^{m-1} \mathscr{D}_m(x_1, k) \cdot \alpha_k
\end{aligned}
$$

However, by Lemma 11, we have $\mathscr{D}_m(x_c, k) = 2^{k-2}(\binom{m-k}{c-1} + \binom{m-k}{c-k}) \geq 2^{k-2}$ for $k \in \{2, \ldots, \gamma\}$, while $\mathscr{D}_m(x_1, k) = 2^{k-2}(\binom{m-k}{0} + \binom{m-k}{1-k}) = 2^{k-2}$ for $k \in \{2, \ldots, \gamma\}$, and $\mathscr{D}_m(x_c, 1) = \binom{m-1}{c-1}$ and $\mathscr{D}_m(x_1, 1) = 1$. Therefore, $\mathscr{D}_m(x_c, k) - \mathscr{D}_m(x_1, k) \geq 0$ for every $k \in [\gamma]$.

Since $\alpha_1 \geq \alpha_2 \geq \dots \alpha_m$, it follows that:

$$\begin{aligned}
&\sum_{k=1}^{\gamma}(\mathscr{D}_m(x_c,k) - \mathscr{D}_m(x_1,k)) \cdot \alpha_k - \sum_{k=\gamma+1}^{m-1} \mathscr{D}_m(x_1,k) \cdot \alpha_k \\
\geq &\sum_{k=1}^{\gamma}(\mathscr{D}_m(x_c,k) - \mathscr{D}_m(x_1,k)) \cdot \alpha_\gamma - \sum_{k=\gamma+1}^{m-1} \mathscr{D}_m(x_1,k) \cdot \alpha_k \\
= &(2^{m-1} - (1 + \sum_{k=2}^{\lfloor m/2 \rfloor+1} 2^{k-2}) \cdot \alpha_\gamma - \sum_{k=\gamma+1}^{m-1} \mathscr{D}_m(x_1,k) \cdot \alpha_k \\
= &(2^{m-1} - 2^{\lfloor m/2 \rfloor}) \cdot \alpha_\gamma - \sum_{k=\gamma+1}^{m-1} \mathscr{D}_m(x_1,k) \cdot \alpha_k \\
\geq &(2^{m-1} - 2^{\lfloor m/2 \rfloor}) \cdot \alpha_\gamma - \sum_{k=\gamma+1}^{m-1} \mathscr{D}_m(x_1,k) \cdot \alpha_\gamma \\
= &(2^{m-1} - 2^{\lfloor m/2 \rfloor}) \cdot \alpha_\gamma - (2^{m-2} - 2^{\lfloor m/2 \rfloor}) \cdot \alpha_\gamma \\
> &0
\end{aligned}$$

Hence, we always have $\sum_{\succ_i \in \Pi^m} \alpha_{r_{\succ_i}(x_c)} > \sum_{\succ_i \in \Pi^m} \alpha_{r_{\succ_i}(x_1)}$, no matter the chosen positional score vector, a contradiction. □

We now turn our attention to the Borda rule and aim to determine under which distribution it remains unbiased. In doing so, we encounter a highly degenerate distribution. Specifically, we consider a distribution that assigns equal positive probability only to the two extreme rankings within the single-peaked domain.

Definition 12 (Polarized distribution)**.** *The polarized single-peaked distribution* $\pi : \Pi^m_{>} \to [0,1]$ *is defined as:*

$$\pi(\succ_i) = \begin{cases} \frac{1}{2} & \text{if } x_1 \succ_i \cdots \succ_i x_m \text{ or } x_m \succ_i \cdots \succ_i x_1 \\ 0 & \text{otherwise} \end{cases}.$$

Although it is degenerate, the polarized distribution is nevertheless symmetric and is the only single-peaked distribution which is unbiased with respect to the Borda rule.

Theorem 27. *A single-peaked distribution is unbiased with respect to the Borda rule iff it is the polarized distribution.*

Proof. The Borda rule is characterized by, e.g., the positional score vector $(m-1, m-2, \dots, 1, 0)$. Under the polarized distribution, each candidate x_j can be ranked either at position j or at position $m-j+1$, with equal probability. It follows that the expected score of each candidate x_j is equal to:

$$\frac{1}{2}(m-j) + \frac{1}{2}(j-1) = \frac{1}{2}(m-1)$$

Therefore, the polarized distribution is unbiased with respect to the Borda rule.

Let us now prove that no other distribution is unbiased with respect to the Borda rule. Suppose that there exists a single-peaked distribution π which is unbiased with respect to the Borda rule. Observe that, globally, all the Borda scores that have been distributed to the candidates are equal to:

$$\sum_{\succ_i \in \Pi_>^m} \pi(\succ_i) \cdot \sum_{x \in M} (m - r_{\succ_i}(x)) = \sum_{\succ_i \in \Pi_>^m} \pi(\succ_i) \cdot \frac{m(m-1)}{2} = \frac{m(m-1)}{2}$$

Therefore, since all m candidates must have the same expected score, it must be equal to $\frac{m-1}{2}$. Let us denote by $\Pi_>^m(1)$ and $\Pi_>^m(m)$ the set of single-peaked orders where candidate x_1 and x_m are ranked last, respectively. We have $\Pi_>^m = \Pi_>^m(1) \sqcup \Pi_>^m(m)$. Candidates x_1 and x_m get zero points in $\Pi_>^m(1)$ and $\Pi_>^m(m)$, respectively. Since the maximum number of points to get is $(m - 1)$, for x_1 and x_m to get an expected score of $\frac{m-1}{2}$, the distribution should be balanced between $\Pi_>^m(1)$ and $\Pi_>^m(m)$, i.e., we must have $\sum_{\succ_i \in \Pi_>^m(1)} \pi(\succ_i) = \sum_{\succ_i \in \Pi_>^m(m)} \pi(\succ_i) = \frac{1}{2}$. Moreover, for x_1 and x_m to reach an expected score of exactly $\frac{m-1}{2}$ on only half of the single-peaked orders, they must get $m-1$ points, i.e., be ranked at the first position, in the orders with positive probability in their half. Since both x_1 and x_m are ranked first in exactly one single-peaked order, i.e., in the extreme orders $x_1 \succ_i x_2 \succ_i \cdots \succ_i x_m$ and $x_m \succ_i \cdots \succ_i x_2 \succ_i x_1$, respectively, π must assign positive equal probability to exactly these two orders, leading to π being the polarized distribution. $\square$

We have addressed the question of how to sample single-peaked preference in a manner that is unbiased with respect to the candidates. We now continue our work with the question of agreement among voting rules when preferences are drawn from other distributions.

3.7 . Agreement under Other Structured Distributions

Finally, we explore structured preference distributions other than single-peaked ones in order to determine whether similar results can be reached. In particular, we study unimodal distributions, including the famous Mallows' distributions [Mallows, 1957], introduced in voting theory by Goldsmith et al. [2014] and recall in Definition 9, and Pólya-Eggenberger urn [Eggenberger and Pólya, 1923] introduced in voting theory by Berg [1985] and recall in Definition 10.

3.7.1 . Unimodal Distributions

The frequency of a preference order $\succ_i \in \Pi^m$ in a preference profile $\mathscr{P} \in (\Pi^m)^n$ is denoted by $f(\succ_i, \mathscr{P})$. A preference profile $\mathscr{P} \in (\Pi^m)^n$ is *unimodal* [Chatterjee and Storcken, 2020] if there exists a mode $\succ^* \in \Pi^m$, i.e., a reference preference order, such that $f(\succ_i, \mathscr{P}) > f(\succ_j, \mathscr{P})$ iff $dist_{KT}(\succ^*, \succ_i) < dist_{KT}(\succ^*, \succ_j)$, for every pair of preference orders $\succ_i, \succ_j \in \mathscr{P}$. *Positively discriminating* rules [Chatterjee and Storcken, 2020] are social welfare functions which always return the mode as the outcome of the election. Both PSRs and Condorcet-consistent rules are positively discriminating.

We adapt the definition of unimodal profile to distributions. A preference distribution $\pi : \Pi^m \to [0, 1]$ is said to be unimodal if there exists a mode $\succ^* \in \Pi^m$ such that $\pi(\succ_i) > \pi(\succ_i')$ iff $dist_{KT}(\succ^*, \succ_i) < dist_{KT}(\succ^*, \succ_i')$, for every pair of preference orders $\succ_i, \succ_i' \in \Pi^m$. We consider independent and identical voter preference drawings. By using the Glivenko-Cantelli theorem [Cantelli, 1935], we deduce that any unimodal distribution will asymptotically generate a unimodal profile, where PSRs and Condorcet-consistent rules agree to select the winner of the mode.

Corollary 28. *Under unimodal distributions, all PSRs and Condorcet-consistent rules asymptotically agree to elect the first-ranked candidate of the mode.*

We go further and give a bound for the speed of convergence toward agreement in terms of election size. Precisely, we would like to know if this convergence can happen in practice or if this is just for theoretical purpose. We first recall the Dvoretzky-Kiefer-Wolfowitz's (DKW) lemma [Dvoretzky et al., 1956].

Lemma 29 (DKW inequality)**.** *Let $X_1, ..., X_n$ be some independent and identical random variables distributed with a law F. Let $F_n(x, \omega) = \frac{1}{n}\sum_{i=1}^{n} \mathbb{1}_{\{X_i(\omega) \leq x\}}$ then $\mathbb{P}(\sup_{x \in \mathbb{R}} |F_n(x) - F(x)| > \varepsilon) \leq 2e^{-2n\varepsilon^2}, \quad \forall \varepsilon > 0.$*

Proposition 30. *For a unimodal preference distribution π, the probability that all PSRs and Condorcet-consistent rules agree is lower bounded by $B_\pi := 1 - 2 \cdot exp(-2n\varepsilon^2)$, for $\varepsilon := \min_{\succ_i, \succ_j \in \Pi^m} |\pi(\succ_i) - \pi(\succ_j)|$.*

Proof. We remark that for ε sufficiently small, i.e., $\varepsilon = \min_{\{\succ_i, \succ_i'\}} |\mathbb{P}_\pi(\succ_i) - \mathbb{P}_\pi(\succ_i')|$, we have $\{||F_n(., \omega) - F||_\infty \leq \varepsilon\} = \{F_n \text{ is unimodal}\}$. Applying Lemma 29 on the contrary event and using theorem 4.1 from Chatterjee and Storcken [2020], we get that the described voting rules agree with probability at least $1 - 2e^{-2n\varepsilon^2}$. □

A typical example of unimodal distributions are Mallows' distributions when $\phi < 1$.

We give below an example of the speed of convergence under Mallows' distributions.

Example 21. *Under a Mallows' distribution $\pi^{\phi,\sigma}$, we get $\varepsilon = \phi^k \cdot (1 - \phi)$ with $k := \max_{\succ} dist_{KT}(\sigma, \succ_i)$ and thus the bound for agreement is $B_{\pi^{\phi,\sigma}} = 1 - 2 \cdot exp(-2n(\frac{\phi^k \cdot (1-\phi)}{Z})^2)$.*

If we consider an election with three candidates, for instance when $\phi = 0.1$ then we need $n = 2,000,000$, (then $k = 3$) to reach $B_{\pi^{\phi,\sigma}} = 0.92$. If $\phi = 0.9$ then $n = 400$ is enough to get $B_{\pi^{\phi,\sigma}} = 0.97$. When more weight is given to orders close to the mode, voting rules agree faster than when the Mallows' distribution gets closer to impartial culture (i.e., $\phi = 1$). This is quite understandable, as voting rules tend to agree more quickly when the culture is concentrated around a single preference. Conversely, when preferences are more dispersed, convergence takes longer to occur. This is also coherent with the impartial culture (i.e., $\phi = 1$) which have a positive probability of disagreement.

3.7.2 . Pólya-Eggenberger Urn

This subsection is dedicated to another structured distribution, namely the Pólya-Eggenberger urn model. The idea is to generate the preference profile using a reinforcement mechanism: we start with an urn containing each possible preference order with equal probability, and after each draw, the selected preference is returned to the urn along with R additional copies of itself. The following result generalizes the asymptotic result from Gehrlein [2002] for three candidates under impartial anonymous culture (when $R = 1$).

For the purpose of the next proposition we need to introduce Dirichlet law which will help us describe the asymptotic limit of the law of the score under Pólya-Eggenberger Urn drawings.

Definition 13 (Dirichlet law)**.** *Let $d \geq 2$ be an integer. Let Σ be the $(d-1)$-dimensional simplex*

$$\Sigma = \left\{ (x_1, \ldots, x_d) \in [0,1]^d \mid \sum_{k=1}^{d} x_k = 1 \right\}$$

then

$$f(x_1, \ldots, x_d)\, d\Sigma(x_1, \ldots, x_d)$$
$$= f\left(x_1, \ldots, x_{d-1}, 1 - \sum_{k=1}^{d-1} x_k \right) \mathbb{1}_{\{x \in [0,1]^{d-1}, \sum_{k=1}^{d-1} x_k \leq 1\}}\, dx_1 \cdots dx_{d-1}$$

for any continuous function f.

Here is the lemma describing the desired convergence.

Lemma 31 (Asymptotic convergence of Pólya-Eggenberger urn [Athreya, 1969])**.** *Let $d \geq 2$ and $R \geq 1$ be an integer. Let also $\beta = (\beta_1, \ldots, \beta_d) \in \mathbb{N}^d \setminus \{0\}$. Let $(P_n)_{n\geq 0}$ be the d-color Pólya-Eggenberger urn random process having R as reinforcement parameter and β as initial composition. Then, almost surely and in any L^t, $t \geq 1$,*

$$\frac{P_n}{nR} \xrightarrow[n\to\infty]{} V$$

where V is a d-dimensional Dirichlet-distributed random vector, with parameters $\left(\frac{\beta_1}{R}, \ldots, \frac{\beta_d}{R}\right)$.

Remark 32. *Let us recall that the convergence in L^t, $t \geq 1$ implies the convergence in law.*

We now have the necessary tools to analyze the asymptotic behavior of this distribution.

Proposition 33. *When the election is drawn with a Pólya-Eggenberger urn culture (named P-E) with $R = m! \cdot r$, the probability that all PSRs asymptotically agree is lower bounded by $\frac{1}{2}$ if $r < \frac{2}{3}$ and $m = 3$, and by $\frac{1}{4}$ if $r < \frac{1}{6}$ and $m = 4$.*

Proof. Let us recall (see Lemma 31) that a Pólya-Eggenberger urn asymptotically converges to the Dirichlet law (see Definition 13). Thus, we can calculate the probability that a specific distribution of preferences occurs.

We now need to describe the event D_m where all positional scoring rules agree. We use the known fact that all positional scoring rules will agree if all k-approval voting rules agree with each other [Saari, 2012] and get for example for $m = 3$: $D_3 = \{(p_1, p_2, p_3, p_4, p_5, p_6) \in \Sigma \mid p_1 + p_2 > p_3 + p_4, p_1 + p_2 > p_5 + p_6, p_2 + p_5 > p_4 + p_6, p_1 + p_3 > p_4 + p_6\}$ using the following notation $(p_1, p_2, p_3, p_4, p_5, p_6)$ for the proportion of each preference in the election in the following order $(x_1 \succ x_2 \succ x_3), (x_1 \succ x_3 \succ x_2), (x_2 \succ x_1 \succ x_3), (x_2 \succ x_3 \succ x_1), (x_3 \succ x_1 \succ x_2), (x_3 \succ x_2 \succ x_1)$.

We now come back to our initial question which is to compute $\lim_{n\to\infty} \mathbb{P}_{P-E}(D_3)$. Using Lemma 31, we are able to identify the limit law and to compute $\mathbb{P}_V(D)$, for every $0 < R \leqslant 4$. Since the analytical is fastidious, we use the Monte-Carlo method with a very high precision ($n = 10{,}000{,}000$) to compute the integral and get the desired result. We recover the result on r by doing the change of variable $R = m! \cdot r$. We follow the exact same steps for $m = 4$, with the only difference being the constraint used to define D_4. The probability of D_4 is again estimated using a Monte Carlo method with very high precision ($n = 10{,}000{,}000$). □

We now specifically examine the agreement between plurality and Borda rule.

Proposition 34. *When the election is drawn with a Pólya-Eggenberger urn culture with $R = m! \cdot r$, the probability that plurality and Borda asymptotically agree is lower bounded by $\frac{3}{4}$ if $r < \frac{2}{3}$ and $m = 3$, and by $\frac{3}{5}$ if $r < \frac{1}{6}$ and $m = 4$.*

Proof. We follow the exact same steps as in the previous proof but we need to construct a different space to find where Plurality and Borda agree. For example for $m = 3$, $D_3 = \{(p_1, p_2, p_3, p_4, p_5, p_6) \in \Sigma \mid p_1 + p_2 > p_3 + p_4, p_1 + p_2 > p_5 + p_6, p_1 + 2 \cdot p_2 + p_5 > p_3 + 2 \cdot p_4 + p_6, 2 \cdot p_1 + p_2 + p_3 > p_4 + p_5 + 2 \cdot p_6\}$ using the following notation $(p_1, p_2, p_3, p_4, p_5, p_6)$ for the proportion of each preference in the election in the following order $(x_1 \succ x_2 \succ x_3), (x_1 \succ x_3 \succ x_2), (x_2 \succ x_1 \succ x_3), (x_2 \succ x_3 \succ x_1), (x_3 \succ x_1 \succ x_2), (x_3 \succ x_2 \succ x_1)$. We again compute the probabilities of D_3 and D_4 using a Monte Carlo method with very high precision ($n = 10{,}000{,}000$). $\square$

To give a comparison, we describe a small example to compare Pólya-Eggenberger urn and Walsh's distribution.

Example 22. *For the agreement plurality and the Borda rule to the election of median candidates C^* under the Walsh's distribution, we have a lower bound given by the plurality rule $\mathcal{F}_1$ (by Theorem 18) which is as follows: if $m = 4$, $\mathbb{P}_{\pi_W}(\mathcal{F}_1(\succ) = C^*) \geqslant 1 - 2e^{-\frac{n}{32}}$ is larger than $\frac{3}{5}$ when $n \geqslant 52$.*

Therefore, we can compare the lower bounds and observe that the lower bound of the Pólya-Eggenberger urn model for $r < \frac{1}{6}$ reaches similar value as that of Walsh's distribution from $n \geq 52$. This implies that, when the number of voters exceeds 52, the lower bounds become comparable.

We finally prove a positive probability of disagreement asymptotically for every pair of PSRs.

Proposition 35. *If the election is drawn with a Pólya-Eggenberger urn culture with $R < 4$ then every pair of positional scoring rules $\mathcal{F}_1$ and $\mathcal{F}_2$ asymptotically disagree with a positive probability, i.e., $\lim_{n\to\infty} \mathbb{P}(\mathcal{F}_1(\mathscr{P}) \neq \mathcal{F}_2(\mathscr{P})) > 0$.*

Proof. Let $\mathcal{F}_1, \mathcal{F}_2$ be two positional scoring rules. There exist two positional score vectors α^1 and α^2 corresponding to these two rules. Since $\mathcal{F}_1$ and $\mathcal{F}_2$ are different, α^1 and α^2 differ on at least one component, i.e., there exists $i \in [m]$ such that $\alpha_i^1 \neq \alpha_i^2$. Let us denote $\varepsilon = \alpha_i^1 - \alpha_i^2 > 0$. We will show that there exists a profile $\mathscr{P}$ such that $\lim_{n\to\infty} \mathbb{P}(\mathcal{F}_1(\mathscr{P}) \neq \mathcal{F}_2(\mathscr{P})) > 0$. Specifically, to build such a profile we consider an arbitrary profile such that the j^{th} candidate has an asymptotic score of 0, then we slowly increase the proportion of one preference such that candidate j is ranked in position i until $\mathcal{F}_1(\mathscr{P}) \neq \mathcal{F}_2(\mathscr{P})$. By doing so, we find that we can still increase this proportion from $\delta < \varepsilon$ and keep the disagreement between the two rules. Thus, there exists a non negligible set where $\mathcal{F}_1(\mathscr{P}) \neq \mathcal{F}_2(\mathscr{P})$. Finally, we identify the limit law of a Pólya-Eggenberger urn as the Dirichlet random variable thanks to Lemma 31 and conclude that $\lim_{n\to\infty} \mathbb{P}(\mathcal{F}_1(\mathscr{P}) \neq \mathcal{F}_2(\mathscr{P})) > 0$ because this is a continuous density on a non negligible set. $\square$

This result means that any pair of positional scoring rules will disagree on a nonempty set asymptotically. Thus, we cannot achieve the same type of convergence results as in single-peaked distributions.

3.8 . Agreement of Voting Rules in Practice

This section focuses on real-world electoral data. Given the natural relevance of this question, it has already been addressed in several studies, including those on American elections [Regenwetter et al., 2007], Romanian elections [Roescu, 2014], and the parliamentary elections in the Austrian federal state of Styria [Darmann et al., 2019]. However, the last one is the most closely related to our work as it analyzes the outcomes of various voting rules, often the same as ours, on real-world data. The study, based on data from the parliamentary elections in the Austrian federal state of Styria and using different bootstrap settings, reports complete agreement between Borda and plurality, almost complete agreement between plurality and the two-round system, and more disagreement with the veto rule.

Another interesting work carried out in France, known as Voter autrement [Bouveret et al., 2018], notably aims to understand the variability induced by the choice of voting rules through data collection. We focus on data from the 2017 presidential election for technical reasons: it is a single-winner election, and we have access to complete rankings, which allows us to avoid making additional assumptions for our experiments.

The authors of Voter Autrement have already worked on the question of agreement among different voting rules with their data. However, they focused on plurality with runoff, (the actual rule used), approval voting, and various graded scales (e.g., $(0, 1, 2)$ or $(-1, 0, +1)$). Instead, we will conduct a similar study focusing on positional scoring rules, as this experiment is intended to illustrate our previous theoretical results.

We will specifically focus our experiment on positional scoring rules and Condorcet-consistent rules, when a Condorcet winner exists.

In the "Voter autrement" dataset, only one file was suitable for our purposes, namely `stv111.csv` with 11 candidates, because it contains 4,068 complete preference orders. This choice has the significant advantage of requiring no additional assumptions about the data, as incomplete data would otherwise necessitate hypotheses to fill in the missing information on the preference profile. However, it may introduce a bias, as we simply discard all incomplete preference orders. For our purpose of comparing voting rules with one another, we assume that this does not interfere with the interpretation of the results. To preserve anonymity, we relabel the candidates using the letters A to K.

Before interpreting the results related to single-peaked distributions in light of real-world data, we first aim to assess whether the data itself satisfies the single-peaked property. To this end, we follow the approach introduced by Faliszewski et al. [2011], Sui et al. [2013] and Elkind and Lackner [2014]; later completed by esc [2021] and further developed in the thesis of Tydrichová [2023]. Indeed, we will rely on the *forbidden triples* measure, as it is the only one that remains computationally feasible for profiles with 11 candidates and 4068 voters. This measure builds upon the definition in Definition 1 , by counting the number of candidate triples that violate the single-peakedness condition with respect to a given axis. The idea is to quantify how far a given profile is from being single-peaked, and then minimize this deviation.

Another commonly used measure of single-peakedness is *Global Swaps (GS)*, introduced by Erdélyi et al. [2013], which minimizes the number of swaps of consecutive candidates required to transform the preference profile into a single-peaked one. Even though both problems are NP-hard, experiments from esc [2021], as well as our own, show that FT is significantly more efficient: for profiles with more than six candidates, the computation remains feasible with FT, whereas it becomes prohibitively slow with GS.

Following the work of Tydrichová [2023], we compute a single-peaked axis by minimizing this measure and inverting the axis, we obtain the following result: “François Asselineau" < “Jacques Cheminade" < “Nathalie Arthaud" < “Philippe Poutou" < “Jean-Luc Mélenchon" < “Benoît Hamon" < “Emmanuel Macron" < “Jean Lassalle" < “Nicolas Dupont-Aignan" < “François Fillon" < “Marine Le Pen". This ordering appears broadly consistent for all major candidates, with the exception of “François Asselineau", whose placement may be affected by his limited popularity in the data set.

We obtain a score, that is the number of forbidden triples, equal to 20,213. However, to make this result interpretable, it is necessary to provide a proportional baseline. To that end, we generated profiles under the impartial culture model with $m = 11$ candidates and $n = 4068$ voters, ran the experiment twice , and observed scores of 59,459 and 59,291, which are quite close. These results indicate that the data set deviates only moderately from single-peakedness. The result might be even stronger if we remove candidates with low support. Therefore, this brief illustration supports the relevance of comparing our experimental data to the theoretical results obtained under single-peaked distributions.

Table 3.1 presents the outcomes obtained from the `stv111.csv` data set, considering only complete preference orders. Since we keep only complete preference orders the data set is highly biased but this is not a problem because we want to compare rules with each other. However, we still anonymous for that part on winners because we want to avoid any political misinterpretation on political results. Since this data set admits a Condorcet winner, all Condorcet-consistent rules are expected to agree on her. The table reports the winners under all positional scoring rules, as well as under Borda and the Condorcet winner.

We then apply the same analysis to two other data sets from Voter autrement [Delemazure and Bouveret, 2024], namely `irv_1.csv` and `irv_2.csv`, from the 2022 French presidential election. We merge these two

Voting Rule	**Winner**
plurality	H
Borda	H
Condorcet winner (if one exists)	H
2-approval	H
3-approval	H
4-approval	H
5-approval	H
6-approval	H
7-approval	H
8-approval	C
9-approval	C
10-approval (veto)	C

Table 3.1: Winners according to various voting rules applied to the `stv111.csv` dataset (only complete preferences orders)

data sets, as they are the only ones that contain complete rankings. We then filter the data to retain only complete preference orders. While this step may introduce a bias, it should not significantly affect the axis. Moreover, we aim to follow the same procedure as before and avoid additional assumptions.

Following the same procedure, we compute a single-peaked axis by minimizing this measure and inverting the axis, we obtain the following result: "Nicolas Dupont-Aignan" > "Nathalie Arthaud" > "Philippe Poutou" > "Fabien Roussel" > "Jean-Luc Mélenchon" > "Yannick Jadot" > "Anne Hidalgo" > "Emmanuel Macron" > "Jean Lassalle" > "Valérie Pécresse" > "Marine Le Pen" > "Éric Zemmour". This ordering appears broadly consistent for all major candidates, with the exception of "Nicolas Dupont-Aignan", whose placement may be affected by his limited popularity in the data set.

We obtain a score, that is the number of forbidden triples, equal to 2,457. However, to make this result interpretable, it is necessary to provide a proportional baseline. To that end, we generated profiles under the impartial culture model with $m = 12$ candidates and $n = 412$ voters, ran the experiment twice , and observed scores of 6,967 and 7,039, which are quite close. These results indicate that the data set deviates only moderately from single-peakedness and in similar proportions as in the previous experiment. Therefore, this brief illustration supports the relevance of comparing our experimental data to the theoretical results obtained under single-peaked distributions.

Table 3.2 presents the outcomes obtained from the `irv_1.csv` and `irv_2.csv` data sets, considering only complete preference orders. Similarly to the last experiment, since we keep only complete preference orders the data set is highly biased but this is not a problem because we want to compare rules with each other. However, we still anonymous for that part on winners because we want to avoid any political misinterpretation on political results. Since this data set admits a Condorcet winner, all Condorcet-consistent rules are expected to agree on her. The table reports the winners under all positional scoring rules, as well as under Borda and the Condorcet winner.

Voting Rule	**Winner**
plurality	K
Borda	K
Condorcet winner (if one exists)	K
2-approval	K
3-approval	K
4-approval	K
5-approval	E
6-approval	E
7-approval	E
8-approval	E
9-approval	E
10-approval	E
11-approval (veto)	E

Table 3.2: Winners according to various voting rules applied to `irv_1.csv` and `irv_2.csv` data sets (only complete preferences orders)

These experimental analysis confirms some theoretical findings observed under structured preference cultures, namely a high degree of agreement between major positional scoring rules and Condorcet-consistent rules. Moreover, we recover a phenomenon previously described under Walsh's distribution: there exists a transition point in the level of agreement between scoring rules, depending on how rapidly their score vectors decrease. In particular, voting rules whose score vectors decrease earlier tend to align with each other, as highlighted in Theorem 16. Indeed, as we approach the veto rule, we observe a change in the winner. However, Theorem 16 predicts that this transition should occur at an index $\ell \in [\lfloor \frac{m}{2} \rfloor + 1]$, which would correspond to index 6 for the first experiment and 7 for the second. In our first experiment, the agreement persists up to the 7-approval rule. The scores are close, 3776 for candidate H and 3630 for candidate C, indicating that the difference is small. In the second one, the change of winner occurs earlier, specifically between the 4-approval and 5-approval rules. Nonetheless, we still observe the phenomenon described in Theorem 16, with an alignment of voting rules up to a certain positional rule, followed by a switch to a different winner that persists until the veto rule. However, a key difference lies in the fact that neither candidate H or K are median candidate. This results from the strong political bias present in our data set, which stems from the exclusion of incomplete rankings. Moreover, our data set does not share the property of Walsh's model, which assigns higher probability weights to median candidate(s). Overall, this experiment allows us to partially recover the theoretical results. Further experiments should be conducted using an unbiased data set to determine whether our findings remain consistent. However, having sufficiently large data sets with complete preference orders remains complicated.

3.9 . Conclusion and Future Works

3.9.1 . Conclusion

In this chapter, we examined the variability of outcomes with different voting rules. The previous results on impartial cultures indicate that the probability of disagreement is significant under impartial cultures, we study several different voting culture commonly used for experiments in social choice, where the agreement is a lot higher. Specifically, we studied two single-peaked cultures, namely Walsh's and Conitzer's distributions. Additionally, we provided some results on the Mallows model and partially describe the outcomes under the Pólya-Eggenberger urn model. Walsh's and Conitzer's distributions tend to favor the election of median candidate(s) in the single-peaked axis, and these candidates also turn out to be (weak) Condorcet winner(s), implying the agreement of several positional scoring rules (PSRs) with all Condorcet-consistent rules. This (weak) Condorcet efficiency holds in general for all symmetric single-peaked distributions, which are natural distributions for experiments when no additional information other than the single-peaked axis is available. This leads to a fast convergence toward agreement among many voting rules. The observed behavior under Walsh's model aligns closely with our experimental findings. We nevertheless observe that Conitzer's distribution is less biased toward the median candidates because it happens to be unbiased with respect to one PSR (namely plurality), contrary to Walsh's distribution. While these single-peaked distributions enable fast convergence to agreement, this is also the case for other structured distributions, such as unimodal ones, where the agreement is very general among voting rules and convergence is rapid. This behavior cannot be extended to Pólya-Eggenberger urns where the probability of disagreement is non-negligible, even if it remains high in some particular cases. In addition, we ran experiments on real-world data and observed some similarities with the theoretical results.

Our findings highlight that particular attention should be taken when using voting cultures for experiments in social choice. Indeed, since we identify cultures in which the agreement of different voting rules rapidly agree as the number of voters increases, conclusions drawn from experiments testing different voting rules for a problem should be interpreted with caution. One could imagine very different conclusions about a problem, not because of the problem itself, but because of the culture used: impartial cultures versus single-peaked cultures, for example.

3.9.2 . Future Works

This work tells us important features on cultures but many questions related to that problem remain open:

- We could consider bounds on the probability to agree in finite elections with Pólya-Eggenberger urn. The difficulty, however, lies in the dependent structure of this distribution.
- Another promising avenue is on impartial culture. To the best of our knowledge, the vast literature on this question tries to give explicit formulas for the probability of disagreement between voting rules. For more than three candidates, they quickly become uninterpretable or impossible to find. An interesting question is whether we can derive useful approximations, perhaps by applying techniques similar to those used in Theorem 8 that yield practical and meaningful bounds.
- One idea could also be to consider nearly single-peaked distributions to bridge the gap between impartial and single-peaked cultures and be closer to real political elections.
- When voting rules asymptotically agree, we might conjecture that the probability of not satisfying certain axioms might also decrease as the election size increases.
- The same study could be done with strategic voters [Meir, 2018].
- One final question that may be of interest is whether there exists a central rule that maximizes agreement with the various classically accepted voting rules. Indeed, if such a rule exists, it could be seen as minimizing the dilemma of choosing the “right" voting rule.

In this chapter, we took a significant step forward in our effort to understand what influences voting outcomes by examining how election results can vary depending on the voting rule used. However, this analysis was conducted from a static perspective, assuming that the ballots themselves remain unchanged. Indeed, the underlying assumption in this chapter is that all voters express truthfully their preferences. Yet, as shown by Gibbard [1973]; Satterthwaite [1975], strategic voting is inevitable. Our next objective, therefore, is to explore how strategic behavior can affect election results, a topic we will begin to address in Chapter 4.

4 - Strategic Voting and its Consequences on Plurality Voting Outcomes

Abstract

This chapter deals with iterative voting under the plurality rule, where voters can strategically perform sequential deviations. Most works in iterative voting focus on convergence properties or evaluate the quality of the outcome. However, the iterative winner depends on the sequence of voters' deviations. We propose to analyze to what extent this impacts the outcome of iterative voting by adopting a qualitative, quantitative and computational approach. In particular, we introduce the notions of possible and necessary iterative winners. We first study the extreme scenario for the existence of a necessary winner, where no voter has an incentive to deviate from her truthful ballot. We show that this phenomenon occurs with high probability under impartial cultures. Then, we explore the computational complexity of determining possible and necessary iterative winners, proving that the two problems fall in different complexity classes. Finally, we prove that the Condorcet efficiency of plurality is increased by considering its iterative voting version.

Résumé

Ce chapitre traite du vote itératif avec la règle de vote pluralité, où les électeurs peuvent effectuer des déviations séquentielles de manière stratégique. La plupart des travaux sur le vote itératif se concentrent sur les propriétés de convergence ou évaluent la qualité du résultat. Cependant, le gagnant itératif dépend de la séquence des déviations des électeurs. Nous proposons d'analyser son impact sur le résultat du vote itératif en adoptant une approche qualitative, quantitative et computationnelle. En particulier, nous introduisons les notions de gagnants itératifs possibles et nécessaires. Nous étudions d'abord le scénario extrême où aucun électeur n'est incité à changer son bulletin sincère, garantissant l'existence d'un gagnant nécessaire. Nous montrons que ce phénomène se produit avec une forte probabilité dans le cas des cultures impartiales. Ensuite, nous explorons la complexité algorithmique de la détermination des gagnants itératifs possibles et nécessaires, en prouvant que les deux problèmes appartiennent à des classes de complexité différentes. Enfin, nous prouvons théoriquement que l'efficacité de Condorcet augmente avec les mouvements stratégiques.

Most of the content of this chapter is based on a paper co-authored with Vincent Mousseau, Magdalena Tydrichová, and Anaëlle Wilczynski, which was accepted at the 10th Workshop on Computational Social Choice (COMSOC 2025) [Mousseau et al., 2025a].

4.1 . Introduction

Our analysis of the variability of voting outcomes continues in this second chapter, which focuses on the impact of strategic voting in Plurality elections. As established in Theorem 2, no voting rule can fully prevent strategic manipulation by voters. Instead, various models allow such manipulation and analyze the resulting outcomes [Meir, 2018]. Iterative voting [Meir, 2017] is a particular voting game where voters are allowed to manipulate by performing successive moves. However, since voters manipulate sequentially, different possible outcomes can arise, depending on which voters' deviations are chosen. A natural question is thus to know which candidates turn out to be winners, for some sequence of deviations, once convergence is reached. We propose to answer this question by adapting the well-known notions of possible and necessary winners [Konczak and Lang, 2005] to the iterative context. More precisely, a possible iterative winner is a candidate for which there exists a sequence of deviations eventually electing her at equilibrium. Analogously, a necessary iterative winner is elected in all possible equilibria that can be reached by the iterative voting process.

In this chapter, we follow the classical initial model of Meir et al. [2010] where the voting rule is plurality, and voters perform direct best responses when they are pivotal, i.e., they vote for the candidate they prefer the most among those they can make the new winner. Under these assumptions, the iterative voting process is guaranteed to converge to a situation of equilibrium. In such a setting, we analyze the iterative voting outcomes quantitatively and qualitatively, and in particular the problems of possible and necessary iterative winner.

For the necessary iterative winner problem, we have two possible scenarios: either there are several deviation sequences but they all eventually elect the same candidate, or a more extreme scenario occurs where no voter can deviate from her truthful ballot and thus the initial winner turns out to be the only iterative winner. We propose to quantify the occurrence of such a phenomenon by analyzing how frequently a preference profile is already an equilibrium. We show, under impartial (anonymous) cultures, a rather high lower bound for the probability of this extreme scenario, and thus for the probability of the existence of a necessary iterative winner.

In general, we investigate the computational complexity of the existence problem of a possible/necessary iterative winner. It turns out that the problems fall into different complexity classes since the possible iterative winner problem is NP-complete, while the necessary one is polynomial-time solvable. Finally, we evaluate the quality of iterative voting outcomes by considering the election of the Condorcet winner, when she exists, as an iterative winner. In particular, we theoretically prove that the Condorcet efficiency, (i.e., the probability to elect a Condorcet winner when she exists) increases through the iterative process. More precisely, under impartial (anonymous) cultures, the iterative version of plurality improves the Condorcet efficiency compared to the one of plurality, confirming and generalizing experimental results [Grandi et al., 2013].

Related Work:

The iterative voting model has been introduced by the seminal work of Meir et al. [2010]. Since then, many articles have investigated iterative voting under different voters' strategic behaviors and voting rules (see Meir [2017] for a recent survey).

In this chapter, we introduce the notion of possible and necessary iterative winners, which are adaptations of the well-known concepts of possible and necessary winners under incomplete preferences [Konczak and Lang, 2005], which have been initially introduced to capture which outcomes can arise when information is incomplete. Up to our best knowledge, these notions have not been used so far to capture iterative voting outcomes. Nevertheless, in the context of manipulation in voting, they have been applied, e.g., to deal with incomplete information of the manipulators [Conitzer et al., 2011], as a list of intermediate results in an iterative elicitation process where voters can answer to the queries strategically [Dery et al., 2019], or to determine the outcome of sequential voting in the context of social networks [Gaspers et al., 2014]. Our computational results, stating a difference in complexity classes between the possible and necessary iterative winner problems, are consistent with the results of the literature regarding the initial notions. Notably, while the necessary winner problem under partial preferences is in P for all positional scoring rules, the possible one is NP-complete on this large class of rules except for the plurality and veto rules [Baumeister and Rothe, 2012; Betzler and Dorn, 2010; Konczak and Lang, 2005; Xia and Conitzer, 2011].

Note that considering all possible iterative outcomes that can arise, depending on the sequence of voters' deviations, is similar in spirit to the notion of "parallel-universe" tie-breaking where the outcome is the set of all candidates who could win using a particular tie breaking method. This has been particularly investigated for multi-stage voting rules where the choice of the candidates to eliminate at a given stage can highly impact the final winner of the voting procedure [Freeman et al., 2015; Tideman, 1987]. The different sequences of eliminated candidates at the different stages can then be represented as a tree [Freeman et al., 2015], and we can use a similar representation for all possible sequences of deviations potentially leading to different winners.

Instead of considering the diversity of iterative voting outcomes, where two equilibria are indistinguishable if they elect the same winner, one can focus more specifically on the possible equilibria that can be reached. This study has notably been conducted by Rabinovich et al. [2015], who establish that checking whether a given ballot profile is a reachable equilibrium is NP-hard, in a similar idea as our NP-completeness proof for the possible iterative winner problem. Following a similar high-level approach to identifying all Nash equilibria, the work of Thompson et al. [2013] conducts extensive experiments to explore a wide range of possible outcomes. However, their study differs in spirit from ours, as it considers an alternative response model incorporating utility-based incentives for deviating from the truthful vote. In particular, they focus on a specific framework in which voters may have an incentive to vote truthfully.

In an orthogonal perspective, one can examine how good or bad are the outcomes of iterative voting. In particular, several works have analyzed the iterative voting outcomes by comparing them to the initial truthful one, following either a worst-case analysis based on an approach similar to the price of anarchy, or an average-case analysis [Brânzei et al., 2013; Kavner and Xia, 2021, 2024]. Mostly, the outcomes have been evaluated via their social welfare, but it is also possible to consider other measures, such as the probability to elect the Condorcet winner when she exists [Gehrlein and Lepelley, 2010]. Grandi et al. [2013] have followed this latter approach by experimentally analyzing the Condorcet efficiency of the iterative voting process. We go a bit further by theoretically demonstrating that indeed the iterative variant of plurality has a higher Condorcet efficiency compared to the initial plurality rule, where we consider the probability of electing a Condorcet winner over all possible deviation sequences with equal weights, under impartial (anonymous) cultures.

4.2 . Preliminaries on the Model

In Section 2.5, we already defined the classical iterative voting model that we will use. We now introduce some concepts and properties of this model that will be used in the following sections.

We recall that $PW(s)$ denote the set of potential winners according to a given score vector s and I_n^m the set of all possible candidates' scores under plurality, i.e., $I_n^m := \{s \in \mathbb{N}^m \mid \sum_{j=1}^m s_j = n\}$.

We introduce the following notion to group the scores by the number of potential winners.

Definition 14. *Let S_n^j be the set of all score vectors in an n-voter election such that the union of potential winner sets over all voters contains exactly j candidates, i.e., $S_n^j = \{s \in I_n^m : |PW(s)| = j\}$.*

Note that $(S_n^j)_{j=1}^m$ forms a partition of I_n^m. Especially, S_n^1 corresponds to all score vectors with a unique potential winner. More precisely, for every score vector s in S^1_n, there exists a candidate which is the unique potential winner for all voters, and thus it is the winner in s.

We then consider a best response dynamics which is defined via deviation sequences. We denote by $DS(\mathscr{P})$ the set of all possible deviation sequences for preference profile $\mathscr{P}$ defined as follows:

Definition 15 (Deviation sequence)**.** *A sequence of strategy profiles $\langle b^0, b^1, \ldots, b^r\rangle$ is a deviation sequence for preference profile $\mathscr{P}$ if:*

- *b^0 corresponds to the initial truthful ballot profile $b^\top$,*
- *for every step $t \in [r]$, state b^t results from a best response by exactly one voter from state b^{t-1}, i.e., for every step $t \in [r]$, there exists one voter $i \in N$ and one candidate $y \in PW_i^{t-1} \setminus \{w^{t-1}\}$ such that $y \succ_i z$ for every $z \in PW_i^{t-1}$, where $b_i^t = y$ and $b_j^t = b_j^{t-1}$ for every voter $j \in N \setminus \{i\}$,*
- *the sequence is maximal, i.e., b^r is an equilibrium where no voter has interest to change her ballot.*

A deviation sequence is said to be *empty* if it is restricted to the initial ballot profile $\langle b^0\rangle$ which is already an equilibrium.

We distinguish two types of strategic moves, one from a potential winner (FPW) (i.e., a deviation by voter i at step t from b_i^{t-1} to b_i^t where $b_i^{t-1} = x$ and $x \in PW^{t-1}$) and one from a non potential winner (FNPW) (i.e., a move by voter i at step t from b_i^{t-1} to b_i^t where $b_i^{t-1} = x$ and $x \notin PW^{t-1}$).

From Meir et al. [2010], we have an upper bound on the number of moves before convergence, in plurality iterative voting, which is given by $O(m \cdot n)$. We state below that this bound can be improved.

Proposition 36. *The number of moves in any deviation sequence is in $O(m + n \cdot log(m))$.*

Proof. We identify the worst-case scenario for the number of strategic moves in a deviation sequence by analyzing the worst-case subsequences. Since the number of potential winners can only decrease throughout the process, we define a subsequence of strategic moves as the set of moves that occur while the number of potential winners remains constant. We then examine the worst-case scenario in which the number of strategic deviations is maximized within each such subsequence. We start with a score vector in S^m_n. By Observation 38, the first move yields a score vector in S^{m-1}_n, in which the (unique) non-potential winner y_1 has less than $\frac{n}{m}$ votes. Since each voter i such that $b_i^1 = y_1$ can deviate to one of the $m-1$ remaining potential winners, we have at most $\frac{n}{m}$ FNPW deviations, each yielding, in the worst case, a new score vector in S^{m-1}_n. These are then followed by a FPW deviation that yields a score vector in S^{m-2}_n. We repeat the process–for each $k \in [m]$, when we reach a score of S^{m-k}_n, we have, in the worst case:

- $m-k$ potential winners, each obtaining approximately $\frac{n}{m-(k-1)}$ votes,
- one non-potential winner obtaining less than $\frac{n}{(m-(k-1))}$ votes,
- $k-1$ additional non-potential winners, each receiving zero votes.

We can thus perform at most FNPW moves and one FPW move before the next decrease of the number of potential winners. By Meir [2022], $|PW^t|$ can only decrease with t, so the process will terminate, and we will have at most

$$1+\sum_{k=2}^{m} 1+\frac{n}{m-k} = m+n\cdot\sum_{k=2}^{m}\frac{1}{m-k}$$

$$= m+n\cdot\sum_{l=0}^{m-2}\frac{1}{l} \leqslant m+n\cdot log(m)$$

strategic moves. □

The following example shows the potential diversity of iterative voting outcomes depending on the choice of the deviation sequence. This example uses the same instance as in Example 9, but models different sequences of moves.

Example 23. *Consider an election with five voters and four candidates, with voters' preferences as follows:*

$$\begin{array}{ccccccc}
x_1 & \succ_1 & x_3 & \succ_1 & x_4 & \succ_1 & x_2 \\
x_2 & \succ_2 & x_1 & \succ_2 & x_3 & \succ_2 & x_4 \\
x_3 & \succ_3 & x_2 & \succ_3 & x_1 & \succ_3 & x_4 \\
x_4 & \succ_4 & x_2 & \succ_4 & x_1 & \succ_4 & x_3 \\
x_4 & \succ_5 & x_3 & \succ_5 & x_1 & \succ_5 & x_2
\end{array}$$

When needed, a lexicographic tie-breaking rule is used. Initially, in the truthful preference profile, x_4 is the winner. We show that each candidate can be the final winner in a different deviation sequence:

(a) If voter 2 deviates from x_2 to x_1, then no other voter has an incentive to deviate afterwards and thus x_1 is finally elected.

(b) If voter 3 deviates to x_2, followed by voter 5 who deviates to x_1 and voter 4 who deviates to x_2, then no other voter has an incentive to deviate afterwards and thus x_2 is finally elected.

(c) If voter 1 deviates to x_3, followed by voter 4 who deviates to x_2 and voter 5 who deviates to x_3, then no other voter has an incentive to deviate afterwards and thus x_3 is finally elected.

(d) If voter 3 deviates to x_2, followed by voter 1 who deviates to x_4, then no other voter has an incentive to deviate afterwards and thus x_4 is finally elected.

One can verify that all described deviations are valid best responses.

Consequently, the notions of possible and necessary iterative winners naturally follow from the fact that different iterative winners can arise from different deviation sequences.

Definition 16 (Possible iterative winner)**.** *A candidate x is a possible iterative winner for preference profile $\mathscr{P}$ if there exists a deviation sequence $\langle b^0, b^1, \ldots, b^r \rangle \in DS(\mathscr{P})$ such that $w^r = x$.*

Definition 17 (Necessary iterative winner)**.** *A candidate x is a necessary iterative winner for preference profile $\mathscr{P}$ if, for every deviation sequence $\langle b^0, b^1, \ldots, b^r \rangle \in DS(\mathscr{P})$ we have $w^r = x$.*

By definition, a necessary iterative winner is also a possible iterative winner.

Let us provide below some observations to make the connections between these two concepts of iterative winner and the best response deviations based on potential winners. First of all, strategic moves are only possible towards potential winners. Thus, once a candidate leaves the set of potential winners, she can never return again.

Observation 37. *If a candidate x is a possible iterative winner for preference profile $\mathscr{P}$, then there exists a deviation sequence $\langle b^0, b^1, \ldots, b^r \rangle \in DS(\mathscr{P})$ such that x is a potential winner all along the sequence: $\forall t \in \{0, 1, \ldots, r\}, x \in PW^t$. In particular, $x \in PW^0$.*

Moreover, from the definition of potential winner, if we remove one vote to a not currently winning potential winner, then she does not fulfill anymore the definition.

Observation 38. *Let us consider a deviation sequence $\langle b^0, b^1, \ldots, b^T \rangle \in DS(\mathscr{P})$ and the potential winner $x \in PW^t \setminus \{w^t\}$ such that the best response at step $t+1$ is a FPW move, i.e., the deviation from state b^t to reach b^{t+1} is performed by a voter $i \in N$ with $b_i^t = x$. Then $x \notin PW^{t+1}$.*

The concepts of possible and necessary iterative winners evaluate the outcomes of iterative voting processes from a qualitative perspective. Indeed, all deviation sequences must reach the same winner for the necessary iterative winner, whereas only one deviation sequence is required for the possible iterative winner. Another perspective is to take a more quantitative point of view. To this end, based on Section 2.6, we will also provide a probabilistic analysis of iterative voting outcomes.

Let us now start our analysis of deviation sequences both from a qualitative and quantitative perspective.

4.3 . Diversity of Iterative Winners

In this section, we will investigate how diverse iterative winners can be. We will first study the number of possible iterative winners and then focus on the extreme case with a necessary iterative winner, by analyzing the particular scenario where the deviation sequence is empty.

4.3.1 . Number of Possible Iterative Winners

We first observe that the iterative winner is determined when there are at most two potential winners.

Observation 39. *For any deviation sequence $\langle b^0, \dots, b^r \rangle$, if $|PW^t| = 2$ for a given step $t \in \{0, 1, \dots, r\}$, then the iterative winner of this sequence will be the winner of the pairwise comparison between the two candidates in PW^t.*

Observation 39 yields directly some straightforward corollaries:

Corollary 40. *If there exists a Condorcet winner c^*, and if $c^* \in PW^0$ with $|PW^0| = 2$, then c^* is the necessary iterative winner.*

Corollary 41. *A Condorcet loser can never be a possible iterative winner.*

Moreover, it can be used to bound the number of possible winners when there are only three candidates.

Proposition 42. *When $m = 3$, there exist at most two possible iterative winners.*

Proof. If there exists a Condorcet winner x then, since $m = 3$, there exists a weak Condorcet loser. In fact, x is winning every pairwise comparison therefore comparing the two other candidates tells us who is the Condorcet loser (resp., the two weak Condorcet losers). By Observation 39, the Condorcet loser (resp., the weak Condorcet loser, which is disadvantaged by the tie-breaking) cannot win. Hence, there can be at most two possible iterative winners.

If there does not exist a Condorcet winner then, since $m = 3$, we have to get either a strict or a weak Condorcet cycle of pairwise comparisons between these three candidates. In the case of a strict cycle, if we name y the initial winner, after the first strategic move we necessarily have a comparison between y and one of the other two candidates. However, with the strict Condorcet cycle one of these two need to lose against y thus, by Observation 39, this candidate cannot be elected. In the second case, the loser of the tie-breaking is also losing, helping us concluding the proof. □

Nevertheless, there exist situations where no candidate can be excluded from the set of possible iterative winners. We generalize below the observation made in Example 23 to show that for any number m of candidates, there exists a preference profile where all m candidates are possible iterative winners.

Proposition 43. *There exist elections where all m candidates are possible iterative winners, for every $m \geq 4$.*

Proof. The case of $m = 4$ has already been shown in Example 23. We will provide here a general construction for every $m \geq 5$.

We will build a preference profile $\mathscr{P}$ with $m+1$ voters and candidates $x_1, \dots, x_m$. To this purpose, we start with a preference profile $\mathscr{P}^0$ where each voter $i \in [m-1]$ has the preferences $x_i \succ_i x_{i+1} \succ_i \cdots \succ_i x_m \succ_i x_1 \succ_i \cdots \succ_i x_{i-1}$, voter m has the preferences $x_m \succ_m x_{m-1} \succ_m \cdots \succ_m x_2 \succ_m x_1$, and voter $m+1$ has the preferences $x_m \succ_{m+1} x_{m-2} \succ_{m+1} \cdots \succ_{m+1} x_{\lfloor \frac{m-1}{2} \rfloor +1} \succ_{m+1} x_1 \succ_{m+1} x_2 \succ_{m+1} \cdots \succ_{m+1} x_{\lfloor \frac{m-1}{2} \rfloor} \succ_{m+1} x_{m-1}$. Then, we obtain our final preference profile $\mathscr{P}$ from $\succ^0$ by swapping the positions of the adjacent candidates x_1 and x_m in agent 3 to agent $m-1$'s preference orders.

For each candidate, we will describe a deviation sequence which leads to her election. When needed, we use the lexicographic tie-breaking rule.

- x_1: Voter $m-1$ deviates to x_1, then voter 2 deviates to x_m, and voter 3 deviates to x_1. Afterwards, the only potential winners are x_1 and x_m and, by construction, more voters prefer x_1 to x_m. It follows that the deviation sequence will finally elect x_1.

- x_i, for $2 \leq i \leq \frac{m-1}{2}$: Voter $i+1$ deviates to x_{i+2}, then voter $i-1$ deviates to x_i, and voter m deviates to x_{i+2}. Afterwards, the only potential winners are x_i and x_{i+2} and, by construction, more voters prefer x_i to x_{i+2}. It follows that the deviation sequence will finally elect x_i.

- x_i, for $\frac{m-1}{2} < i < m-1$: Voter $i-1$ deviates to x_i, then voter $m-1$ deviates to x_1, and then voter m deviates to x_i. Afterwards, the only potential winners are x_i and x_1 and, by construction, more voters prefer x_i to x_1. It follows that the deviation sequence will finally elect x_i.

- x_{m-1} if $m > 5$: Voter $m-2$ deviates to x_{m-1}, then voter 1 deviates to x_2, and then voter m deviates to $m-1$. Afterwards, the only potential winners are x_2 and x_{m-1} and, by construction, more voters prefer x_{m-1} to x_2. It follows that the deviation sequence will finally elect x_{m-1}.

- x_{m-1} if $m = 5$: Voter $m-2$ deviates to x_{m-1}, then voter $m+1$ deviates to x_1 (this is a best response because $x_{m-2} = x_{\lfloor \frac{m-1}{2} \rfloor +1}$ when $m = 5$ and x_{m-2} is not a potential winner anymore because of the first deviation). Then, voter 2 deviates to x_{m-1}. Afterwards, the only potential winners are x_1 and x_{m-1} and, by construction, more voters prefer x_{m-1} to x_1. It follows that the deviation sequence will finally elect x_{m-1}.

- x_m: Voter 1 deviates to x_2, then voter 3 deviates to x_m. Afterwards, the only potential winners are x_2 and x_m and, by construction, more voters prefer x_m to x_2. It follows that the deviation sequence will finally elect x_m. □

A natural question is how often this situation occurs or, more generally, what is the typical number of possible iterative winners. To get quickly some first insights, we have drawn 1,000 elections, under impartial culture, where the preference profile is not an equilibrium, for each couple (m, n) with $m \in \{3, 4, 5\}$, and $5 \leq n \leq 15$, and we have computed the average number of possible iterative winners, represented in Figure 4.1. We note that, regardless the value of m, this average is rather low (less than 1.6 for all cases studied), and suggests a decreasing trend.

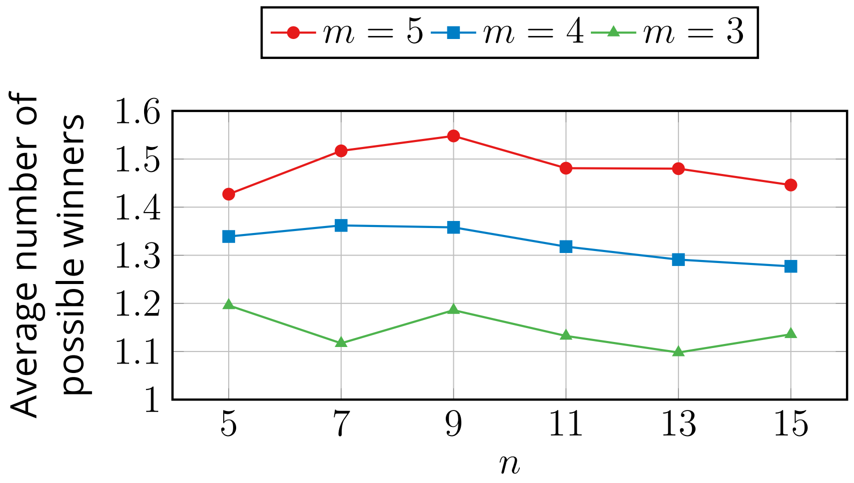

Figure 4.1: Average number of possible iterative winners in function of n (for $m \in \{3, 4, 5\}$) under IC

For a more in-depth view, we also provide in Section 4.3.1 the distribution of the number of possible iterative winners of these randomly generated elections. We indeed observe that the vast majority of instances have a unique possible (and thus necessary) iterative winner. While for each m, there are still about 20% of instances with two possible iterative winners, the situations with more than two possible iterative winners, and in particular the extreme situation from Theorem 43, seem to be extremely rare.

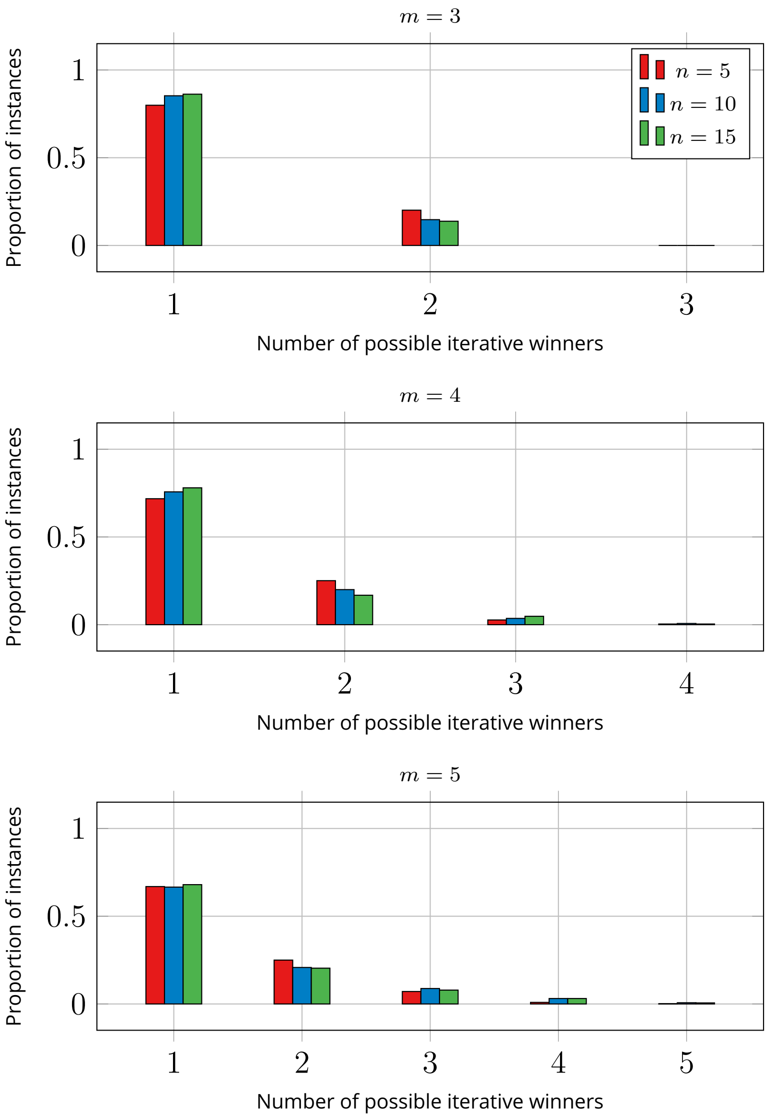

Figure 4.2: Distribution of the number of possible iterative winners as a function of n (for $m \in \{3, 4, 5\}$) under IC

Let us present the same experiment under impartial anonymous culture following the same parameters (i.e., 1,000 elections where the preference profile is not an equilibrium, for each couple (m, n) with $m \in \{3, 4, 5\}$, and $5 \leq n \leq 15$). We have computed the average number of possible iterative winners, represented in Figure 4.3.

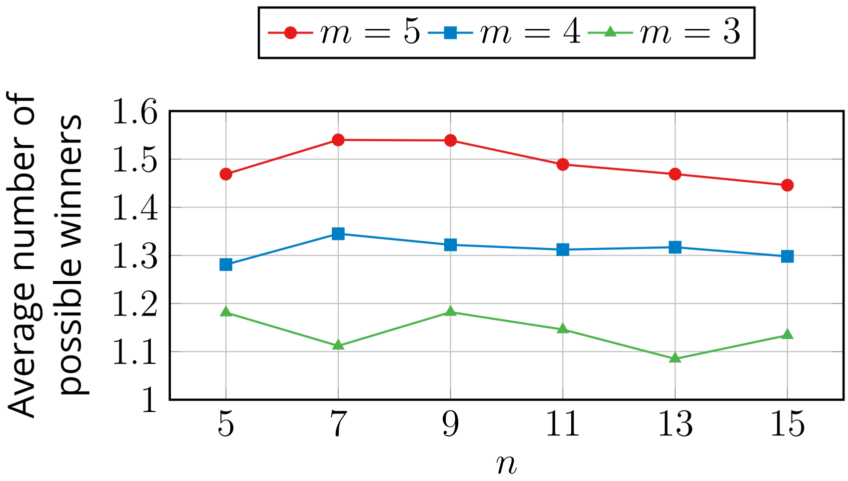

Figure 4.3: Average number of possible iterative winners in function of n (for $m \in \{3, 4, 5\}$) under IAC

Section 4.3.1 provides a detailed view of the number of winners in each case, under the same parameters as in the IC model but assuming the IAC model instead.

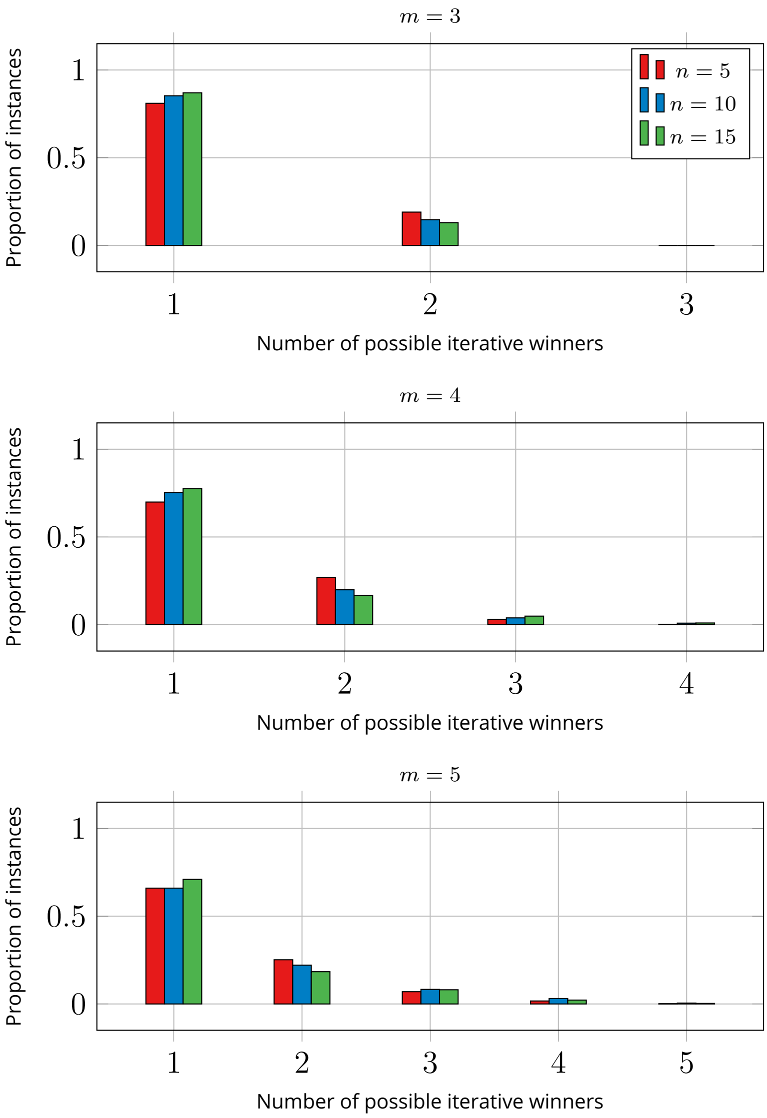

Figure 4.4: Distribution of the number of possible iterative winners under IAC for various values of m and n

The results are very similar under both impartial cultures. However, we sometimes observe small differences in the number of winners, which tends to be slightly lower under IAC. This is expected to become more pronounced as the number of voters increases, since IAC assigns lower probability weights to balanced score profiles.

4.3.2 . Extreme Case of Necessary Iterative Winner

Let us now examine how frequently the initial ballot profile is already an equilibrium, leading hence to an empty deviation sequence, where the initial winner turns out to be the only possible iterative winner, and thus the necessary iterative winner. From Proposition 73 and an adaptation of Theorem 14 in [Xia, 2012], we know that, for each m, the proportion of truthful ballot profiles from which no voter has an incentive to deviate, tends to 1 as n increases. To better understand the behavior of iterative voting processes, even in small elections, we are particularly interested here in the rate of this convergence. While deriving an exact formula seems challenging, we propose, for each pair (m, n), an increasing lower bound in n for the proportion of equilibrium profiles. Let E_n^m be the set of all preference profiles $\mathscr{P}$ that are equilibria. We start by providing some general results on the set of potential winners that will be used to establish the above-mentioned lower bound. Indeed, one way to deal with iterative voting is to track the set of potential winners over time t, i.e., PW^t.

The next lemma provides a characterization of potential winners:

Lemma 44. *Given a score vector $s \in I_n^m$, a candidate y is a potential winner for at least one voter $i \in [n]$, i.e., $y \in PW(s)$, if and only if all conditions (i) - (v) hold:*

(i) $\forall x \rhd y, s_x \leq s_y + 1$

(ii) $\forall x \rhd y, z \rhd y, s_x \leq s_y$ *or* $s_z \leq s_y$

(iii) $\forall x, z$ *such that* $y \rhd x, z, s_x \leq s_y + 1$ *or* $s_z \leq s_y + 1$

(iv) $\forall x$ *such that* $y \rhd x, s_x \leq s_y + 2$

(v) $\forall x, z$ *such that* $x \rhd y \rhd z, s_x > s_y \Rightarrow s_z \leq s_y + 1$

Proof. $\Leftarrow$ We suppose the conditions $(i) - (v)$ all hold, and we show that they are sufficient for y to be a potential winner for at least one voter. Conditions (i) and (iv) together imply that for each candidate x, we have $s_x \leq s_y + 2$:

- Suppose that there exists a candidate z such that $s_z = s_y + 2$. Condition (i) implies that $y \rhd z$. Therefore, y is a potential winner for each voter voting for z: indeed, we have $s_y + 1 = s_z - 1$, and y beats z by tie-breaking. Moreover:
  - condition (iii) ensures that for each x such that $y \rhd x$, $s_x \leq s_y + 1$, and in case of equality, y beats x by tie-breaking.

- condition (v) implies that for each $x \rhd y$, we have $s_x \leq s_y < s_y + 1$, so y wins over x.

- Suppose now that for each candidate z, $s_z < s_y + 2$, and that there exists a candidate $x \rhd y$ such that $s_x = s_y + 1$. Therefore, y is a potential winner for each voter voting for x. Indeed, $s_x - 1 < s_y + 1$, so y wins over x if ever it receives one vote from x. Moreover:
  - condition (ii) implies that for each $x' \rhd y$, $x' \neq x$, we have $s_{x'} \leq s_y$, so $s_{x'} < s_y + 1$.
  - condition (v) implies that for each z such that $y \rhd z$, $s_z \leq s_y + 1$, and in case of equality, y beats z by tie-breaking.

- Finally, it is easy to see that whenever $s_y \geq s_x$ for all $x \rhd y$, and $s_z \leq s_y + 1$ for each z such that $y \rhd z$, y is a potential winner for all voters.

$\Rightarrow$ Now we need to prove that each of these conditions is actually necessary:

- if (i) does not hold, then there is a candidate x such that $s_x > s_y + 1$. Hence, even if one voter of x deviates to y, we will still have $s_x - 1 \geq s_y + 1$, and since x wins over y by tie-breaking, y can not be a potential winner for any voter.

- if (ii) does not hold, then there exist two candidates x and z with $x, z \rhd y$, such that $s_x > s_y$ and $s_y > s_y$. Therefore, $s_x \geq s_y + 1$ and $s_z \geq s_y + 1$, and as both x and z win over y by tie-breaking, y can not be a potential winner.

- if (iii) does not hold, there exist two different candidates x, z such that $y \rhd x, z$ and $s_x > s_y + 1$, $s_z > s_y + 1$. In other words, even if y obtains one more vote (possible from one of the candidates x and z), there will be at least one of y, z having a strictly higher score than y, and therefore y can not be a potential winner.

- if (iv) does not hold, there exists a candidate x such that $s_x > s_y + 2$, in other words, $s_x - 1 > s_y + 1$, so y can not be a potential winner.

- if (v) does not hold, there exist $x \rhd y$ and z such that $y \rhd z$ such that $s_x > s_y$ and $s_z > s_y + 1$. If y obtains one extra vote from z, we will still have $s_x \geq s_y + 1$, so x wins over y by tie-breaking. Otherwise, z wins over y. Hence, y can not be a potential winner. □

Lemma 44 allows to determine the size of S_n^m, as stated below.

Lemma 45. *The number of score vectors in I_n^m with m potential winners is equal to m, i.e., $|S_n^m| = m$.*

Proof. Let $n = qm + r, r \in \{0, \ldots, m-1\}$, and $s \in S_n^m$. Let us denote by $\min_s$, resp. $\max_s$, the minimum, resp. maximum, score value in s. Without loss of generality, we can rename the candidates as $1, 2, \ldots, m$ so that $i \rhd j$ iff $i < j$, and the score of candidate i corresponds to the i-th component s_i of s.

Since $s \in S_n^m$, the conditions $(i) - (v)$ of Lemma 44 must be satisfied for each component s_i of s. In particular, we can make the following three observations:

O_1 : $\min_s \geq q-1$: let us assume for contradiction that $\min_s \leq q-2$. Then the condition (iv) of Lemma 44 implies that $\max_s \leq q$ for each $i \in [m]$, so

$$\sum_{i=1}^{m} s_i \leq (q-2) + (m-1)q < n.$$

O_2 : $\max_s \leq q+2$: similarly to the previous case, let us assume for contradiction that $\max_s \geq q+3$. Then $\min_s \geq q+1$, and

$$\sum_{i=1}^{m} s_i \geq (q+3) + (m-1)(q+1) = qm + q + 2 > n.$$

O_3 : It is easy to see that $\min_s \leq q$ and $\max_s \geq q$.

We are now ready to prove the statement by case distinction on r:

- $\mathbf{r = 0}$: There are two possible values of $\min_s$:
  - $\min_s = q-1$. Then we necessarily have $\max_s = q+1$, otherwise, the sum of all components of s would be strictly less than n. Conditions (ii) and (iii) of Lemma 44 imply that there is a unique component of score $\max_s$, which implies that there is also a unique component of score $\min_s$ (to ensure that $\sum_{i=1}^{m} s_i = n$). The condition (v) implies that $s_1 = q-1$. We then need to choose the candidate $i \in \{2, \ldots, m\}$ such that $s_i = q+1$, all remaining candidates achieving the score of q–we note that for each possible value of i, the resulting vector satisfies Lemma 44. This yields $m-1$ vectors of S_n^m.
  - $\min_s = q$. We have then $\max_s = q$–otherwise, the sum of all components is greater than $mq = n$. There is a unique vector of this type, where all components are of value q.

  Put together, we have $|S_n^m| = (m-1) + 1 = m$.

- $\mathbf{r \geq 1}$: The previous case implies that there is no $s \in S_n^m$ such that $\min_s = q-1$, and it is easy to see that $\max_s > q$. Hence, the above observations imply that $\min_s = q$, and $\max_s \in \{q+1, q+2\}$:

- $\max_s = q+1$: there are r components of s of value $q+1$, and $(m-r)$ components of value q. The condition (ii) of Lemma 44 implies that for each $i \in [m]$ such that $s_i = q$, there is at most one $j < i$ such that $s_j = q+1$. Therefore, for each $i > (m-r)+1, s_i = q+1$–in other words, the $(r-1)$ last components of s equal $q+1$. There is one remaining component of value $q+1$ to be placed to one of the $(m-r)+1$ first positions. It is easy to check that regardless the choice, the score will satisfy all conditions of Lemma 44. Hence, there are $(m-r+1)$ scores of this type in S_n^m. Note that if $r = 1$, we are done, and $|S_n^m| = m$.
- $\max_s = q+2$ - note that this can only occur for $r \geq 2$. There are then $(r-2)$ components of value $q+1$, and $(m-r+1)$ components of value q. Note that there always exist (at least two) components of value q, so conditions (ii) and (iii) of Lemma 44 imply that there is a unique component of score $\max_s = q+2$. The condition (v) of Lemma 44 implies that for each pair i, j such that $s_i = q, s_j = q+1$, we have $i < j$. Similarly, the condition (i) implies that for each pair i, j such that $s_i = q, s_j = q+2$, we have $i < j$. In other words, the $(m-r+1)$ first components of s are all of value q, and we need to place the unique component of value $q+2$ to one of the remaining $(r-1)$ places. As previously, it is easy to check that each possible choice yields a score satisfying Lemma 44. Hence, there are $(r-1)$ scores of this type in S_n^m.

Putting both types together, we have $|S_n^m| = (m-r+1)+(r-1) = m$. $\square$

Using the result of Lemma 45 as base case, we can finally determine the size of S_n^j for each $j \in [m]$.

Lemma 46. *For each* $k \in [m]$, $|S_n^{m-k}| = (m-k) \cdot \binom{n+k-2}{k}$.

Proof. Let us start by defining the set of partial scores $\tilde{S}_n^{m-k}$ as follows: for each $s \in S_n^{m-k}$, we define $\tilde{s} \in \tilde{S}_n^{m-k}$ such that $\tilde{s}_i = s_i$ for each $i \notin PW(s)$, and for each $j \in PW(s)$, $\tilde{s}_j$ is a variable such that we have

$$\sum_{j \in PW(s)} \tilde{s_j} = n - \sum_{i \notin PW(s)} s_i.$$

Note that two (or more) scores $s, s' \in S_n^{m-k}$ can yield the same partial score of $\tilde{S}_n^{m-k}$ - this happens if $PW(s) = PW(s')$ and each non-potential winner gets the same number of votes in both s and s'. We remove these duplicates from $\tilde{S}_n^{m-k}$. Lemma 45 implies that each partial score of $\tilde{S}_n^{m-k}$ can be completed into $(m-k)$ scores of S_n^{m-k}. Therefore, we have $|S_n^{m-k}| = (m-k) \cdot |\tilde{S}_n^{m-k}|$, and it remains to prove that $|\tilde{S}_n^{m-k}| = \binom{n+k-2}{k}$.

We proceed by induction on k. If $k = 0$, Lemma 45 implies that $|S_n^{m-k}| = (m-k)$, and $|\tilde{S}_n^{m-k}| = 1 = \binom{n-2}{0}$. Let us now suppose that for given k, we have, for each $n \geqslant 0, \tilde{S}_n^{m-k} = \binom{n+k-2}{k}$, and let us prove that, for every $n \geqslant 0, \tilde{S}_n^{m-(k+1)} = \binom{n+k-1}{k+1}$.

We have:

$$\begin{aligned}
\tilde{S}_{n+1}^{m-k} &= \binom{n+1+k-2}{k} = \frac{(n+k-1)!}{k!(n-1)!} = \\
&= \frac{(n+k-1)!(k+1)}{k!(k+1)(n-2)!(n-1)} = \\
&= \binom{n+(k+1)-2}{k+1} \cdot \frac{k+1}{n-1} = \\
&= \tilde{S}_n^{m-(k+1)} \cdot \frac{k+1}{n-1}
\end{aligned}$$

Therefore, we get, for every $n \geqslant 0, \tilde{S}_n^{m-(k+1)} = \tilde{S}_{n+1}^{m-k} \cdot \frac{n-1}{k+1} = \binom{(n+1)+k-2}{k} \cdot \frac{n-1}{k+1} = \frac{(n+k-1)!(n-1)}{k!(k+1)(n-1)!} = \frac{(n+k-1)!(n-1)}{k!(k+1)(n-1)!} = \binom{n+k-1}{k+1}$, which ends the proof. □

We are now ready to present the main results of this section, which establish a lower bound on the probability that a preference profile (under impartial anonymous culture or impartial culture) is an equilibrium. We begin with the case of impartial anonymous culture.

Theorem 47. *Under impartial anonymous culture (IAC), $\mathbb{P}_{IAC}(E_n^m) \geq \mathbb{P}_{IAC}(S_n^1)$, where $\mathbb{P}_{IAC}(S_n^1)$ increases with respect to n.*

Proof. As $S_n^1 \subset E_n^m$, we have $\mathbb{P}_{IAC}(E_n^m) \geqslant \mathbb{P}_{IAC}(S_n^1)$. Under IAC, we have $\mathbb{P}_{IAC}(S_n^1) = \frac{|S_n^1|}{|I_n^m|}$. By Lemma 46 (applied for $k = m-1$), we get $|S_n^1| = \binom{n+(m-1)-2}{m-1}$, and we have $|I_n^m| = \binom{n+m-1}{m-1}$.

Put together, we obtain, after simplification, $\mathbb{P}_{IAC}(S_n^1) = \frac{n \cdot (n-1)}{(n+m-1)(n+m-2)}$. It remains to be proven that $\mathbb{P}_{IAC}(S_n^1)$ increases with respect to n. Indeed, we have $\mathbb{P}_{IAC}(S_{n+1}^1) - \mathbb{P}_{IAC}(S_n^1) = \frac{2m-2}{(n+1)\cdot(n+2)\cdot(n+3)} > 0$ whenever $n > 0$ and $m > 2$. □

We provide below a brief illustration of the growth rate of this lower bound.

Example 24. *In an election with 3 candidates, the probability for a preference profile to be at equilibrium under IAC is at least 0.68 for 10 voters and at least 0.82 for 20 voters. In an election with 5 candidates, this probability is at least 0.49 for 10 voters and at least 0.69 for 20 voters.*

Figure 4.5 is an illustration for $m = 3$ and $m = 5$:

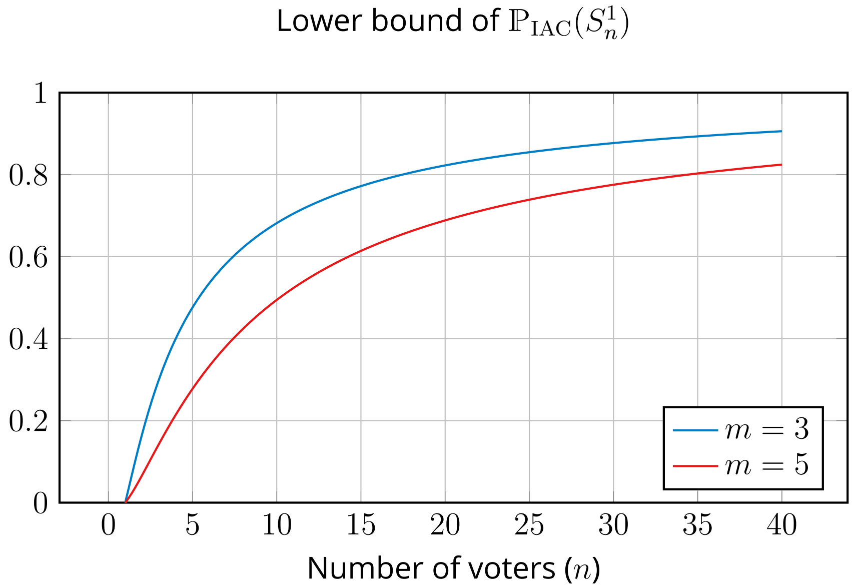

Figure 4.5: Plot of $\mathbb{P}_{\text{IAC}}(S_n^1) = \frac{n(n-1)}{(n+2)(n+1)}$ from Theorem 47

We now establish an analogous result under impartial culture, starting with the following observation, based on the fact that there are n voters' preferences independently sampled from the same distribution, and we have m possibilities for the most preferred candidate of each voter.

Observation 48. *Whenever all voters' preferences are sampled with independent and identical random variables, then the resulting score vector $s^\top$ follows a multinomial law Multi(q, n) where $q = (q_1, \ldots, q_m)$ and $q_j := \mathbb{P}_C(\{\mathcal{W}_P(s^\top) = j\})$, for every $j \in M$.*

Under impartial culture, computing explicitly $\mathbb{P}_{IC}(S_n^1)$ becomes much more harder. Instead, we prove the existence of an increasing lower bound in n.

Theorem 49. *Under impartial culture (IC), $\mathbb{P}_{IC}(E_n^m) \geq 1 + \frac{m\cdot(m-1)}{2} \cdot (\phi(\frac{-2}{\sigma\cdot\sqrt{n}}) - \phi(\frac{2}{\sigma\cdot\sqrt{n}}))$, where ϕ is the cumulative distribution function of a standard Gaussian, $\sigma = \sqrt{\frac{2}{m}}$ and this probability is increasing with respect to n.*

Proof. As in the proof of Theorem 47, we use $\mathbb{P}_{IC}(E_n^m) \geqslant \mathbb{P}_{IC}(S_n^1)$. We start by the following remark: "if for each pair of candidates $i, j \in M$, $|s_i - s_j| \geq 2$, then there is a unique potential winners in score s". Therefore,

$$\mathbb{P}_{IC}(S_n^1) \geq \mathbb{P}_{IC}(\forall i, j \in M, |s_i - s_j| \geq 2)$$

There exist $\binom{m}{2} = \frac{m\cdot(m-1)}{2}$ pairs of candidates. We denote $X_k^{(i)}$ the random variable that equals if the k-th voter has voted for the i-th candidate, and 0 otherwise. We then denote

$$Y_k^{i,j} = X_k^{(i)} - X_k^{(j)}, \forall i \neq j$$

the difference of those random variables such that $s_i - s_j = \sum_{k=0}^{n} Y_k^{i,j}$. Note that $(Y_k^{i,j})_{1\leqslant k\leqslant n}$ are independent,

$$\mathbb{P}_{IC}(Y_k = 1) = \mathbb{P}_{IC}(Y_k = -1) = \frac{1}{m}$$

and

$$\mathbb{P}_{IC}(Y_k = 0) = 1 - \frac{2}{m}$$

Therefore,

$$\mathbb{P}_{IC}(\forall i, j \in M, |S_i - S_j| \geq 2) = \mathbb{P}_{IC}(\forall i, j \in M, |\sum_{k=0}^{n} Y_k^{i,j}| \geq 2)$$

. By Bonferroni's inequality,

$$\mathbb{P}_{IC}(\forall i, j \in M, |\sum_{k=0}^{n} Y_k^{i,j}| \geq 2) \geq \sum_{k=0}^{\frac{m\cdot(m-1)}{2}} \mathbb{P}_{IC}(|\sum_{k=0}^{n} Y_k^{i,j}| \geq 2) - (\frac{m\cdot(m-1)}{2} - 1)$$

As all $Y_k^{i,j}$ follow the same law we have:

$$\sum_{k=0}^{\frac{m\cdot(m-1)}{2}} \mathbb{P}_{IC}(|\sum_{k=0}^{n} Y_k^{i,j}| \geq 2) - (\frac{m\cdot(m-1)}{2} - 1)$$

$$= 1 + \frac{m\cdot(m-1)}{2} \cdot (\mathbb{P}_{IC}(|\sum_{k=0}^{n} Y_k^{i,j}| \geq 2) - 1)$$

It remains to find a lower bound to

$$\mathbb{P}_{IC}(|\sum_{k=0}^{n} Y_k^{i,j}| \geq 2)$$

$$= \mathbb{P}_{IC}(\sum_{k=0}^{n} Y_k^{i,j} \geq 2) + \mathbb{P}_{IC}(\sum_{k=0}^{n} Y_k^{i,j} \leqslant -2)$$

Using Berry-Essen's theorem [Berry, 1941; Esseen, 1942] we get the following lower bounds:

$$\mathbb{P}_{IC}(\sum_{k=0}^{n} Y_k^{i,j} \leqslant -2) \geq \phi(\frac{-2}{\sigma\cdot\sqrt{n}}) - \frac{C\cdot\rho}{\sigma^3\cdot\sqrt{n}}$$

and

$$\mathbb{P}_{IC}(\sum_{k=0}^{n} Y_k^{i,j} \geq 2) \geq \mathbb{P}_{IC}(\sum_{k=0}^{n} Y_k^{i,j} > 2)$$

$$= 1 - \mathbb{P}_{IC}(\sum_{k=0}^{n} Y_k^{i,j} \leqslant 2) \geq 1 - \phi(\frac{2}{\sigma \cdot \sqrt{n}}) - \frac{C \cdot \rho}{\sigma^3 \cdot \sqrt{n}}$$

, where σ is the standard deviation, C is a constant, ρ the moment of order 3, ϕ is the repartition of a standard gaussian and $t \cdot \sigma\sqrt{n} = 2$, t from the original formula. However, if we compute ρ for $\sum_{k=0}^{n} Y_k^{i,j}$, we obtain $\rho = 0$, therefore simplifying the inequality as follows:

$$\mathbb{P}_{IC}(|\sum_{k=0}^{n} Y_k^{i,j}| \geq \phi(\frac{-2}{\sigma \cdot \sqrt{n}}) + 1 - \phi(\frac{2}{\sigma \cdot \sqrt{n}})$$

As expected, this converges to 1 asymptotically in n since $\phi(0) = 0$. We finally get:

$$\mathbb{P}_{IC}(E_n^m) \geq 1 + \frac{m \cdot (m-1)}{2} \cdot (\phi(\frac{-2}{\sigma \cdot \sqrt{n}}) - \phi(\frac{2}{\sigma \cdot \sqrt{n}}))$$

Note that

$$\phi(\frac{-2}{\sigma \cdot \sqrt{n}}) - \phi(\frac{2}{\sigma \cdot \sqrt{n}})$$

can be verified to be negative since this the repartition of a standard gaussian. Therefore, this bound is increasing and goes to 1 asymptotically in n. □

Thanks to Theorem 49, the lower bound of IC increases slowly compared to that of IAC.

Example 25. *In an election with 3 candidates, 70 voters are needed for the probability to exceed 0.30, and 137 voters for it to exceed 0.5. In an election with 5 candidates, 1000 voters are needed for the probability to exceed 0.2.*

Figure 4.7 are some illustrations for $m = 3$ and $m = 5$:

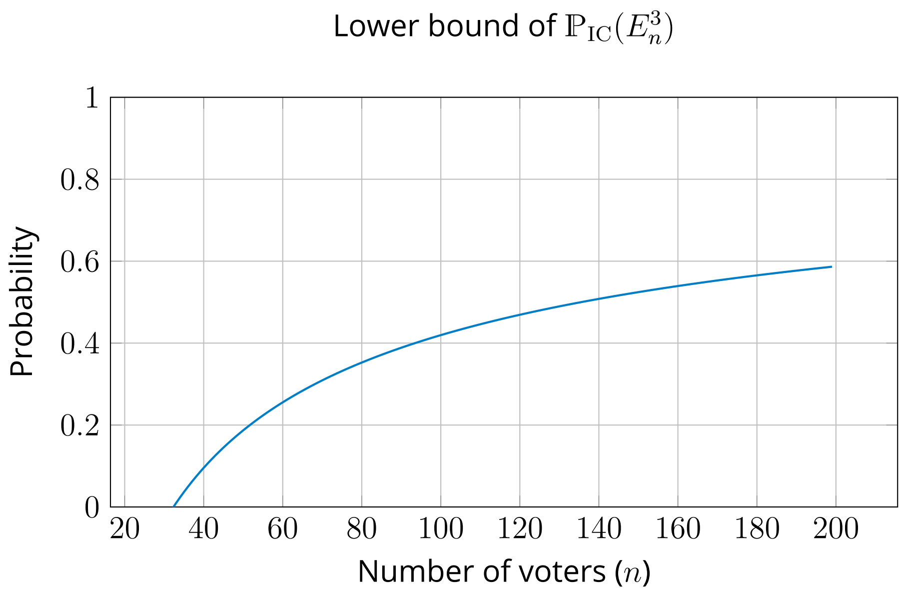

Figure 4.6: Lower bound on $\mathbb{P}_{\mathrm{IC}}(E_n^3)$ with respect to n from Theorem 49

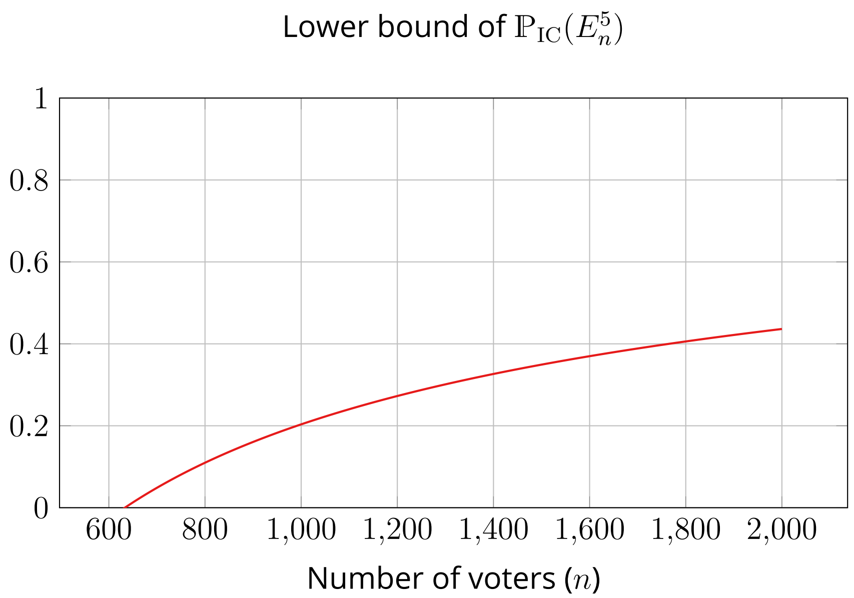

Figure 4.7: Lower bound on $\mathbb{P}_{\mathrm{IC}}(E_n^5)$ with respect to n from Theorem 49

4.4 . Possible and Necessary Winner Problems

The situation analyzed in the previous section, where no deviation can occur from the initial ballot profile, is an extreme case of a scenario with a necessary iterative winner. In this section, we aim to go further on the recognition of scenarios where given candidates are possible or necessary iterative winners, by investigating the complexity of the associated existence problems. More precisely, we will study the following decision problem PossibleIterativeWinner (resp., NecessaryIterativeWinner): *Given an election $(N, M, \mathscr{P}, \rhd)$ and a candidate $x \in M$, is x a possible (resp., necessary) iterative winner?*

First of all, the two problems turn out to be equivalent when the initial potential winner set is limited to at most two candidates.

Proposition 50. PossibleIterativeWinner *and* NecessaryIterativeWinner *are equivalent and can be solved in polynomial time when $|PW^0| \leq 2$.*

Proof. If $|PW^0| = \{x\}$ then, by Observation 37, x is the unique possible–and thus necessary–winner.

If $PW^0 = \{x, y\}$ with $x = w^0$ then, by Observation 37, only x or y can be iterative winners. Since voters can only deviate to favor x or y and voters in $N^x \cup N^y$ have no incentive to deviate, candidate x (resp., y) is the unique possible–and thus necessary–iterative winner iff $|(N \setminus (N^x \cup N^y))^{x \succ y}| \geq |(N \setminus (N^x \cup N^y))^{y \succ x}|$ (resp., $|(N \setminus (N^x \cup N^y))^{y \succ x}| > |(N \setminus (N^x \cup N^y))^{x \succ y}|$). ☐

Note that the equivalence between the two problems does not hold starting with three candidates in the initial potential winner set. Let us illustrate that through the following example.

Example 26. *Consider the following preference profile with $n = 3$ voters and $m = 3$ candidates where $x_1 \succ_1 x_2 \succ_1 x_3$, $x_2 \succ_2 x_3 \succ_2 x_1$, and $x_3 \succ_3 x_2 \succ_3 x_1$, and a is the initial winner. If voter 2 (resp., voter 3) first deviates then x_3 (resp., x_2) is the iterative winner. It follows that x_2 and x_3 are the two possible iterative winners, but none of them is a necessary iterative winner.*

In addition of the non-equivalence of the two problems, even their complexity class differs. We first establish below that the necessary iterative winner problem can be solved in polynomial time.

Theorem 51. NecessaryIterativeWinner *is in P.*

Proof. We will provide a polynomial number of conditions, which can be checked in polynomial time, on the preference profile $\mathscr{P}$ to determine whether a given candidate y is a necessary winner. We distinguish the cases where y is the initial truthful winner w^0 or not.

Is candidate $y \neq w^0$ a necessary winner? Trivially, by Observation 37, if $y \notin PW^0$, then she is not a possible, and thus not a necessary iterative winner. Therefore, we assume from now on that $y \in PW^0$. Let us give some necessary conditions for y to be a potential winner along each possible deviation sequence:

(i) For all $z \in PW^0 \setminus \{w^0, y\}$ and all $i \in N^y$, we have $x \succ_i z$: Otherwise, there exists a candidate $z \in PW^0 \setminus \{w^0, y\}$ and a voter $i \in N^y$ such that $z \succ_i x$. There exists then a deviation sequence where i is the first voter to deviate, and she will do so from her initial ballot for y to a ballot for z. It follows from Observation 38 that y is not a potential winner anymore after this first step and thus, by Observation 37, y will not be the iterative winner in this deviation sequence, implying that y is not a necessary iterative winner.

(ii) Assume that (i) holds. For every candidate $z_1 \in M \setminus \{w^0, y\}$ and voter $i \in N^{z_1}$, we must have either $w^0 \succ_i z$ for every $z \in PW^0 \setminus \{w^0, z_1\}$, or $y \succ_i w^0$. Otherwise, there exist a candidate $z_1 \in M \setminus \{w^0, y\}$, a potential winner $z_2 \in PW^0 \setminus \{w^0, y, z_1\}$ and a voter $i \in N^{z_1}$ such that $z_2 \succ_i z$, for every $z \in PW^0 \setminus \{z_1, z_2\}$. There exists then a deviation sequence where i is the first voter to deviate, and she will do so from her initial ballot for z_1 to a ballot for z_2 (that she prefers to w^0). Since w^0 was the initial winner, she is still a potential winner after this deviation. Therefore, there exists a second deviation in which a voter $j \in N^y$ deviates from her initial ballot for y to a ballot for w^0 (that she prefers over all potential winners other than y, by (i)). Thus, by Observations 37 and 38, y will not be the iterative winner in this deviation sequence, implying that y is not a necessary iterative winner.

Assume that the conditions (i) and (ii) hold, and let us look closer to point (ii) where there are two cases to distinguish:

- If, for every candidate $z_1 \in M \setminus \{w^0, y\}$ and voter $i \in N^{z_1}$, we have $w^0 \succ_i z$ for every $z \in PW^0 \setminus \{w^0, z_1\}$, then no deviation can occur. It follows that the initial winner w^0 will be the unique possible–and thus necessary–iterative winner, implying that y cannot be a necessary iterative winner.

- Otherwise, there exist a candidate $z_1 \in M \setminus \{w^0, y\}$ and a voter $i \in N^{z_1}$ such that $y \succ_i w^0$. In that case, by Observation 39, y is the unique possible–and thus necessary–iterative winner iff

$$|(\bigcup_{z \in M \setminus \{w^0, y\}} N^z)^{y \succ w^0}| > |(\bigcup_{z \in M \setminus \{w^0, y\}} N^z)^{w^0 \succ y}|$$

Is candidate w^0 a necessary winner? If, for every candidate $z_1 \in M \setminus \{w^0\}$ and voter $i \in N^{z_1}$, we have $w^0 \succ_i z$ for every $z \in PW^0 \setminus \{w^0, z_1\}$, then no deviation can occur and w^0 is a necessary winner. Therefore, we assume from now on that there exists a candidate $z_1 \in M \setminus \{w^0\}$, a voter $i \in N^{z_1}$ and a candidate $z_2 \in PW^0 \setminus \{w^0, z_1\}$ such that $z_2 \succ_i w^0$, i.e., there is a voter with an incentive to deviate from the initial truthful profile b^0.

For a candidate $y \in PW^0 \setminus \{w^0\}$, let $PW^{1,y} \subseteq PW^0$ denote the set of potential winners in the strategy profile $b^{1,y}$ resulting from a best response electing candidate y, where a voter changes her initial ballot to a ballot for y, performed from the initial truthful profile b^0. Suppose that there exist a candidate $z_1 \in M \setminus \{w^0\}$, a potential winner $z_2 \in PW^0 \setminus \{w^0, z_1\}$, a voter $i \in N^{z_1}$ such that $z_2 \succ_i z$ for every $z \in PW^0 \setminus \{z_1, z_2\}$, a candidate $y \in PW^{1,z_2} \setminus \{w^0, z_2\}$ and a voter $j \in N^{w^0}$ such that $y \succ_i z$ for every $z \in PW^{1,z_2} \setminus \{w^0, y\}$. It means that there exists a deviation sequence where voter i is the first voter to deviate and she does so from her initial ballot for z_1 to a ballot for z_2 (that she prefers to w^0), and then voter j is the second voter to deviate and she does so from her initial ballot for w^0 to a ballot for y (that she prefers to the current winner z_2). It follows from Observations 37 and 38 that w^0 will not be the iterative winner in this deviation sequence, implying that w^0 is not a necessary iterative winner. Therefore, we assume from now on that, for every candidate $z_1 \in M \setminus \{w^0\}$, potential winner $z_2 \in PW^0 \setminus \{w^0, z_1\}$, voter $i \in N^{z_1}$ such that $z_2 \succ_i z$ for every $z \in PW^0 \setminus \{z_1, z_2\}$, we have all voters $j \in N^{w^0}$ who prefer z_2 over all potential winners in $PW^{1,z_2} \setminus \{w^0, z_2\}$.

Let Z denote the set of all potential winners to which there is a voter who has an incentive to deviate and $A(y)$ the set of voters having an incentive to deviate to $y \in Z$, i.e., $Z := \{y \in PW^0 \setminus \{w^0\} : \exists z_1 \in M \setminus \{w^0, y\},\ i \in N^{z_1} \text{ s.t. } y \succ_i z,\ \forall z \in PW^0 \setminus \{y, z_1\}\}$ and $A(y) := \{i \in N : \exists z_1 \in M \setminus \{w^0\} \text{ s.t. } i \in N^{z_1},\ y \succ_i z,\ \forall z \in PW^0 \setminus \{y, z_1\}\}$. By definition, we have $|A(y)| > 0$ for every $y \in Z$. If $|Z| = 1$ with $Z = \{y\}$, then the only first deviations that can occur are towards candidate y and no further deviation can then occur for a candidate different from w^0 or y and, by assumption, voters in N^{w^0} are satisfied by both candidates w^0 and y and thus do not deviate. It follows that w^0 is the unique possible–and thus necessary–iterative winner iff $(\bigcup_{z \in M \setminus \{w^0, y\}} N^z)^{w^0 \succ y} \geq (\bigcup_{z \in M \setminus \{w^0, y\}} N^z)^{y \succ w^0}$. Let us thus assume, from now on, that $|Z| > 1$.

By assumption, for every potential winner $z \in Z$, every voter $j \in N^{w^0}$ prefers z to any other potential winner $y \in Z \cap PW^{1,z}$. It follows that, for every candidates $z_1, z_2 \in Z$ such that $z_1 \neq z_2$, we have either $z_1 \notin PW^{1,z_2}$ or $z_2 \notin PW^{1,z_1}$. Note that both cannot hold simultaneously because for $z_2 \notin PW^{1,z_1}$ to hold, since $z_2 \in PW^0$, we need that $z_1 \rhd z_2$ or that $z_2 \rhd w^0 \rhd z_1$ while z_2 has one vote less than both z_1 and w^0 in the initial scores; under either condition z_1 is still a potential winner in the ballot profile b^{1,z_2} resulting from a best response from the truthful initial profile where z_2 gets one additional vote. Consequently, for every $z_1, z_2 \in Z$, we have either $z_1 \notin PW^{1,z_2}$ and $z_2 \in PW^{1,z_1}$ and all voters in N^{w^0} prefer z_1 to z_2, or $z_2 \notin PW^{1,z_1}$ and $z_1 \in PW^{1,z_2}$ and all voters in N^{w^0} prefer z_2 to z_1. We can thus assume, w.l.o.g., that $Z = \{z_1, \dots, z_\ell\}$, with $z_t \notin PW^{1,z_{t'}}$, $z_{t'} \in PW^{1,z_t}$, and $z_t \succ_j z_{t'}$ for every voter $j \in N^{w^0}$ and every $1 < t < t' < \ell$.

For given indices $t_1 < t_2 < t_3 \in [\ell]$, let $A^{t_2,t_3}(t_1)$ denote the set of voters in $A(z_{t_1})$ who prefer z_{t_2} to w^0 and to z_t for all $t_3 \leq t \leq \ell$, i.e., $A^{t_2,t_3}(t_1) := \{i \in A(z_{t_1}) : z_{t_2} \succ_i w^0$ and $z_{t_2} \succ_i z_t, \forall t_3 \leq t \leq \ell\}$. If there exist $t, t' \in [\ell]$ such that $t < t'$ with $|A(z_t) \cup \bigcup_{t'' \in [t'-1]} A^{t,t'}(t'')| > 1$, then there exists a deviation sequence where a voter $i_1 \in A(z_t)$ first deviates to a ballot for z_t, then a voter $j \in A(z_{t'})$ deviates to a ballot for $z_{t'}$, and another voter $i_2 \in A(z_t) \cup \bigcup_{t'' \in [t'-1]} A^{t,t'}(t'')$ then deviates to a ballot for z_t, creating a gap too important between the score of the current winner and the score of w^0, which thus cannot be a potential winner anymore. Consequently, by Observation 37, w^0 will not be the iterative winner in this deviation sequence, implying that w^0 is not a necessary winner.

Otherwise, it means that, for every candidate $z_t \in Z$, all voters in $A(z_t)$ prefer z_ℓ or w^0 to every $z_{t'} \in Z \setminus \{z_t, z_\ell\}$ (if not, $|A^{t',\ell}(t)| > 0$, and the previous condition would hold). Since, by definition, all voters in $N \setminus \bigcup_{t \in [\ell-1]} A(z_t)$ prefer w^0 to all candidates in Z, it follows that w^0 will be the unique possible–and thus necessary–iterative winner iff $(\bigcup_{z \in M \setminus \{w^0, z_\ell\}} N^z)^{w^0 \succ z_\ell} \geq (\bigcup_{z \in M \setminus \{w^0, z_\ell\}} N^z)^{z_\ell \succ w^0}$. □

In contrast, the possible iterative winner problem is NP-complete.

Theorem 52. PossibleIterativeWinner *is NP-complete.*

Proof. The problem belongs to NP because, given a sequence of voter strategic deviations, we can check in polynomial time whether it is valid and eventually elects a target candidate t at equilibrium because the length of such a sequence is polynomially bounded (see Theorem 36).

For hardness, we perform a reduction from Exact Cover by 3-Sets (X3C), a problem known to be NP-complete [Garey and Johnson, 1979]. In an instance of X3C, we are given a set $X = \{x_1, x_2, \dots, x_{3q}\}$ and a set $S = \{S_1, S_2, \dots, S_r\}$ of 3-element subsets of X and we ask whether there exists an exact cover, i.e., a subset $S' \subseteq S$ of size $|S'| = q$ such that every element of X occurs in exactly one member of S', in other words, S' is a partition of X. We consider the variant of the problem, that is still hard, where each element x_i occurs in exactly three subsets of S, implying that $r = 3q$.

For each element $x_i \in X$, we create a corresponding element-candidate y_i. For each subset $S_j \in S$, we create one candidate d_j and three subset-candidates s_j^1, s_j^2, and s_j^3 associated with the three elements of subset S_j. For each $\ell \in [2q]$, we create an candidate z_ℓ, supposed to correspond to the $2q$ elements of S which are not chosen for the partition of X. We additionally create five candidates, namely a, b, c, e, and t. The tie-breaking rule is given by the following linear order over the candidates: $a \rhd b \rhd c \rhd z_1 \rhd \cdots \rhd z_{2q} \rhd y_1 \rhd \cdots \rhd y_{3q} \rhd t \rhd d_1 \rhd \cdots \rhd d_{3q} \rhd e \rhd s_1^1 \rhd s_1^2 \rhd s_1^3 \rhd \cdots \rhd s_{3q}^1 \rhd s_{3q}^2 \rhd s_{3q}^3$.

For each element $x_i \in X$, we create $3q$ element-voters Y_i^ℓ, for $\ell \in [3q]$, whose preferences are as follows for each $i \in [3q]$, where $s^\ell(x_i)$ stands for the subset-candidate s_j^k such that the k^{th} element of subset S_j is the ℓ^{th} occurrence of element x_i, when $\ell \in [3]$:

$$\begin{array}{lll} Y_i^\ell: & y_i \succ s^\ell(x_i) \succ a \succ t \succ [\dots] & \text{if } \ell \in [3] \\ Y_i^\ell: & y_i \succ a \succ t \succ [\dots] & \text{if } 4 \leq \ell \leq 3q \end{array}$$

For each $\ell \in [2q]$, we create $3q$ voters Z_ℓ^j, for $j \in [3q]$, with the following preferences:

$$Z_\ell^j: \quad z_\ell \succ c \succ y_1 \succ \cdots \succ y_{3q} \succ s_j^1 \succ s_j^2 \succ s_j^3 \succ d_j \succ a \succ t \succ [\dots]$$

To allow all candidates to be potential winners, we create the voters A^ℓ, B^ℓ, C^ℓ, D_j^ℓ, E^ℓ, $S_{j,k}^\ell$, and T^ℓ, for $j, \ell \in [3q]$ and $k \in [3]$, with the following preferences:

$$\begin{array}{ll} A^\ell: & a \succ b \succ t \succ [\dots] \\ B^\ell: & b \succ a \succ t \succ [\dots] \\ C^\ell: & c \succ e \succ a \succ t \succ [\dots] \\ U^\ell: & u \succ a \succ t \succ [\dots] \\ & \text{for } (U, u) \in \bigcup_{j \in [3q]} \{(D_j, d_j), (S_{j,k}, s_j^k)\} \cup \{(E, e)\} \\ T^\ell: & t \succ a \succ b \succ [\dots] \end{array}$$

We finally create an candidate f and a voter F with the following preferences:

$$F: \quad f \succ z_1 \succ \cdots \succ z_{2q} \succ y_1 \succ \cdots \succ y_{3q} \succ t \succ a \succ b \succ [\dots]$$

By construction, in the truthful initial profile, there are exactly $3q$ votes for each candidate except f, and thus candidate a is winning, thanks to the tie-breaking rule.

We claim that there exists a subset $S' \subseteq S$ which is a partition of X iff there exists a sequence of voter strategic deviations which leads to the victory of candidate t.

$\Longrightarrow$: Suppose first that there exists a subset $S' \subseteq S$ which is a partition of X, say $S' = \{S_{j'_1}, \dots, S_{j'_q}\}$ where $j'_1 < \cdots < j'_q$. By definition, each element x_i is covered by exactly one element of S', say that x_i is covered by the element of S' which contains the k_i^{th} occurrence of element x_i, for $k_i \in [3]$. We will thus let voter $Y_i^{k_i}$ deviate to subset-candidate $s^{k_i}(x_i)$. We will schedule these deviations with respect to the tie-breaking order $\rhd$, i.e., voter $X_i^{k_i}$ deviates before voter $X_{i'}^{k_{i'}}$, with $s_j^k := s^{k_i}(x_i)$ and $s_{j'}^{k'} := s^{k_{i'}}(x_{i'})$, iff $j > j'$, or $j = j'$ and $k > k'$. It follows that each candidate y_i loses one vote, while each candidate $s_{j'_\ell}^k$ gains one vote, for each $\ell \in [q]$ and $k \in [3]$, by decreasing order of indices.

Then, voter C^1 deviates from her vote for candidate c to a vote for candidate e, and thus c loses one vote. It follows that none of the candidates $y_1 \ldots, y_{3q}$ and c are potential winners anymore, nor are any of the subset-candidates associated with elements of $S \setminus S'$.

Let us consider the set of non-chosen elements of S, i.e., $S \setminus S' = \{S_{j_1}, \ldots, S_{j_{2q}}\}$ where $j_1 < \cdots < j_{2q}$. For $\ell = 2q$ to $\ell = 1$, we let voter $Z_\ell^{j_\ell}$ deviate from candidate z_ℓ to candidate d_{j_ℓ}. This is a best response because none of the candidates $y_1 \ldots, y_{3q}$, c, and $s_{j_\ell}^1$, $s_{j_\ell}^2$ and $s_{j_\ell}^3$ are potential winners.

Afterwards, voter F deviates from her vote for candidate f to a vote for candidate t. This is a best response because none of the candidates $z_1, \ldots, z_{2q}$ and $y_1, \ldots, y_{3q}$ are potential winners. Now let voter A^1 deviate from her vote for candidate a to a vote for candidate b. It follows that candidate a is not a potential winner anymore. If we then let, e.g., voter D_1^1 deviate from her vote for candidate d_1 to a vote for candidate t, then b and t are the only remaining potential winners, with $3q + 1$ and $3q + 2$ votes, respectively, while the other candidates which are less (resp., more) favored than t (resp., except b) have at most $3q + 1$ (resp., $3q - 1$) votes. Since there are more voters preferring t to b than the reverse, among the voters who do not currently vote for any of them, it thus leads to a sequence of voter deviations eventually electing candidate t at the equilibrium.

$\Longleftarrow$ **:** Suppose now that there exists a sequence of voter strategic deviations which leads to the victory of candidate t. First observe that, since all candidates (except f) have initially the same score, any iterative winner must gain at least one vote and thus must have at least $3q + 1$ votes. Therefore, candidate t must gain at least one vote. Since candidates a, b, c, $z_1, \ldots, z_{2q}$, $y_1, \ldots, y_{3q}$ are more favored by the tie-breaking order $\rhd$ than t, none of them can gain a new vote before t gets one, because otherwise t would not be a potential winner anymore. All voters prefer a to t, except voters T^ℓ, who already vote for t, and voter F. Moreover, the only possibility for a to not be a potential winner before t can gain one vote, would be that some voter A^ℓ deviates, and the only possible deviation would be towards b, a contradiction. Therefore, we need that voter F deviates to t, and this is the only possible first deviation to t.

However, voter F prefers all candidates $z_1, \ldots, z_{2q}$ and $y_1, \ldots, y_{3q}$, initially potential winners, to candidate t. Therefore, we need for F to deviate to t as a first deviation to t, that none of the candidates $z_1, \ldots, z_{2q}$ and $y_1, \ldots, y_{3q}$ are potential winners, while t is still a potential winner. The only possible way to achieve this situation, is that every candidate z_ℓ and y_i, for $\ell \in [2q]$ and $i \in [3q]$, loses at least one vote.

Therefore, we need that at least one voter Y_i^ℓ, for some $\ell \in [3q]$, deviates from her current vote for y_i, for each $i \in [3q]$. The only possible deviation which can still enable the future election of t is by a voter Y_i^k for $k \in [3]$ towards $s^k(x_i)$. Let us construct the subset $S' \subseteq S$ such that all elements of S' correspond to subset-candidates $s^k(x_i)$ to which some voter Y_i^k deviates to, so that y_i is not a potential winner anymore, for each $i \in [3q]$. By definition of $s^k(x_i)$, it follows that S' covers all elements of X.

We also need that at least one voter Z_ℓ^j, for some $j \in [3q]$, deviates from her current vote for z_ℓ, for each $\ell \in [2q]$. To enable the first deviation of F to t, such a voter Z_ℓ^j should not deviate to c or $y_1, \ldots, y_{3q}$, and thus none of these candidates should be a potential winner. It follows that all previously described deviations of voters Y_i^ℓ should occur before those of Z_ℓ^j. Moreover, the only possibility for c not being a potential winner anymore is that it loses one vote, with a deviation by a voter C^ℓ, for some $\ell \in [3q]$. Such a voter must deviate to candidate e. Then, by the tie-breaking order, none of the subset-candidates not chosen for deviation by voters Y_i^ℓ can be a potential winner anymore. Voter Z_ℓ^j can thus deviate to a subset-candidate s_j^k for $k \in [3]$ which has previously been chosen for deviation by a voter Y_i^ℓ or, if none of them has been chosen, to candidate d_j if not already the winner. However, it is not possible for the future election of t that Z_ℓ^j deviates to a subset-candidate s_j^k or to a candidate d_j which has already gained votes because, otherwise, such candidates would get at least $3q+2$ votes and t would not be a potential winner anymore. It follows that each such voter Z_ℓ^j deviates to a different candidate d_j, and that no subset-candidate s_j^k, associated with the same element $S_j \in S$, has been chosen for deviation by voters Y_i^ℓ. Since there are $2q$ different such voters Z_ℓ^j associated with different elements $S_j \in S$ which are not part of S', it means that $|S'| = q$ and thus it is an exact cover of X. □

Now that we have investigated whether an arbitrary candidate can be a possible or necessary iterative winner, it makes sense to focus on particularly desirable candidates.

4.5 . About Electing the Condorcet Winner

Since electing the Condorcet winner is a desirable property for a voting rule, we now investigate the ability of the iterative voting process to elect it.

4.5.1 . The Condorcet Winner as an Iterative Winner

If a Condorcet winner exists, the natural question is whether she is guaranteed to be a possible or even a necessary iterative winner. We first study the question of a necessary iterative winner.

Proposition 53. *If $m = 3$ and the Condorcet winner is the initial winner, then she is also a necessary iterative winner.*

Proof. Let c^* be the Condorcet and initial winner. If no strategic move can be performed, we are done. Otherwise, the first strategic move of each deviation sequence cannot be neither towards nor from c^*, and by Observation 38, there are at most two potential winners after this move, c^* being one of them. Observation 39 implies that c^* is the winner of each sequence, hence the necessary winner. □

The following example shows that if the Condorcet winner is not initially winning, then she is not guaranteed to be the necessary iterative winner, even if $m = 3$.

Example 27. *Let us consider the profile $\mathscr{P} = \{x_2 \succ_1 x_3 \succ_1 x_1, x_1 \succ_2 x_2 \succ_2 x_3, x_3 \succ_3 x_2 \succ_3 x_1\}$ where x_2 is the Condorcet but not initial winner (x_1 initially wins by tie-breaking), and $PW^0 = M$. If voter 1 deviates from x_2 to x_3, we get $PW^1 = \{x_1, x_3\}$. Since $x_2 \notin PW^1$, she cannot win in this deviation sequence, therefore, she is not the necessary winner.*

Similarly, the following example shows that if $m > 3$, then the Condorcet winner is not guaranteed to be the necessary iterative winner, and this is true even if she is the initial winner:

Example 28. *Let us consider the profile $\mathscr{P} = \{x_4 \succ_1 x_3 \succ_1 x_1 \succ_1 x_2, \quad x_1 \succ_2 x_4 \succ_2 x_3 \succ_2 x_2, \quad x_3 \succ_3 x_2 \succ_3 x_4 \succ_3 x_1, \quad x_2 \succ_4 x_3 \succ_4 x_1 \succ_4 x_4, \quad x_4 \succ_5 x_1 \succ_5 x_2 \succ_5 x_3\}$ with x_4 the Condorcet and initial winner, and $PW^0 = M$. Let us exhibit a deviation sequence in which x_4 is not winning. First, voter 4 deviates from x_2 to x_3, making x_3 the current winner and $PW^1 = \{x_1, x_3, x_4\}$. Then voter 5 deviates from x_4 to x_1, yielding $PW^2 = \{x_1, x_3\}$. Since $x_4 \notin PW^2$, she cannot win in this deviation sequence and is not the necessary winner.*

On the other hand, the Condorcet winner is always guaranteed to be a possible iterative winner:

Proposition 54. *If the Condorcet winner is an intial winner of the truthful ballot b^0 (given a profile $\mathscr{P}$), then she is a possible iterative winner.*

Proof. Let c^* be the Condorcet winner of given profile $\mathscr{P}$, $c^* \in PW^0$. We show by construction that there exists a deviation sequence $\langle b^0, b^1, \ldots, b^r \rangle \in DS(\mathscr{P})$ such that $w^r = c^*$.

If $|PW^0| \leq 2$, then, by Theorem 40, c^* is a necessary and thus possible winner. Let us assume from now that $|PW^0| \geq 3$. In order to build a deviation sequence in which c^* is elected, we repeatedly use Observation 38 to rule out potential winners one by one, until we reach the situation where there are only two potential winners including c^* (hence, c^* is guaranteed to be elected). For each iteration t of the deviation sequence,

If no strategic move is possible, we are done. Let us now assume the opposite. If there exists a voter i such that $b_i^t = x \in PW^t$, and $y \in PW_i^t$ such that $y \succ_i c^*$, then i can change her ballot for x to a ballot for y, and by Observation 38, $x \notin PW^{t+1}$. Otherwise, each voter casting her ballot for a potential winner at iteration t prefers c^* to any other potential winner. Then only FNPW moves are possible. Let j be a voter such that $b_j^t = z \notin PW^t$, and $y \in PW_j^t$ such that $y \succ_j c^*$. Then j can change her ballot for z to a ballot for y. If after this FNPW move, $PW^{t+1} = \{y, c^*\}$, c^* is a necessary (and thus possible) winner. Otherwise, there exists a candidate $x \in PW^{t+1}$, and we have assumed that each voter of x prefers c^* over all the other potential winners (different from x). In particular, there is a voter k such that $b_k^{t+1} = x$ who prefers c^* to the current winner y. k will then move to c^*, and by Observation 38, $x \notin PW^{t+2}$. $\square$

4.5.2 . Condorcet Efficiency of the Iterative Rule

We have previously examined the conditions under which a Condorcet winner is a necessary or possible iterative winner. In this section, we go further by investigating how the iterative voting process affects the probability of electing the Condorcet winner.

More formally, we model iterative voting (under plurality) as a randomized voting rule, called randomized iterative plurality. Given the initial truthful score vector $s \in I_n^m$, we enumerate all possible deviation sequences and define the outcome as a probability distribution π^s over candidates, where for each $x \in M$, $\pi^s(x)$ denotes the proportion of sequences in which x is elected. Any branch has the same weight whatever its length. In particular, for a given score vector s, a candidate x is a possible iterative winner iff $\pi^s(x) > 0$, and a necessary iterative winner iff $\pi^s(x) = 1$. Note that this is not the only way to fairly weight branches, as we could also distribute the weight equally among the subbranches.

For any given voting rule, the Condorcet efficiency (CE) is defined as the probability of electing the Condorcet winner when one exists:

Definition 18 (Condorcet efficiency)**.** *When the Condorcet winner exists, we define the Condorcet efficiency as the probability to elect the Condorcet winner with respect to a voting rule.*

Note that for plurality, the Condorcet efficiency corresponds to $\mathbb{P}_C(c^* = \mathcal{W}_P(b^0) \mid c^* \text{ exists})$ while the Condorcet efficiency under randomized iterative plurality corresponds to $\mathbb{P}_C(c^* = \mathcal{W}_P(b^r) \mid c^* \text{ exists})$, for any deviation sequence $\langle b^0, \dots, b^r \rangle$. To study whether the iterative voting increases the Condorcet efficiency it remains thus to study the sign of the value $\Delta CE = \mathbb{P}_C(c^* = \mathcal{W}_P(b^r) \mid c^* \text{ exists}) - \mathbb{P}_C(c^* = \mathcal{W}_P(b^0) \mid c^* \text{ exists})$.

This question has already been studied empirically by Grandi et al. [2013]. However, it has been done for a particular turn function which arbitrarily selects the voter allowed to deviate at each step. In contrast, our proof does not assume any turn function and considers all possible deviation sequences, via randomized iterative plurality.

In practice, we draw a preference profile under a certain culture C, and denote by C^*, similarly as Gehrlein and Lepelley [2010], the culture associated with C that is reduced to preference profiles where the Condorcet winner exists.

Lemma 55. *Let C be a culture and c^* the Condorcet winner, when c^* exists, we have the following decomposition: $\Delta CE = \mathbb{P}_{C^*}(c^* = \mathcal{W}_P(b^r) \cap c^* \neq \mathcal{W}_P(b^0)) - \mathbb{P}_{C^*}(c^* \neq \mathcal{W}_P(b^r) \cap c^* = \mathcal{W}_P(b^0))$.*

Proof.

$$\begin{aligned}
\Delta CE &= \mathbb{P}_{C^*}(c^* = \mathcal{W}_P(b^r)) - \mathbb{P}_{C^*}(c^* = \mathcal{W}_P(b^0)) \\
&= \mathbb{P}_{C^*}(c^* = \mathcal{W}_P(b^r) \ \cap \ c^* = \mathcal{W}_P(b^0)) \\
&+\mathbb{P}_{C^*}(c^* = \mathcal{W}_P(b^r) \ \cap \ c^* \neq \mathcal{W}_P(b^0)) \\
&-\mathbb{P}_{C^*}(c^* = \mathcal{W}_P(b^0) \ \cap \ c^* = \mathcal{W}_P(b^r)) \\
&-\mathbb{P}_{C^*}(c^* = \mathcal{W}_P(b^0) \ \cap \ c^* \neq \mathcal{W}_P(b^r)) \\
&= \mathbb{P}_{C^*}(c^* = \mathcal{W}_P(b^r) \ \cap \ c^* \neq \mathcal{W}_P(b^0)) \\
&-\mathbb{P}_{C^*}(c^* \neq \mathcal{W}_P(b^r) \ \cap \ c^* = \mathcal{W}_P(b^0))
\end{aligned}$$

□

Lemma 55 enables us to simplify the computation of the difference in the Condorcet efficiency. We now use this result to prove the increase in Condorcet efficiency under the Impartial Anonymous Culture (IAC) assumption.

Theorem 56. *Under IAC, the iterative voting process increases the Condorcet efficiency of plurality for any m, and n sufficiently larger than m.*

Proof. To prove that $\Delta CE > 0$ whenever c^* exists, it suffices by Lemma 55 to show that

$$\begin{aligned}
&\mathbb{P}_{IAC}(c^* = \mathcal{W}_P(b^r) \cap c^* \neq \mathcal{W}_P(b^0) \mid c^* \text{ exists}) \\
&> \mathbb{P}_{IAC}(c^* \neq \mathcal{W}_P(b^r) \cap c^* = \mathcal{W}_P(b^0) \mid c^* \text{ exists})
\end{aligned}$$

To simplify the notations, we denote $\mathbb{P}_{IAC}(\cdot \mid c^* \text{ exists})$ by $\mathbb{P}_{IAC^*}(\cdot)$. Also, to shorten formulas and thus improve the readability of the proof, we use interchangeably the notations $\{|PW^0(s)| = k\}$ (resp. $|PW^0| = k$) and $s \in S^k$.

Upper bound on $\mathbb{P}_{IAC^*}(c^* \neq \mathcal{W}_P(b^r) \cap c^* = \mathcal{W}_P(b^0))$: We first note that $\{c^* \neq \mathcal{W}_P(b^r) \ \cap \ c^* = \mathcal{W}_P(b^0)\} \subset \{\cup_{k=4}^m S_n^k\}$. Indeed, if $s(b^0) \in S_n^i$ for $i \leq 2$, then by Theorem 40, $c^* = \mathcal{W}_P(b^r)$. By Theorem 53, we also have $c^* = \mathcal{W}_P(b^r)$ when $s(b^0) \in S_n^3$ and $c^* = \mathcal{W}_P(b^0)$. Therefore,

$$\mathbb{P}_{IAC^*}(c^* \neq \mathcal{W}_P(b^r) \ \cap \ c^* = \mathcal{W}_P(b^0))$$

$$\leqslant \mathbb{P}_{IAC^*}(\cup_{k=4}^m \{|PW^0| = k\}) = \sum_{k=4}^m \mathbb{P}_{IAC^*}(|PW^0| = k)$$

The last equality is obtained because $\{|PW^0| = k\}_{4 \leqslant k \leqslant m}$ is a partition.

Lower bound on $\mathbb{P}_{IAC^*}(c^* = \mathcal{W}_P(b^r) \ \cap \ c^* \neq \mathcal{W}_P(b^0))$:

We have:

$$\begin{aligned}
&\mathbb{P}_{IAC^*}(c^* = \mathcal{W}_P(b^r) \ \cap \ c^* \neq \mathcal{W}_P(b^0)) \geqslant \\
&\geqslant \mathbb{P}_{IAC^*}(c^* = \mathcal{W}_P(b^r) \ \cap \ c^* \neq \mathcal{W}_P(b^0) \ \cap \ s \in S^2) \\
&= \mathbb{P}_{IAC^*}(c^* = \mathcal{W}_P(b^r) \ \cap \ c^* \neq \mathcal{W}_P(b^0) \mid s \in S^2) \\
&\cdot \mathbb{P}_{IAC^*}(s \in S^2) \\
&\geqslant \mathbb{P}_{IAC^*}(c^* = \mathcal{W}_P(b^r) \ \cap \ c^* \neq \mathcal{W}_P(b^0) \ \cap \ c^* \in PW^0(s) \mid s \in S^2) \\
&\cdot \mathbb{P}_{IAC^*}(s \in S^2) \geqslant \\
&\geqslant \mathbb{P}_{IAC^*}(c^* = \mathcal{W}_P(b^r) \ \cap \ c^* \neq \mathcal{W}_P(b^0) \mid c^* \in PW^0(s), s \in S^2) \\
&\cdot \mathbb{P}_{IAC^*}(c^* \in PW^0(s) \mid s \in S^2) \cdot \mathbb{P}_{IAC^*}(s \in S^2)
\end{aligned}$$

Let us now look closer to the two first terms of the last product:

(i) $\mathbb{P}_{IAC^*}(c^* = \mathcal{W}_P(b^r) \ \cap \ c^* \neq \mathcal{W}_P(b^0) \mid c^* \in PW^0(s), s \in S^2)$:
As the distribution over scores is uniform under IAC, if $PW^0 = \{c, c'\}$, then each of these two candidates has the same probability to be the initial winner. In other words, $\mathbb{P}_{IAC}(c = \mathcal{W}_p(b^0)) = \mathbb{P}_{IAC}(c' = \mathcal{W}_p(b^0)) = \frac{1}{2}$. Under IAC^*, the distribution over scores is biased in favor of the Condorcet winner c^* - we have

$$\mathbb{P}_{IAC^*}(c^* \neq W_p(b^0) \mid c^* \in PW^0(s), s \in S^2) = \frac{1}{2} - \epsilon,$$

with ϵ going to 0 when n grows and m is fixed. In addition, under assumptions that $s \in S^2$ and $c^* \in PW^0(s)$, by Theorem 40, c^* is the necessary winner, thus

$$\{c^* = \mathcal{W}_P(b^r) \ \cap \ c^* \neq \mathcal{W}_P(b^0)\} = \{c^* \neq \mathcal{W}_P(b^0)\}$$

and hence

$$\begin{aligned}
&\mathbb{P}_{IAC^*}(c^* = \mathcal{W}_P(b^r) \ \cap \ c^* \neq \mathcal{W}_P(b^0) \mid c^* \in PW^0(s), s \in S^2) = \\
&= \frac{1}{2} - \epsilon.
\end{aligned}$$

(ii) $\mathbb{P}_{IAC^*}(c^* \in PW^0(s) \mid s \in S^2)$:
Again by the uniformity of scores under IAC, we have, for any m and any candidate c, for any m, $\mathbb{P}_{IAC}(c \in PW^0(s) \mid s \in S^2) = \frac{2}{m}$. Indeed, among the $\binom{m}{2}$ equally likely pairs of potential winners, c appears in $m-1$ of them. Under IAC^*, this distribution is again biased in favor of the Condorcet winner c^*, which yields

$$\mathbb{P}_{IAC^*}(c^* \in PW^0(s) \mid s \in S^2) \geqslant \frac{2}{m}$$

Put together, we get:

$$\mathbb{P}_{IAC^*}(c^* = \mathcal{W}_P(b^r) \ \cap \ c^* \neq \mathcal{W}_P(b^0))$$
$$\geqslant (\frac{1}{2} - \epsilon) \cdot \frac{2}{m} \cdot \mathbb{P}_{IAC^*}(s \in S^2)$$

Intermediate step: Implication between IAC and IAC^*: To conclude the proof, we now need to prove that:

$$\sum_{k=4}^{m} \mathbb{P}_{IAC^*}(s \in S^k) \leqslant (\frac{1}{2} - \epsilon) \cdot \frac{2}{m} \cdot \mathbb{P}_{IAC^*}(s \in S^2) \tag{4.1}$$

As working directly under the IAC^* distribution seems challenging, we will rather prove the analogous inequality under IAC:

$$\sum_{k=4}^{m} \mathbb{P}_{IAC}(s \in S^k) \leqslant (\frac{1}{2} - \epsilon) \cdot \frac{2}{m} \cdot \mathbb{P}_{IAC}(s \in S^2) \tag{4.2}$$

We can actually prove that Equation (4.2) implies Equation (4.1). Indeed, let us assume that Equation (4.2) holds. We note that

$$\mathbb{P}_{IAC}(c^* \text{ exists} \mid s \in \cup_{k \leqslant 4} S^k) \leqslant \mathbb{P}_{IAC}(c^* \text{ exists} \mid s \in S^2)$$

since the probability of Condorcet winner existence increases as the score becomes unbalanced. Therefore, we obtain

$$\sum_{k=4}^{m} \mathbb{P}_{IAC}(s \in S^k) \cdot \mathbb{P}_{IAC}(c^* \text{ exists} \mid \cup_{k \leqslant 4} S^k)$$
$$\leqslant (\frac{1}{2} - \epsilon) \cdot \frac{2}{m} \cdot \mathbb{P}_{IAC}(s \in S^2) \cdot \mathbb{P}_{IAC}(c^* \text{ exists} \mid s \in S^2)$$

Dividing by $\mathbb{P}_{IAC}(c^* \text{ exists})$, we get:

$$\frac{\sum_{k=4}^{m} \mathbb{P}_{IAC}(s \in S^k) \cdot \mathbb{P}_{IAC}(c^* \text{ exists} \mid \cup_{k \leqslant 4} S^k)}{\mathbb{P}_{IAC}(c^* \text{ exists})}$$

$$\leqslant \frac{(\frac{1}{2}-\epsilon)\cdot \frac{2}{m}\cdot \mathbb{P}_{IAC}(s\in S^2)\cdot \mathbb{P}_{IAC}(c^* \text{ exists} \mid s\in S^2)}{\mathbb{P}_{IAC}(c^* \text{ exists})}$$

By the conditional Bayes' formula, we end up having:

$$\sum_{k=4}^{m} \mathbb{P}_{IAC}(s\in S^k \mid c^* \text{ exists})$$
$$\leqslant (\frac{1}{2}-\epsilon)\cdot \frac{2}{m}\cdot \mathbb{P}_{IAC}(s\in S^2 \mid c^* \text{ exists})$$

which is nothing but Equation (4.1):

Putting the bounds together under IAC: It remains to prove that Equation (4.2) holds. Using Lemma 46, we get:

$$\frac{\sum_{k=0}^{m-4}(m-k)\cdot \binom{n+k-2}{k}}{\binom{n+m-1}{m-1}}$$
$$\leqslant (\frac{1}{2}-\epsilon)\cdot \frac{2}{m}\frac{2\cdot \binom{n+m-4}{m-2}}{\binom{n+m-1}{m-1}}$$

After some algebraic simplifications, using the identity

$$k\cdot \binom{n+k-2}{k} = (n-1)\cdot \binom{n+k-2}{k-1},$$

and performing a change of variable, we obtain:

$$m\cdot \sum_{k=0}^{m-4}\binom{n+k-2}{k} - (n-1)\cdot \sum_{k=0}^{m-5}\binom{n+k-1}{k}$$
$$\leqslant (\frac{1}{2}-\epsilon)\cdot \frac{2}{m}\frac{2\cdot \binom{n+m-4}{m-2}}{m}$$

Using the following inequality (that can be easily proven by mathematical induction)

$$\sum_{k=0}^{M}\binom{A+k}{k} = \binom{A+M+1}{M}$$

with $A=n-2$ and $M=m-4$ for the first sum and $A=n-1$ and $M=m-5$ for the second, we get:

$$m\cdot \binom{n+m-5}{m-4} - (n-1)\cdot \binom{n+m-5}{m-5}$$
$$\leqslant (\frac{1}{2}-\epsilon)\cdot \frac{2}{m}\frac{2\cdot \binom{n+m-4}{m-2}}{m}$$

In other words:

$$m^2 \cdot \binom{n+m-5}{m-4} - m \cdot (n-1) \cdot \binom{n+m-5}{m-5}$$

$$\leqslant (\frac{1}{2} - \epsilon) \cdot \frac{2}{m} 2 \cdot \binom{n+m-4}{m-2}$$

If we increase n, we will see that the inequality has to become true at some point. Indeed, ϵ becomes small as n grows and for m fixed and n large enough the left hand side becomes negative thanks to the second term while the second is increasing in n and is positive. □

We now state the analogous result under impartial culture (IC).

Theorem 57. *Under IC, the iterative voting process increases the Condorcet efficiency of plurality for any m, and n sufficiently larger than m.*

Proof. Following the same steps as in the proof of Theorem 56 but for IC, it remains to show:

$$\sum_{k=4}^{m} \mathbb{P}_{IC}(s \in S^k) \leqslant (\frac{1}{2} - \epsilon) \cdot \frac{2}{m} \cdot \mathbb{P}_{IC}(s \in S^2) \tag{4.3}$$

Since ϵ is going to 0 when n is large then we can just remove it.

To prove Equation (4.3), we first prove the case of $m = 4$, and then we generalize its idea to $m > 4$.

Case of $m = 4$:

We need to prove that

$$\mathbb{P}_{IC}(s \in S^4) \leqslant \frac{1}{4} \cdot \mathbb{P}_{IC}(s \in S^2) \tag{4.4}$$

Let us denote by $S^{4 \to 2}$ the set of scores with 2 potential winners obtained from some score of S^4 by transferring at most two votes between candidates. More formally, $S^{4 \to 2} = \{s \in S^2 | \exists s' \in S^4$ such that s differs from s' in 2 votes$\}$. Also, for $s' \in S^4$, we denote by $S^{4 \to 2}(s')$ all scores de $S^{4 \to 2}$ built from s', ie., $S^{4 \to 2}(s') = \{s \in S^2 | s$ differs from s' in 2 votes$\}$. To prove Equation (4.4), it is sufficient to prove that for each score $s \in S^4$, there exists a function $f^4 : S^4 \longrightarrow [S^{4 \to 2}]^8$ association each score $s \in S^4$ with 8 different scores from $S^{4 \to 2}(s)$ in a way that:

- $\forall s' \in f^4(s), \mathbb{P}_{IC}(s') \geqslant \frac{1}{2} \mathbb{P}_{IC}(s)$
- for each couple $s, s' \in S^4$, $f^4(s) \cap f^4(s') = \emptyset$.

We define below the function f^4. Let $s^4 \in S^4$. It remains to find 8 scores of $S^{4\to 2}(s)$ such that for each $s^2 \in S^{4\to 2}(s^4)$, $\frac{\mathbb{P}_{IC}(s^2)}{\mathbb{P}_{IC}(s^4)} \geqslant \frac{1}{2}$. We define f^4 so that all scores of f^4 are of the two following types:

- Type 1: $s^2 \in f^4(s^4)$ was built from s^4 by transferring two votes from a unique candidate j to two different candidates i and k.
- Type 2: $s^2 \in f^4(s^4)$ was built from s^4 by transferring one vote from candidates j, l to two remaining candidates i, k

We denote $s^4 = (s_1^4, s_2^4, s_3^4, s_4^4), s^4 \in S_n^4$, and $s^2 = (s_1^2, s_2^2, s_3^2, s_4^2), s^2 \in S_n^2$, and we have We have:

$$\mathbb{P}_{IC}(s^4) = \frac{n!}{s_1^4! \cdot s_2^4! \cdot s_3^4! \cdot s_4^4!}(\frac{1}{m})^n$$

and

$$\mathbb{P}_{IC}(s^2) = \frac{n!}{s_1^2! \cdot s_2^2! \cdot s_3^2! \cdot s_4^2!}(\frac{1}{m})^n.$$

Let us show that for each of these types, we have $\frac{\mathbb{P}_{IC}(s^2)}{\mathbb{P}_{IC}(s^4)} \geqslant \frac{1}{2}$ for n sufficiently large.

- Type 1: in all cases where we don't change the winner (resp. the winner changes), $|s_j^4 - s_i^4| \leq 1$ (resp. $|s_j^4 - s_i^4| \leq 2$) and $s_k^4 \geq q-1$ for each $k \in \{1,2,3,4\}$.

  Then we get:

  $$\frac{\mathbb{P}_{IC}(s^2)}{\mathbb{P}_{IC}(s^4)} = \frac{s_j^4(s_j^4 - 1)}{(s_i^4 + 1)(s_k^4 + 1)}.$$

  The smallest ratio is reached when $s_j^4 = q-1$, $s_i^4 = q$ and $s_k^4 = q$ if we don't change the winner and for $s_j^4 = q+2$, $s_i^4 = q$ and $s_k^4 = q+1$ otherwise.

  Therefore,

  $$\frac{\mathbb{P}_{IC}(s^2)}{\mathbb{P}_{IC}(s^4)} \geqslant \frac{(q-1)(q-2)}{(q+1)^2}$$

  We find this ratio is greater than $\frac{1}{2}$ for $q \geq 8$, i.e., $n \geq 32$.

- Type 2: two votes are transferred from two different candidates j, l to two different candidates i, k. We get

  $$\frac{\mathbb{P}_{IC}(s^2)}{\mathbb{P}_{IC}(s^4)} = \frac{(s_j^4 - 1)(s_l^4 - 1)}{(s_i^4 + 1)(s_k^4 + 1)}.$$

  The same as in the previous cas, $s_p^4 \geq q-1$ for each $p \in \{1,2,3,4\}$, and for each $p, p' \in \{1,2,3,4\}$, we have $|s_p^4 - s_{p'}^4| \leq 2$. Therefore,

  $$\frac{\mathbb{P}_{IC}(s^2)}{\mathbb{P}_{IC}(s^4)} \geqslant \frac{(q-2)^2}{(q+2)^2},$$

which is greater than $\frac{1}{2}$ for $q \geq 12$, ie., $n \geq 48$.

We will now build 8 scores of $f(s^4)$ as follows:

- We will create 5 scores of type 1, by distinguishing three sub-types:
  - 2 scores where the winner of s^4 gets one more vote and the looser of s^4 is not modified. We can then choose arbitrary which of the two remaining candidates x and y will get one more vote, and which one will loose two votes - indeed, each of these both choices yields a score of two potential winners, namely the winner of s^4 and the other candidate that gets one vote.
  - 1 scores where the the winner of s^4 gets one more vote, the third-ranked candidate of s^4 is not modified, and the second-ranked candidate looses two votes.
  - 2 scores vectors where the winner looses two votes. The candidate that is not modified needs to be the third or fourth ranked candidate of s^4 in order to ensure that the resulting score has two potential winners.
- Finally, we will create 3 more scores of type 2. The winner can not loose a vote because, depending of the number of votes of remaining candidates, we might reach a score with 3 potential winners. Therefore, the winner will get one more vote, and we have 3 choices for the second candidate to get one more vote, each of these yielding a score with 2 potential winners.

Moreover, all scores built by this construction are different, i.e $|f^4(s)| = 8$. Indeed, when starting from the same winner and adding one to her then the subtraction part differentiates the score of case 1 and 2.

The last thing to check is that $f^4(s) \cap f^4(s') = \emptyset$. It is easy to see that for all cases where the winner gets one more vote (and in particular remains the winner), we can not have duplicates. Indeed, every score in S_n^4 yields a different winner, so winners will also be different in new scores. In the case we allow the winner to lose points (only for type 1) then this candidate who is now outside the potential winner set is last and characterized different new scores also.

General case: $m > 4$

We now explain how the construction of f^4 can be generalized for any $m > 5$. Let $m = 5$, for the case where the number of potential winners is 5, we can apply the same reasoning and we will have more cases to enumerate. For instance, for type 1, there is one more candidate that can loose 2 votes.

Therefore, we can apply exactly the same idea of transformation as previously, namely $f^5 : S^5 \longrightarrow [S^{5\to 2}]^h$, where $h > 8$ and $f^5(s) \cap f^4(s') = \emptyset$ because we start from different scores.

Let us show that $h - 8 \geq 2$ to preserve our probability ratio greater than $\frac{1}{2}$. Indeed, in case 2 we are able to build 3 more scores by taking a point to the fifth candidate, i.e., we now have $\binom{3}{2} = 3$ choices to subtract a point to two candidates. The case where this candidate has no vote can be treated separately. We see that these scores do not intersect. By recurrence, we apply the same reasoning when considering one more candidate and see that all the difficulty remains in the case of four potential winners for any m. □

These results provide formal guarantees that Condorcet efficiency increases under impartial cultures.

4.6 . Conclusion and Future Works

4.6.1 . Conclusion

In this chapter, we have examined the outcomes of iterative voting for the plurality rule under different aspects. We have particularly investigated the potential diversity of outcomes via the concepts of possible and necessary iterative winners. Although we may find instances where all candidates can be elected in some sequence of voters' deviations, we have experimentally seen that this scenario rarely occurs. Indeed, the most frequent situations are when a few different candidates turn out to be possible iterative winners. This is partly due to the existence of a necessary iterative winner, an event which is itself "biased" by the extreme scenario where no deviation is initially possible. We show that this extreme situation actually often occurs in our setting under impartial (anonymous) cultures.

In a computational point of view, the existence problem for a possible iterative winner is harder than for the necessary variant. It shows in a way that the kind of robustness created by the election of the same candidate at every sequence is easily detectable while more fluctuating scenarios are difficult to predict. Beyond quantitative or computational results on possible outcomes, our analysis also helps provide theoretical insights on how beneficial manipulation can be. Indeed, we show that the frequency of election of the Condorcet winner is increased, when considering all possible iterative sequences with equal weights, under impartial (anonymous) cultures, compared to the single outcome of the initial plurality rule. This confirms and generalizes previous observations that were only made experimentally.

4.6.2 . Future Works

Our work opens several avenues for future work.

- While we have focused on a specific iterative voting setting, one could examine the impact of other types of strategic behaviors. Indeed, one of the main weaknesses of this model is that a voter deviates only when she is pivotal. A useful way to address this limitation is to consider a more general model parameterized by the number of voters required to be pivotal. The idea is that voters may change their ballot when they are close to being pivotal [Wilczynski, 2019], either in an absolute sense, such as the number of votes, or in a relative sense, such as the proportion of voters in the election. Thus, one could examine the same type of question under this strategic behavior, for example.

- Another direction for generalizing this work is to consider different voting rules. For instance, Condorcet efficiency is the same for all Condorcet-consistent rules, but one could ask how close positional scoring rules are to being Condorcet-consistent. Nevertheless, one possible difficulty is to deal with potential problems of convergence [Meir, 2018].

- Another natural direct extension would be to consider other–more realistic–voting cultures for probabilistic analyses, such as single-peaked ones, Mallows distributions, or even Polya-Eggenberger urns [Boehmer et al., 2024]. Following the ideas developed in Chapter 3, it would not be surprising if adding more structure to the cultures led to a decrease in the diversity of winners.

- Finally, another more subtle study would be to analyze the strategic power of the voters (or their coalitions) on the iterative outcome, with respect to their position of deviation in the sequence or their preferences. More precisely, this could take the form of a Shapley-Shubik index adapted to our setting of sequence-dependent winners [Shapley and Shubik, 1954].

In this chapter, we study how strategic voting affects plurality elections by analyzing the variability of outcomes under impartial cultures, investigating the complexity of determining possible and necessary winners, and examining the impact on Condorcet efficiency. This work represents a first step toward a better understanding of strategic voting, highlighting also the crucial role of information. Our next question is whether controlling the dissemination of such information can confer power to the one who broadcast it. We address this issue in Chapter 5.

5 - The Influence of Poll Manipulation on Elections Outcomes

Abstract

In this chapter, we consider the problem of poll manipulation in political elections. In the context of strategic voting, we are interested in whether a polling institute can manipulate the information it communicates to voters in order to influence the outcome of the election. We start with a version of the problem where the polling institute is allowed to send any score to voters. Then, for realistic reasons, we investigate a restricted version in which the polling institute cannot announce scores which are too far from the truthful ones. While we show that both decision problems are computationally hard, we go beyond this worst-case complexity analysis by using probabilistic tools to address the possibility of successful and efficient manipulation in practice, with respect to several natural preference distributions.

Résumé

Dans ce chapitre, nous étudions le problème de la manipulation des sondages dans les élections politiques. Dans le contexte du vote stratégique, nous nous intéressons à la question de savoir si un institut de sondage peut manipuler les informations qu'il communique aux électeurs afin d'influencer le résultat de l'élection. Nous commençons par une version du problème dans laquelle l'institut est autorisé à envoyer n'importe quel score aux électeurs. Puis, pour des raisons de réalisme, nous considérons une version restreinte dans laquelle l'institut ne peut pas annoncer des scores trop éloignés des scores réels. Bien que nous montrons que ces deux problèmes de décision sont computationnellement difficiles, nous allons au-delà de cette analyse en pire cas en utilisant des outils probabilistes pour étudier la possibilité d'une manipulation efficace et réussie en pratique, par rapport à plusieurs distributions naturelles de préférences.

Much of the content of this chapter is based on a paper co-authored with Vincent Mousseau and Anaëlle Wilczynski, which was accepted at the 27th European Conference on Artificial Intelligence (ECAI 2024) [Mousseau et al., 2024].

5.1 . Introduction

In this chapter, we study another source of outcomes' variability arising from poll information through strategic voting.Indeed, the question of the information available to the voters is key and has a strong impact on the manipulability of voting processes [Endriss et al., 2016; Reijngoud and Endriss, 2012], asking the question of the power granted to those who disseminate it, particularly polling institutes. To deal with partial information in voting, one can naturally follow a Bayesian approach by considering a probability distribution over a set of possible preference orders for other voters [Myerson and Weber, 1993; Hazon et al., 2012]. Alternatively, a set of possible preference profiles can be derived from partial votes [Conitzer et al., 2011; Dey et al., 2018] or from a given maximum distance to the voters' actual preferences [Anand and Dey, 2021]. Another possibility is to assume local information for the voters, which is captured by a social network [Grandi, 2017]. Finally, an aggregated global information coming from opinion polls can be communicated to voters [Baumeister et al., 2020; Endriss et al., 2016; Reijngoud and Endriss, 2012; Wilczynski, 2019].

Following this latter line of research, in this chapter, inspired by political elections, we assume that voters receive only a global information about the voting intentions within the population, which is communicated through *opinion polls*. Voters trust the information communicated in the polls and compute their best response ballot on the basis of this information. This confidence in the polls grants an important power to the polling institute which disseminates it, raising the natural question of *poll manipulation*. Indeed, a polling institute might have its own interests in the election and try to orient votes toward them. This problem is close to the question of election control [Faliszewski and Rothe, 2016], where an external agent aims to alter the outcome of the election, but here no structural change is made on the election. Indeed, election control examines this issue from a structural perspective, investigating whether the winner can be modified through the addition or deletion of candidates or voters.

In the line of seminal works analyzing the complexity of voter manipulation [Bartholdi et al., 1989], one can analyze the complexity of the poll manipulation problem. However, computational intractability may not constitute a relevant barrier to manipulation, as it relies on worst-case analysis [Faliszewski and Procaccia, 2010]. Therefore, to complement complexity results, an average-case study using a probabilistic approach is relevant, as it has been widely investigated for voter manipulation (see, e.g., [Friedgut et al., 2008; Isaksson et al., 2012; Procaccia and Rosenschein, 2007; Xia and Conitzer, 2008]). In particular, the asymptotic study is meaningful since political elections are characterized by a large number of voters. Considering election control problems, as far as we know, this approach has been surprisingly neglected. A notable exception is a recent work by Xia [2023] which investigates the likelihood of manipulability for several coalition influence problems, including control by adding or deleting votes. Up to our best knowledge, no such study has been conducted so far for the poll manipulation problem. In this chapter, we study the constructive poll manipulation problem where the polling institute wishes to favor a specific candidate by broadcasting manipulated candidates' scores. This problem has been introduced by Wilczynski [2019] and further extended by Baumeister et al. [2020], who also consider the destructive variant where the polling institute aims to prevent the election of a given candidate. While both works consider a framework where voters are embedded in a social network and analyze the complexity of the problem with respect to the structure of the graph, we consider a simpler model with no social network, which clarifies the role of the opinion polls. In particular, we analyze the following two versions of the problem. In the unrestricted problem, the polling institute is free to send any score information. The restricted problem considers a more realistic context where only score information that would be close enough to truthful scores are allowed. The idea for this second problem is for the polling institute to lie in a reasonable manner, by submitting realistic scores, not too far from a ground truth that may correspond to the results of a past election, or another poll. Such restrictions help to gain the trust and confidence from the voters. We prove that both versions of the problem are computationally hard, answering an open question from Baumeister et al. [2020], but also analyze the probability of existence of a successful and efficiently computable poll manipulation. For this latter purpose, we introduce a natural condition on statistical cultures which is satisfied by most natural preference distributions [Szufa et al., 2020]. In fact, we exhibit a simple heuristic and prove its success for the unrestricted problem, which means that, without restriction, the polling institute can almost always efficiently influence large elections. For the restricted manipulation problem, we prove that if the allowed distance is negligible with respect to the number of voters, then no manipulation is possible. However, when this distance becomes significant, e.g., when it is a fixed proportion of the number of voters, easy manipulation is almost always successful in large elections. Finally, we show that most results still hold when assuming a more general strategic behavior for voters [Wilczynski, 2019].

5.2 . A Poll Manipulation Problem

We consider the same iterative voting model as presented in Chapter 2 and used in Chapter 4, namely the one introduced by Meir et al. [2010]. Since the goal is to study strategic voting in large political elections, we naturally assume that $n > m > 2$ (by the Gibbard-Satterthwaite theorem, voting rules are susceptible to manipulation only when there are more than two candidates).

However, there is a slight difference in this model: we assume that the information (i.e., the poll) is given at the beginning, after which the strategic process unfolds and the winner is elected. There are no successive updates of poll information as in the previous model. Each voter can deviate at most once since she only gets the information about the scores s provided by the polling institute, and cannot see the deviations from other voters (thus the order of voters' deviations does not matter, they could even be simultaneous). Hence, the deviation process ends after at most n steps; it is then obviously converging and converges to a final ballot profile denoted by b^s. In this model, a polling institute sends out a score at the beginning of the process and then each voter votes strategically w.r.t. that information, and finally the winner of the election is computed.

We recall some notations presented in Chapter 2 to facilitate the reading. A candidate y is a potential winner for voter i, at a given step where the current score vector is s, if i believes that voting for y will make candidate y the new winner, i.e., $s^{-i}_{\mathcal{W}_P(s^{-i})} - s^{-i}_y + \mathbb{1}_{\mathcal{W}_P(s^{-i}) \rhd y} \leqslant 1$, where s^{-i} denotes the score vector s without counting the current ballot b_i of voter i. Let PW^t_i denote the set of potential winners for voter i at step t, and PW^t the set of all potential winners at step t, i.e., $PW^t := \bigcup_{i \in N} PW^t_i$.

We then want to describe the behavior of the polling institute who may have its own interest in the election. Let x^* be the target candidate of the polling institute, i.e., it wants x^* to be elected. Let I be the space of all possible scores that the polling institute can announce, i.e., $I := \{s \in \mathbb{N}^m \mid \sum_{j=1}^m s_j = n\}$.

We consider the following poll manipulation problem by the polling institute:

Unrestricted manipulation problem	
Instance:	Election $(N, M, \mathscr{P}, \rhd)$, target candidate $x^* \in M$
Question:	Does there exist a score $s \in I$ to announce such that $\mathcal{W}_P(b^s) = x^*$?

However, the fact that the polling institute is allowed to send any score is not very realistic. We use a restricted version of the decision problem where the distance between the truthful poll and the one sent by the polling institute is bounded. We use the number of vote changes to evaluate the distance between possible scores, i.e., $d(s, s') = \frac{1}{2} \sum_{j \in M} |s_j - s'_j|$, for every scores $s, s' \in I$. Note that this distance is equivalent to the restriction of the ℓ_1 distance on I divided by 2 and sometimes called the "earth mover distance" in the literature [Meir, 2018]. We let $I_k := \{s \in \mathbb{N}^m \mid d(s, s^T) \leq k$ and $\sum_{j \in M} s_j = n\}$ be the restricted space of action of the polling institute. We then analyze the following poll manipulation problem by the polling institute:

Restricted manipulation problem	
Instance:	Election $(N, M, \mathscr{P}, \rhd)$, target candidate $x^* \in M$, integer k
Question:	Does there exist a score $s \in I_k$ to announce such that $\mathcal{W}_P(b^s) = x^*$?

A poll manipulation is illustrated in the next example.

Example 29. *Let us consider an election* $(N, M, \mathscr{P}, \rhd)$ *where* $N = \{1, \ldots, 8\}$, $M = \{x_1, x_2, x_3, x_4, x_5\}$*, the tie-breaking* $\rhd$ *follows the lexicographic order and the preference profile* $\mathscr{P}$ *is as follows:*

1:	$x_1 \succ x_3 \succ x_4 \succ x_5 \succ x_2$	*5:*	$x_2 \succ x_4 \succ x_5 \succ x_3 \succ x_1$
2:	$x_5 \succ x_2 \succ x_4 \succ x_3 \succ x_1$	*6:*	$x_2 \succ x_4 \succ x_5 \succ x_3 \succ x_1$
3:	$x_1 \succ x_4 \succ x_2 \succ x_3 \succ x_5$	*7:*	$x_3 \succ x_4 \succ x_1 \succ x_2 \succ x_5$
4:	$x_1 \succ x_4 \succ x_2 \succ x_3 \succ x_5$	*8:*	$x_4 \succ x_5 \succ x_1 \succ x_2 \succ x_3$

The initial truthful scores are given by $s^0 = (3, 2, 1, 1, 1)$*. If the score communicated by the polling institute is the truthful one, then no voter has an incentive to deviate and* x_1 *remains the winner.*

Suppose that the polling institute communicates the following score vector $s^M = (0, 2, 2, 3, 1)$*, at distance 3 to the truthful one. The set of potential winners w.r.t.* s^M *is equal to* $PW_i^{s^M} = \{x_2, x_3, x_4\}$ *for every voter* $i \in \{1, 2, 3, 4, 8\}$*, while* $PW_5^{s^M} = PW_6^{s^M} = \{x_3, x_4\}$ *and* $PW_7^{s^M} = \{x_2, x_4\}$*. Voters 3, 4, 5, 6, 7, and 8 do not have an incentive to deviate since the announced winner* x_4 *is their most preferred candidate among the potential winners. However, voters 1 and 2 have an incentive to deviate to a ballot supporting* x_3 *and* x_2*, respectively. After their deviations, we reach the final scores* $s^2 = (2, 3, 2, 1, 0)$ *where* x_2 *is the winner. Hence, the polling institute can enforce the election of* x_2*, whereas* x_1 *would remain the winner without poll manipulation.*

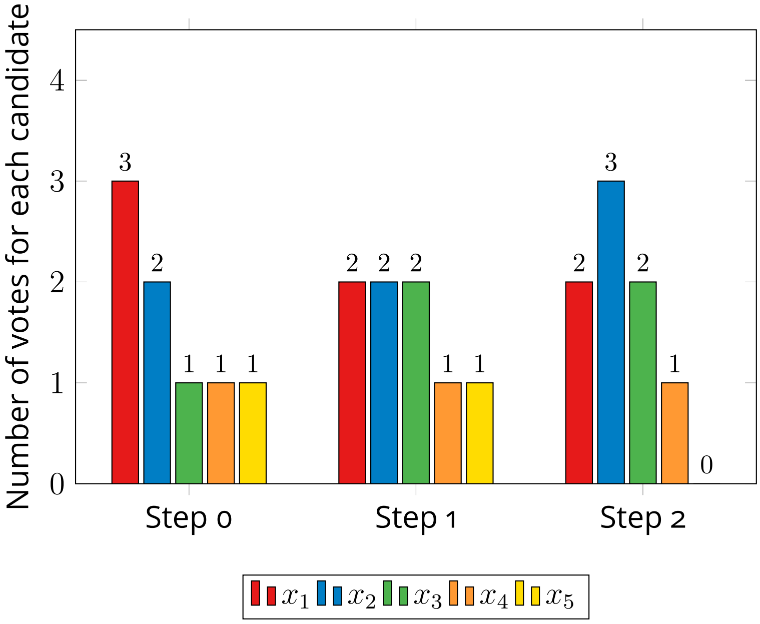

Figure 5.1: Example of strategic moves for the manipulation score $s^M = (0, 2, 2, 3, 1)$.

5.3 . Balanced Culture

We have already presented some cultures in Section 2.6, in particular impartial culture in Definition 4, Mallows culture in Definition 9, Walsh's culture in Definition 7 and Conitzer's culture in Definition 8. We will now show that most of them satisfy a general condition on cultures, which is very useful for the purpose of our chapter. We use $\mathbb{P}_C(a \succ_i b)$ when it is clear from the context instead of $\mathbb{P}_C(\succ_i \mid a \succ_i b)$. In the following of this chapter, we consider independent and identical drawings of voters' preferences such that we can either look at the distribution $C(n, \Pi^m_{sub})$ as a whole object or n drawings of preferences $\succ_i$. For technical reasons, we assume that the considered culture assigns positive probability to ranking more than two different candidates in the top positions. Note that this assumption is also natural since we focus on strategic voting and manipulation only occurs with at least three candidates [Gibbard, 1973; Satterthwaite, 1975].

We introduce below a simple property on cultures which will be key in the poll manipulation analysis.

Definition 19 (Balanced culture)**.** *A distribution $C(n, \Pi^m)$ is said to be balanced for a given candidate $c \in M$ if there exists another sufficiently worst candidate $\ell \in M \setminus \{c\}$, in the sense that $\mathbb{P}_C(c \succ_i \ell) \geq \frac{1}{2}$. The set of such candidates ℓ for x is denoted by $B_C(x)$. In general, a distribution C is said to be balanced if it is balanced for every candidate $c \in M$.*

It turns out that all cultures that we consider are balanced.

Proposition 58. *The impartial culture is balanced.*

Proof. Let $i \in N$ be a voter. The impartial culture is balanced for every candidate because for any pair of candidates x and y, we have $\mathbb{P}_{IC}(x \succ_i y) = \mathbb{P}_{IC}(y \succ_i x) = \frac{1}{2}$ since each preference order in Π^m has the same probability to be drawn. □

For a given axis $>$ on M, let $e_1^>$ and $e_2^>$ denote the two extreme candidates of $>$.

Proposition 59. *If $x \in M \setminus \{e_1^>, e_2^>\}$, then every single-peaked culture $C(n, \Pi_>^m)$ is balanced for x. If $x \in \{e_1^>, e_2^>\}$, then every single-peaked culture $C(n, \Pi_>^m)$ which also satisfies $\mathbb{P}_C(\{\succ_i | worst_{\succ_i} = x\}) \leqslant \frac{1}{2}$, is balanced for x.*

Proof. Let $>$ be an axis on M. Observe first that, by definition, every preference order $\succ_i$ which is single-peaked w.r.t. $>$, must rank last an extreme candidate of $>$. Therefore, for every single-peaked culture $C(n, \Pi_>^m)$, we must have $\mathbb{P}_C(\{\succ_i \mid worst_{\succ_i} = e_1^>\} \cup \{\succ_i \mid worst_{\succ_i} = e_2^>\}) = 1$. Moreover, since $\mathbb{P}_C(\{\succ_i \mid worst_{\succ_i} = e_1^>\} \cup \{\succ_i \mid worst_{\succ_i} = e_2^>\}) \leqslant \mathbb{P}_C(\{\succ_i \mid worst_{\succ_i} = e_1^>\}) + \mathbb{P}_C(\{\succ_i \mid worst_{\succ_i} = e_2^>\})$, this implies that there exists an extreme candidate $e_\ell^>$, for $\ell \in [2]$, such that $\mathbb{P}_C(\{\succ_i \mid worst_{\succ_i} = e_\ell^>\}) \geqslant \frac{1}{2}$. It follows that, for every candidate $x \in M \setminus \{e_1^>, e_2^>\}$, $\mathbb{P}_C(x \succ_i e_\ell^>) \geqslant \frac{1}{2}$, proving the first part of the statement. Consider now a candidate $e_\ell^>$ for $\ell \in [2]$. Assuming that $\mathbb{P}_C(\{\succ_i | worst_{\succ_i} = e_\ell^>\}) \leqslant \frac{1}{2}$, implies that $\mathbb{P}_C(\{\succ_i | worst_{\succ_i} = e_{3-\ell}^>\}) \geqslant \frac{1}{2}$ and thus $\mathbb{P}_C(e_\ell^> \succ_i e_{3-\ell}^>) \geqslant \frac{1}{2}$, proving the second part. □

In particular, the previous proposition shows that both Walsh's [2015] and Conitzer's [2007] cultures are balanced.

Proposition 60. *Any Mallows culture $\mathcal{M}^{\phi,\sigma}$ is balanced for every candidate $x \in M \setminus \{worst_\sigma\}$.*

Proof. Consider any candidate $x \in M \setminus \{worst_\sigma\}$ and the candidate $\ell := worst_\sigma$. Let $\Pi_{y \succ z}^m$ denote the set of all preferences orders where y is ranked before z, i.e., $\Pi_{y \succ z}^m := \{\succ_i \in \Pi^m : y \succ_i z\}$.

Consider the bijection $\tau : \Pi_{\ell \succ x}^m \to \Pi_{x \succ \ell}^m$, where for every preference order $\succ_i \in \Pi_{\ell \succ x}^m$, we construct the preference order $\tau(\succ_i) \in \Pi_{x \succ \ell}^m$ which is the same as $\succ_i$ except that the positions of x and ℓ are swapped. We will show that $\mathbb{P}_{\mathcal{M}^{\phi,\sigma}}(\tau(\succ_i)) \geq \mathbb{P}_{\mathcal{M}^{\phi,\sigma}}(\succ_i)$ for every $\succ_i \in \Pi_{\ell \succ x}^m$. For this purpose, we will show that $d_{KT}(\sigma, \succ_i) \geq d_{KT}(\sigma, \tau(\succ_i))$, by analyzing the differences between $\succ_i$ and $\tau(\succ_i)$ in terms of agreement on pairwise comparisons with σ.

By definition, for any arbitrary preference order $\succ_i'$, we have that:

$$d_{KT}(\sigma, \succ_i') = d_{KT}([\sigma]_{|M \setminus \{x,\ell\}}, [\succ_i']_{|M \setminus \{x,\ell\}}) + |\{y \in M : \ell \succ_i' y\}|$$

$$+|\{y \in M \setminus \{\ell\} : x \succ_i' y \text{ and } y\sigma x\}| + |\{y \in M \setminus \{\ell\} : y \succ_i' x \text{ and } x\sigma y\}|$$

where $[\succ_i'']_{|Y}$ denotes the restriction of the preference order $\succ_i''$ on $Y \subseteq M$. Observe that, by construction, for any $\succ_i \in \Pi^m_{\ell \succ x}$, $\succ_i$ and $\tau(\succ_i)$ agree on all pairwise comparisons within $M \setminus \{\ell, x\}$.

Therefore, we have:

$$d_{KT}([\sigma]_{|M\setminus\{x,\ell\}}, [\succ_i]_{|M\setminus\{x,\ell\}}) = d_{KT}([\sigma]_{|M\setminus\{x,\ell\}}, [\tau(\succ_i)]_{|M\setminus\{x,\ell\}})$$

Moreover, by construction:

- For all candidates y such that $\ell\, \tau(\succ_i)\, y$ it implies that $\ell \succ_i y$,
- For all candidates $y \in M \setminus \{\ell\}$ such that $x \succ_i y$ it implies that $x\, \tau(\succ_i)\, y$,
- For all candidates $y \in M \setminus \{\ell\}$ such that $y\, \tau(\succ_i)\, x$ it implies that $y \succ_i x$.

It follows that we have:

$$d_{KT}(\sigma, \succ_i) - d_{KT}(\sigma, \tau(\succ_i)) = |\{y \in M : \ell \succ_i y\, \tau(\succ_i)\, \ell\}|$$

$$-|\{y \in M\setminus\{\ell\} : y\sigma x \text{ and } x\, \tau(\succ_i)\, y \succ_i x\}| + |\{y \in M\setminus\{\ell\} : x\sigma y \text{ and } x\, \tau(\succ_i)\, y \succ_i x\}|$$

By construction, it holds that

$$|\{y \in M : \ell \succ_i y\, \tau(\succ_i)\, \ell\}| = r_{\succ_i}(x) - r_{\succ_i}(\ell)$$

Moreover,

$$|\{y \in M \setminus \{\ell\} : y\sigma x \text{ and } x\, \tau(\succ_i)\, y \succ_i x\}|$$

$$+|\{y \in M \setminus \{\ell\} : x\sigma y \text{ and } x\, \tau(\succ_i)\, y \succ_i x\}| = r_{\succ_i}(x) - r_{\succ_i}(\ell) - 1$$

which implies that

$$-(r_{\succ_i}(x) - r_{\succ_i}(\ell) - 1) \leqslant -|\{y \in M \setminus \{\ell\} : y\sigma x \text{ and } x\, \tau(\succ_i)\, y \succ_i x\}|$$

$$+|\{y \in M \setminus \{\ell\} : x\sigma y \text{ and } x\, \tau(\succ_i)\, y \succ_i x\}| \leqslant r_{\succ_i}(x) - r_{\succ_i}(\ell) - 1$$

Therefore, in total, we have

$$1 \leq d_{KT}(\sigma, \succ_i) - d_{KT}(\sigma, \tau(\succ_i)) \leq 2(r_{\succ_i}(x) - r_{\succ_i}(\ell)) - 1$$

and thus $d_{KT}(\sigma, \succ_i) \geq d_{KT}(\sigma, \tau(\succ_i))$. By definition of the Mallows culture $\mathcal{M}^{\phi,\sigma}$, we thus have $\mathbb{P}_{\mathcal{M}^{\phi,\sigma}}(\tau(\succ_i)) \geq \mathbb{P}_{\mathcal{M}^{\phi,\sigma}}(\succ_i)$. Hence, we conclude that $\mathbb{P}_{\mathcal{M}^{\phi,\sigma}}(\{\succ_i' : x \succ_i' \ell\}) \geq \mathbb{P}_{\mathcal{M}^{\phi,\sigma}}(\{\succ_i' : \ell \succ_i' x\})$, and thus $\mathbb{P}_{\mathcal{M}^{\phi,\sigma}}(\{\succ_i' : x \succ_i' \ell\}) \geq \frac{1}{2}$, implying that $\mathcal{M}^{\phi,\sigma}$ is balanced for x. □

5.4 . The Unrestricted Poll Manipulation Problem

This section is devoted to the study of the unrestricted manipulation problem where the polling institute can send any score in I. We first give some results on the computational complexity of the problem then we continue our work with a probabilistic approach of the problem to capture what can happen in practice.

We first prove that, even in the unrestricted case, the poll manipulation problem is NP-complete. Our result answers an open question from Baumeister et al. [2020].

Theorem 61. *The unrestricted manipulation problem is NP-complete.*

Proof. Membership to NP is straightforward: given communicated scores, we can efficiently derive the possible unique deviation of each voter and compute the winner in the deviating profile.

For hardness, we perform a reduction from a variant of Exact Cover by 3-Sets (X3C) known to be NP-complete [Garey and Johnson, 1979]: Given a set $X = \{x_1, x_2, \ldots, x_{3q}\}$ and a set $S = \{S_1, S_2, \ldots, S_{3q}\}$ of 3-element subsets of X, where each element x_i occurs in exactly three subsets of S, we ask whether there exists an exact cover, i.e.,., a subset $S' \subseteq S$ such that every element of X occurs in exactly one member of S'. From an instance (X, S) of X3C, we construct an instance of our unrestricted manipulation problem as follows. For each element x_i, for $i \in [3q]$, we create a candidate y_i, and for each subset S_j where $j \in [3q]$, we create a candidate c_j. We add three candidates w, z, and t where t is our target candidate.

There are $12q + 7$ voters: for each element x_i, for $i \in [3q]$, we create one voter Y_i, for each subset S_j, for $j \in [3q]$, we create three voters C_j^r where $r \in [3]$, and we finally add two voters T^ℓ, two voters Z^ℓ, two voters W^ℓ for $\ell \in [2]$, and one voter D.

Their preferences are defined below, where $y(s_j^r)$ denotes the candidate y_i associated with the r^{th} element of subset S_j, and when a subset of candidates is mentioned, the candidates are ranked according to the increasing order of their indices.

$$
\begin{array}{rll}
Y_i: & w \succ y_i \succ z \succ \{y_{i'}\}_{i' \neq i} \succ \{c_j\}_j \succ t & \text{for } i \in [3q] \\
C_j^r: & y(s_j^r) \succ c_j \succ z \succ w \succ \{y_{i'}\}_{i' \neq i} \succ \{c_j\}_j \succ t & \text{for } j \in [3q], r \in [3] \\
T^\ell: & t \succ z \succ w \succ \{y_i\}_i \succ \{c_j\}_j & \text{for } \ell \in [2] \\
Z^\ell: & z \succ w \succ \{y_i\}_i \succ \{c_j\}_j \succ t & \text{for } \ell \in [2] \\
W^\ell: & w \succ z \succ \{y_i\}_i \succ \{c_j\}_j \succ t & \text{for } \ell \in [2] \\
D: & w \succ t \succ z \succ \{y_i\}_i \succ \{c_j\}_j &
\end{array}
$$

Finally, the tie-breaking rule is as follows: $w \rhd t \rhd z \rhd y_1 \rhd \cdots \rhd y_{3q} \rhd c_1 \rhd \cdots \rhd c_{3q}$.

Table 5.1: Candidates' scores in the complexity proof of Theorem 61

candidate	initial score	announced score	score after manipulation
y_i ($i \in [3q]$)	3	3	3
c_j ($j \in [3q]$)	0	3 if $S_j \in S'$ 0 otherwise	3 if $S_j \in S'$ 0 otherwise
w	$3q+3$	2	2
t	2	2	3
z	2	3	2
winner	w	z	t

The winner of the election with the truthful ballot profile is candidate w. The details of the scores for this truthful ballot profile are given in the second column of Table 5.1.

We claim that there exists an exact cover in (X, S) iff we can force the election of candidate t in the constructed instance.

$\Longrightarrow$ **:** Suppose first that there exists a subset $S' \subseteq S$ such that every element of X occurs in exactly one subset of S'. Since $|X| = 3q$ and all elements of S are subsets of X of size 3, we have $|S'| = q$. Let us consider manipulated communicated scores which differ from the sincere ones by taking $3q+1$ votes initially given to w to give one additional vote to z and three votes to c_j for each $S_j \in S'$. These scores are summarized in the third column of Table 5.1. By the tie-breaking rule, candidate z is the announced winner.

It follows from these communicated scores that all candidates are potential winners except the candidates c_j such that $S_j \notin S'$. Therefore, each voter Y_i will deviate from ballot w to ballot y_i, for $i \in [3q]$, all voters C_j^r such that $S_j \in S'$ will deviate from ballot $y(s_j^r)$ to ballot c_j, and voter D will deviate from ballot w to ballot t. Since S' is an exact cover, each additional vote for y_i by voter Y_i will be balanced by the removal of one vote for y_i by the voter C_j^r, such that $S_j \in S'$ and $y(s_j^r) = y_i$, who deviates from y_i to c_j. Therefore, in total, these deviations will remove $3q+1$ votes from w, give three votes to q candidates c_j and add one vote to t, leading to the victory of t, as summarized in the fourth column of Table 5.1.

$\Longleftarrow$ **:** Suppose now that there exist communicated scores such that the target candidate t becomes the winner after deviations from the voters. The global idea of the proof is that the only possibility for communicated scores to lead to the victory of the target candidate t is to announce candidate z the winner and, as potential winners, the target candidate t and exactly q candidates c_j which correspond to subsets S_j forming an exact cover of X.

We will first show by disjunction case that the announced winner can only be candidate z.

Observe that t cannot win if it does not gain any additional vote. Indeed, for t to win with at most two votes, w cannot get more than one vote, and all the other candidates more than two votes, which sums to at most $12q+3$ votes for other candidates, whereas there would be $12q+5$ voters who do not vote for t, a contradiction. It follows that t cannot be announced as the winner, and must be a potential winner. However, by construction of the preferences, the only voter who can deviate to a ballot t is voter D. Therefore, in the deviating profile, t can get at most three votes.

If w is announced the winner, then the $3q$ voters Y_i will keep their vote for w, therefore t can never win with its maximum score of three, a contradiction.

Let us now analyze the case where the announced winner is a candidate y_i or c_j, by considering the candidates that can be announced potential winners:

- If candidate z is a potential winner, then at least voters T^ℓ and W^ℓ will deviate to it, which leads to at least four votes for z, whereas t can get at most three votes. Therefore, z cannot be a potential winner.
- Now, if candidate w is a potential winner, then at least voters T^ℓ and Z^ℓ will deviate to it, leading to at least four votes for w, whereas t can get at most three votes. Therefore, w cannot be a potential winner.
- Now, if a candidate $y_{i'}$ is a potential winner, for $i' < i$ or when c_j is the winner, then at least voters T^ℓ, Z^ℓ, and W^ℓ, for $\ell \in [2]$, will deviate to the candidate $y_{i'}$, that we call y^*, which is declared potential winner with the smallest index i', by construction of their preferences. Therefore, y^* would get at least six votes, whereas t can get at most three votes. Thus, such $y_{i'}$ cannot be a potential winner.
- Now, if a candidate $y_{i'}$ or $c_{j'}$ is a potential winner, for $i' > i$ and y_i winner, then voter Y_i will keep her vote for w as well as voters W^ℓ for $\ell \in [2]$, which leads to at least three votes for w whereas w is preferred to t in the tie-breaking rule. Therefore, such $y_{i'}$ or $c_{j'}$ cannot be a potential winner.
- Now, if a candidate $c_{j'}$ is a potential winner, for $j' < j$ and c_j winner, then at least voters T^ℓ, Z^ℓ, and W^ℓ, for $\ell \in [2]$, will deviate to the candidate $c_{j'}$, that we call c^*, which is declared potential winner with the smallest index j', by construction of their preferences. Therefore, c^* would get at least six votes, whereas t can get at most three votes. Thus, such $c_{j'}$ cannot be a potential winner.
- Now, finally, if a candidate $c_{j'}$ is a potential winner, for $j' > j$ and c_j winner, then all voters Y_i will keep their vote for w, which leads to at least $3q$ votes for w, whereas t can get at most three votes. Therefore, such $c_{j'}$ cannot be a potential winner.
- It follows that t is the only potential winner, and thus all voters Y_i keep

their vote for w, which leads to at least $3q$ votes for w, and thus t cannot win, a contradiction.

Consequently, the announced winner must be candidate z. Since t can get at most three votes, and w initially gets $3q+3$ votes, at least $3q+1$ votes must be removed from w (w is preferred to t in the tie-breaking). Voters W^ℓ will not deviate from w if z is the announced winner, therefore all voters Y_i and D must deviate from w. It follows that each candidate y_i must be a potential winner as well as t. However, for each candidate y_i, we need that at least one of the three voters C_j^r such that $y(s_j^r) = y_i$ deviates from her initial vote to y_i, otherwise y_i would get four votes and t could not win. For such a voter C_j^r to deviate, the only solution is to make candidate c_j a potential winner. By construction, it follows that we need to find a subset of candidates c_j (to make them potential winners) such that the associated subsets S_j entirely cover the elements in X. Thus, we need to make at least q candidates c_j potential winners.

Let us now analyze the compatible scores that can be communicated. If z is the announced winner with at most two votes, then by the tie-breaking rule, candidates w and t can get at most one vote, and all the other candidates at most two votes, which sums to at most $12q+2$ votes for other candidates, whereas there would be $12q+5$ voters who do not vote for z, a contradiction. If z is the announced winner with at least four votes, then to be potential winners, t should get at least three votes, all candidates y_i at least four votes, and at least q candidates c_j at least four votes, which sums to at least $16q+3$ votes for other candidates, whereas there are $12q+7$ voters in total, a contradiction. Consequently, z must be announced the winner with exactly three votes, and thus t must be announced with two votes, all candidates y_i with three votes, and at least q candidates c_j with three votes. The only possibility to announce such scores is to take $3q+1$ votes from w and to distribute them to give three votes to q candidates c_j and one vote to candidate z. The only possible margin then is given by the two remaining votes for w, however they are not sufficient to make another candidate c_j a potential winner. Hence, there are exactly q candidates c_j which are potential winners such that the associated subsets of S_j entirely cover X, which means that the union of such subsets is an exact cover. □

Note that even though we have proved that the problem is NP-complete, we know from Baumeister et al. [2020] that it is FPT when parameterized by the number of candidates m. Another way to go beyond the NP-hardness result, which focuses on worst-case complexity, is to analyze the actual possibility of poll manipulation using a probabilistic approach which works even when m is large. We will see that the poll manipulation problem is often easy to tackle in a probabilistic point of view, following natural statistical cultures as defined in Section 2.6. We will start by considering a balanced culture.

For a given target candidate x^* the polling institute wants to elect, we say its poll manipulation is successful if after all strategic moves from voters, the desired candidate x^* is elected. Let us denote by S the associated event of success, which corresponds to the yes-instances of the unrestricted poll manipulation problem.

Let 2PW-H(x^*, ℓ) be the heuristic detailed in Algorithm 2 which announces a score with exactly two potential winners x^* and ℓ, with x^* the target candidate and ℓ the announced winner. For realistic conditions, one point is given to candidates with a positive score in the truthful ballot profile. Assuming $n > m+5$ is sufficient to guarantee the possibility of making any pair of candidates the only potential winners (this hypothesis is rather weak since we focus on large elections in terms of voters). It then suffices to check whether the associated communicated polling score leads to the victory of x^*. This heuristic can be called by a global heuristic detailed in Algorithm 1, which tests it with different candidates ℓ.

Algorithm 1: Global Heuristic

Input: $(N, M, \mathscr{P}, \rhd)$, Target candidate x^*
foreach $\ell \in M \setminus \{x^*\}$ **do**
  $(is\_successful, s) \leftarrow$ 2PW-H(x^*, ℓ);
  if $is\_successful$ **then return** $(True, s)$;
return $(False, None)$

Algorithm 2: 2PW-H(x^*, ℓ)

Input: $(N, M, \mathscr{P}, \rhd)$, Target candidate x^*, Candidate ℓ
$s \leftarrow m$-vector with zeros; $R \leftarrow n$;
foreach $j \in M \setminus \{x^*, \ell\}$ **do**
  if $\exists i \in N$ *such that* $top_{\succ_i} = j$ **then** $s_j \leftarrow 1$; $R \leftarrow R - 1$;
$s_{x^*} \leftarrow \lfloor \frac{R}{2} \rfloor$; $s_\ell \leftarrow \lfloor \frac{R}{2} \rfloor$; $j^* \leftarrow \arg\min_{j \in M \setminus \{x^*, \ell\}} s_j$;
if $x^* \rhd \ell$ *and* R *is even* **then** $s_{x^*} \leftarrow s_{x^*} - 1$; $s_{j^*} \leftarrow s_{j^*} + 1$;
if $x^* \rhd \ell$ *and* R *is odd* **then** $s_\ell \leftarrow s_\ell + 1$;
if $\ell \rhd x^*$ *and* R *is odd* **then** $s_{j^*} \leftarrow s_{j^*} + 1$;
if $\mathcal{W}_P(b^s) = x^*$ **then return** $(True, s)$;
else return $(False, None)$;

We present a slightly modified version of Example 29 to illustrate this heuristic:

Example 30. *Let us consider an election* $(N, M, \mathscr{P}, \rhd)$ *where* $N = \{1, \dots, 6\}$, $M = \{x_1, x_2, x_3\}$, *the tie-breaking* $\rhd$ *follows the lexicographic order and the preference profile* $\mathscr{P}$ *is as follows:*

1:	$x_1 \succ x_2 \succ x_3$	*4:*	$x_2 \succ x_1 \succ x_3$
2:	$x_1 \succ x_2 \succ x_3$	*5:*	$x_3 \succ x_2 \succ x_1$
3:	$x_2 \succ x_1 \succ x_3$	*6:*	$x_3 \succ x_2 \succ x_1$

The initial truthful scores are given by $s^0 = (2, 2, 2)$. *If the score communicated by the polling institute is the truthful one, then voters 5 and 6 will vote for* x_2 *and she will be the winner. Suppose that the polling institute communicates the following score vector* $s^M = (2, 1, 3)$. *Then, voters 3 and 4 will deviate to* x_1 *and she will be the winner. Hence, the polling institute can enforce the election of* x_1, *whereas* x_2 *would be the winner without poll manipulation. In this example,* $x^* = x_1$ *and* $l = x_3$, *thus we play the heuristic 2PW-H(*x_1, x_3*).*

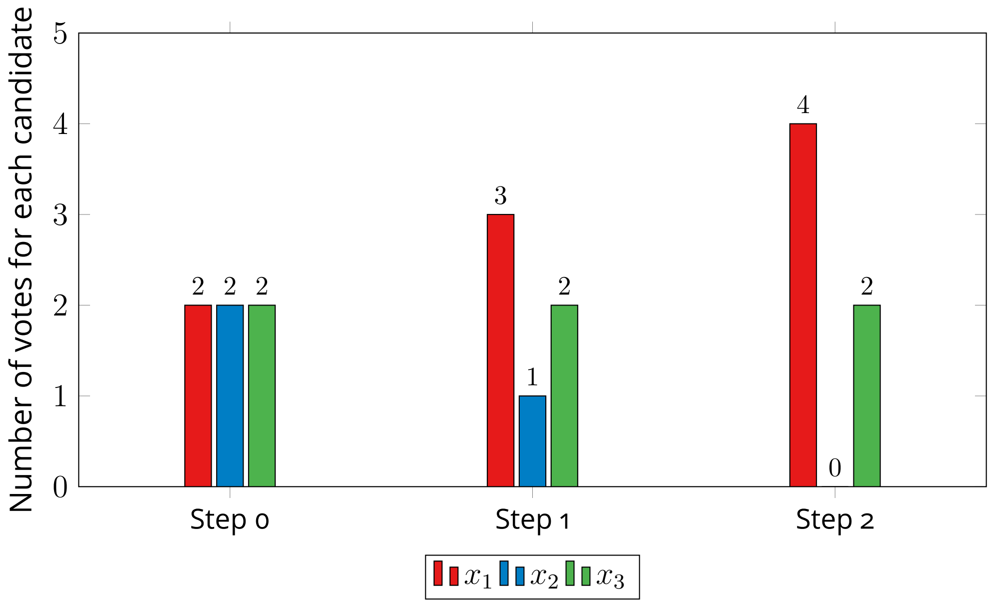

Figure 5.2: Example of strategic moves for the manipulation score $s^M = (2, 1, 3)$.

Our heuristics are computable in polynomial time and are inspired from the heuristics of Wilczynski [2019] and Baumeister et al. [2020], where the idea is to find a candidate ℓ, which is a threatening winner, i.e., enough voters prefer x^* to ℓ, while x^* is the only credible alternative to ℓ, in order to incentivize voters to deviate to x^*.

Let $S_{\text{2PW-H}(x^*,\,\ell)}$ denote the event of success for heuristic 2PW-H(x^*, ℓ). Let $X_{-\ell}$ be the random variable which counts the number of voters who prefer x^* over ℓ, i.e., $X_{-\ell} = |A^{x^* \succ \ell}|$.

Similarly, let $Y_{\ell,j}$ be the random variable which counts the number of voters who prefer ℓ over x^* while their most preferred candidate is j, i.e., $Y_{\ell,j} = |\{i \in A^{\ell \succ x^*} : top_{\succ_i} = j\}|$. If our heuristic 2PW-H(x^*, ℓ) indeed succeeds to announce exactly two potential winners x^* and ℓ with ℓ as a winner, then only voters who prefer x^* over ℓ and currently vote for another candidate, will deviate and they will do so in favor of x^*. Note that voters already having x^* as their top choice would keep this vote because there is no other potential winner. Therefore, in total, after deviations, x^* obtains a number of votes which is equal to the numbers of voters who prefer x^* over ℓ. It follows that x^* would win only if the number of voters preferring x^* over ℓ is greater than the number of voters who keep their vote for another candidate, implying that for a given culture $C(n, \Pi^m)$, $\mathbb{P}_C(S_{\text{2PW-H}(x^*,\,\ell)}) = \mathbb{P}_C(\forall j \in M \setminus \{x^*\}, Y_{\ell,j} \leqslant X_{-\ell})$.

Our first theorem provides a high lower bound on the probability of success of the poll manipulation heuristic.

Theorem 62. *For a balanced culture $C(n, \Pi^m)$, there exists $\ell \in B_C(x^*)$ such that the probability of success of the sub-heuristic 2PW-H(x^*, ℓ), is as follows: $\mathbb{P}_C(S_{\text{2PW-H}(x^*,\,\ell)}) \geqslant 1 - 2(m-2)(e^{-2n(p_{x^*,\ell} - q_{x^*,\ell,j^*})^2})$ where:*

- $p_{x^*,\ell} := \mathbb{P}_C(x^* \succ_i \ell)$,
- $r_{x^*,\ell,j} := \mathbb{P}_C(\{\ell \succ_i x^*\} \cap \{j = top_{\succ_i}\})$, *for* $j \neq \ell$,
- $q_{x^*,\ell,j} := \frac{p_{x^*,\ell} + r_{x^*,\ell,j}}{2}$,
- $j^* := \arg\min_{j \in M \setminus \{x^*,\ell\}} \mathbb{P}_C(Y_{\ell,j} \leqslant X_{-\ell})$.

In particular, the probability of success of the global heuristic satisfies the same lower bound.

Proof. For our target candidate x^* and a balanced culture $C(n, \Pi^m)$, let us consider a candidate $\ell \in B_C(x^*)$. Since $\ell \in B_C(x^*)$, ℓ will never be better and we can simplify the equality: $\mathbb{P}_C(S_{\text{2PW-H}(x^*,\,\ell)}) = \mathbb{P}_C(\forall j \in M \setminus \{x^*, \ell\}, Y_{\ell,j} \leqslant X_{-\ell})$. Our goal is to show a lower bound to $\mathbb{P}_C(\forall j \in M \setminus \{x^*, \ell\}, Y_{\ell,j} \leqslant X_{-\ell})$. For this purpose, the following lemma will be useful. We deduce from Bonferroni's inequality (Lemma 6) that:

$$\mathbb{P}_C(\forall j \in M \setminus \{x^*, \ell\}, Y_{\ell,j} \leqslant X_{-\ell}) \geqslant \sum_{j \in M \setminus \{x^*, l\}} \mathbb{P}_C(Y_{\ell,j} \leqslant X_{-\ell}) - (m-2-1)$$

$$\geqslant \sum_{j \in M \setminus \{x^*, \ell\}} \min_{j \in M \setminus \{x^*, \ell\}} \mathbb{P}_C(Y_{\ell,j} \leqslant X_{-\ell}) - (m-3)$$

$$\geqslant (m-2) \cdot \min_{j \in M \setminus \{x^*, \ell\}} \mathbb{P}_C(Y_{\ell,j} \leqslant X_{-\ell}) - (m-3)$$

By considering $j^* := \arg\min_{j \in M \setminus \{x^*,\ell\}} \mathbb{P}_C(Y_{\ell,j} \leqslant X_{-\ell})$, we then have that $\mathbb{P}_C(\forall j \in M \setminus \{x^*, \ell\}, Y_{\ell,j} \leqslant X_{-\ell}) \geqslant (m-2) \cdot (\mathbb{P}_C(Y_{\ell,j^*} \leqslant X_{-\ell}) - 1) + 1$.

Let us now treat the term $\mathbb{P}_C(Y_{\ell,j^*} \leqslant X_{-\ell})$. We remark that $X_{-\ell}$ follows a binomial distribution of parameters n and $p_{x^*,\ell}$ and Y_{ℓ,j^*} follows a binomial distribution of parameters n and r_{x^*,ℓ,j^*}, where $p_{x^*,\ell}$ and $r_{x^*,\ell,j}$ are defined as $p_{x^*,\ell} = \mathbb{P}_C(x^* \succ_i \ell)$ and $r_{x^*,\ell,j} = \mathbb{P}_C(\{\ell \succ_i x^*\} \cap \{j = top_{\succ_i}\})$, for every $j \neq \ell$. We introduce $q_{x^*,\ell,j^*} := \frac{p_{x^*,\ell}+r_{x^*,\ell,j^*}}{2}$ to lower bound our probability as follows:

$$\mathbb{P}_C(Y_{\ell,j^*} \leqslant X_{-\ell}) \geqslant \mathbb{P}_C(\{Y_{\ell,j^*} < q_{x^*,\ell,j^*} \cdot n\} \cap \{X_{-\ell} > q_{x^*,\ell,j^*} \cdot n\})$$

We use again Bonferroni's inequality (Lemma 6) to get: $\mathbb{P}_C(Y_{\ell,j^*} \leqslant X_{-\ell}) \geqslant \mathbb{P}_C(\{Y_{\ell,j^*} < q_{x^*,\ell,j^*} \cdot n\}) + \mathbb{P}_C(\{X_{-\ell} > q_{x^*,\ell,j^*} \cdot n\}) - 1$.

Applying the inequality from Lemma 4 on Bernoulli variables X_k with $a_k = 0$ and $b_k = 1$, for every $k \in [n]$, and $t = x \cdot \sqrt{n}$, we get: $\mathbb{P}(S_n - \mathbb{E}(S_n) \geqslant x \cdot \sqrt{n}) \leqslant e^{-2x^2}$, where $S_n = \sum_{i=1}^{n} X_k$. Now, by taking $x = \sqrt{n} \cdot (q_{x^*,\ell,j^*} - r_{x^*,\ell,j^*})$ and applying Lemma 4 to our sum of Bernoulli variables Y_{ℓ,j^*} (i.e., a binomial of parameters n and r_{x^*,ℓ,j^*}), we get:

$$\begin{aligned}
&\mathbb{P}_C(Y_{\ell,j^*} < q_{x^*,\ell,j^*} \cdot n) = 1 - \mathbb{P}_C(Y_{\ell,j^*} \geqslant q_{x^*,\ell,j^*} \cdot n) \\
&= 1 - \mathbb{P}_C(Y_{\ell,j^*} - r_{x^*,\ell,j^*} \cdot n \geqslant q_{x^*,\ell,j^*} \cdot n - r_{x^*,\ell,j^*} \cdot n) \\
&= 1 - \mathbb{P}_C(Y_{\ell,j^*} - r_{x^*,\ell,j^*} \cdot n \geqslant \sqrt{n}(\sqrt{n}(q_{x^*,\ell,j^*} - r_{x^*,\ell,j^*}))) \\
&\geqslant 1 - e^{-2n(q_{x^*,\ell,j^*} - r_{x^*,\ell,j^*})^2}
\end{aligned}$$

Now, we want to apply a similar treatment to variables $X_{-\ell}$. Let us denote $X'_{-\ell} = n - X_{-\ell}$ the random variable following a binomial distribution of parameters n and $1 - p_{x^*,\ell}$. We have:

$$\begin{aligned}
&\mathbb{P}_C(X_{-\ell} > q_{x^*,\ell,j^*} \cdot n) = \mathbb{P}_C(n - X'_{-\ell} > q_{x^*,\ell,j^*} \cdot n) \\
&= \mathbb{P}_C(X'_{-\ell} < n - q_{x^*,\ell,j^*} \cdot n) = 1 - \mathbb{P}_C(X'_{-\ell} - (1 - p_{x^*,\ell}) \cdot n \\
&\geqslant (1 - q_{x^*,\ell,j^*}) \cdot n - (1 - p_{x^*,\ell}) \cdot n) = 1 - \mathbb{P}_C(X'_{-\ell} - (1 - p_{x^*,\ell}) \cdot n \\
&\geqslant \sqrt{n}\sqrt{n}(p_{x^*,\ell} - q_{x^*,\ell,j^*})) \geqslant 1 - e^{-2n(p_{x^*,\ell} - q_{x^*,\ell,j^*})^2}
\end{aligned}$$

Putting the last two inequalities together we get:

$$\mathbb{P}_C(Y_{\ell,j^*} \leqslant X_{-\ell}) \geqslant 1 - e^{-2n(p_{x^*,\ell} - q_{x^*,\ell,j^*})^2} - e^{-2n(q_{x^*,\ell,j^*} - r_{x^*,\ell,j^*})^2}$$

Coming back to the first work of the proof we have:

$$\mathbb{P}_C(S_{\text{2PW-H}(x^*,\,\ell)}) \geqslant 1 - (m-2)(e^{-2n(p_{x^*,\ell} - q_{x^*,\ell,j^*})^2} + e^{-2n(q_{x^*,\ell,j^*} - r_{x^*,\ell,j^*})^2})$$

Finally, since q_{x^*,ℓ,j^*} is defined as the middle between $p_{x^*,\ell}$ and r_{x^*,ℓ,j^*}, we can simplify the inequality:

$$\mathbb{P}_C(S_{\text{2PW-H}(x^*,\,\ell)}) \geqslant 1 - 2(m-2)e^{-2n(p_{x^*,\ell} - q_{x^*,\ell,j^*})^2}$$

□

We can thus deduce the same lower bound for the probability of existence of a successful unrestricted poll manipulation.

Corollary 63. *For a balanced culture $C(n, \Pi^m)$, the probability of success of an unrestricted poll manipulation is as follows:* $\mathbb{P}_C(S) \geqslant 1 - 2(m - 2)(e^{-2n(p_{x^*,\ell} - q_{x^*,\ell,j^*})^2})$.

Our next theorem considers the asymptotic case and shows the convergence of the lower bound probability toward 1 when n becomes large. Since the number of voters is typically large in political elections, this shows an important susceptibility to poll manipulation.

Theorem 64. *For a balanced culture $C(n, \Pi^m)$, there exists $\ell \in B_C(x^*)$ such that the probability of success of the sub-heuristic 2PW-H(x^*, ℓ), and thus of the Global Heuristic, tends toward 1, i.e.,* $\lim_{n\to\infty} \mathbb{P}_C(S_{\text{2PW-H}(x^*,\ \ell)}) = 1$ *and thus* $\lim_{n\to\infty} \mathbb{P}_C(S) = 1$.

Proof. We use the lower bound from Theorem 62 to deduce the convergence toward 1 of this probability. In fact it is enough to pass to the limit on both sides in n the number of voters. The only tricky point might be when $p_{x^*,\ell} = q_{x^*,\ell,j^*}$. However, this situation can happen only when the culture puts positive probability only on preference orders whose top can only be x^* or j and in an equal manner, which is not possible by natural assumption on the culture. □

Observe that the quantities $p_{x^*,\ell}$ and $q_{x^*,\ell,j}$ from Theorem 62 are constants and different, we thus have exponentially fast convergence toward 1 for the probability of success of 2PW-H(x^*, ℓ) w.r.t. the number of voters.

Example 31. *To give a quick intuition, observe that for $m = 5$ and $n = 50$ under impartial culture, we get a lower bound of 0.89 and for $m = 5$ and $n = 100$, we already have a lower bound of 0.99 which is very fast! Example 31 is an illustration with different values of m:*

Beyond this general result on balanced cultures, the goal would be to capture realistic cultures regarding real elections [Boehmer et al., 2024]. From Propositions 58–60, we can derive the following corollary which shows that our general result covers very natural concrete cultures.

Lower bounds of $\mathbb{P}_{IC}(S_{\text{2PW-H}(x^*,\,\ell)})$ from Theorem 62

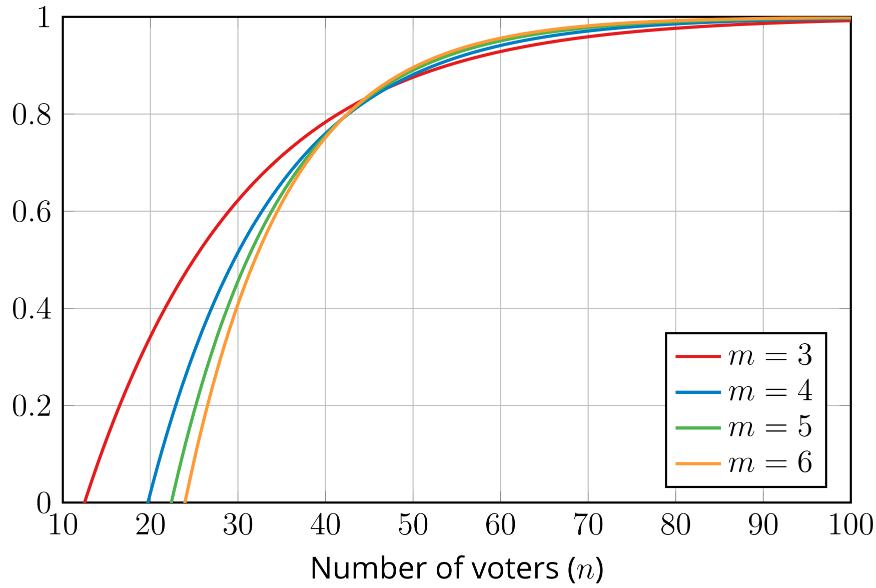

Figure 5.3: Plot of Lower Bounds with respect to n and $m = \{3, 4, 5, 6\}$

Corollary 65. *For a culture $C(n, \Pi^m)$, there exists $\ell \in B_C(x^*)$ such that the probability of success of 2PW-H(x^*, ℓ), and thus of the Global Heuristic, tends toward 1, i.e., $\lim_{n\to\infty} \mathbb{P}_C(S_{\text{2PW-H}(x^*,\,\ell)}) = 1$ and thus $\lim_{n\to\infty} \mathbb{P}_C(S) = 1$, when:*

- *C corresponds to the impartial culture, or*
- *C is a single-peaked culture and x^* is not an extreme candidate or x^* is extreme but $\mathbb{P}_C(\succ_i \mid\ worst_{\succ_i} = x^*) \leqslant \frac{1}{2}$, which includes Walsh's and Conitzer's cultures, or*
- *C corresponds to a Mallows culture $\mathcal{M}^{\phi,\sigma}$ where $x^* \neq worst_\sigma$.*

Our results show that even if the poll manipulation problem is hard, it is very likely for the polling institute to efficiently and successfully control the election, under natural preference distributions. However, the hypothesis that allows to send any score is questionable since the polling institute might be forced to meet some legal quality standards or to maintain voter trust by sending a reasonable score.

5.5 . The Restricted Poll Manipulation Problem

This section is devoted to the study of the manipulation problem in its restricted version i.e., the polling institute is restricted in its ability to lie about the scores and can only send a score vector from I_k.

The restricted poll manipulation problem is known to be NP-hard [Baumeister et al., 2020]. However, one could hope to get a fixed-parameter tractable algorithm w.r.t. the maximum allowed distance k to the truthful scores. We show below that such an efficient algorithm is unlikely to exist since we prove that the problem is W[1]-hard.

Theorem 66. *The restricted manipulation problem is W[1]-hard.*

Proof. From an instance $(G = (V, E), k)$ of k-Clique where $n := |V|$, $m := |E|$ and, w.l.o.g., $2 < k < n - 1$, we construct an instance of our restricted poll manipulation problem as follows.

For each vertex $v_i \in V$, we create a candidate v_i, and for each edge $\{v_i, v_j\} \in E$, we create a candidate e_{ij} (we suppose $i < j$ for this notation). We add three other candidates w, t, and z. In total, we thus have $n + m + 3$ candidates.

Let $K := (n-k)k$. For each vertex $v_i \in V$, we create k voters U_i^ℓ for $\ell \in [k]$, and $K - 1 - \delta(v_i)$ voters D_i^ℓ for $\ell \in [K - 1 - \delta(v_i)]$ (by our assumption on k, this quantity cannot be negative), where $\delta(v_i)$ denotes the degree of vertex v_i in G.

For each edge $\{v_i, v_j\} \in E$, we create two voters F_{ij}^i and F_{ij}^j, and $K - 2$ voters E_{ij}^ℓ for $\ell \in [K-2]$. Finally, we add K voters T^ℓ for $\ell \in [K]$ and $K - 1$ voters Z^ℓ for $\ell \in [K-1]$.

The preferences of the voters over the candidates are described below, for each $i \in [n]$, and each $\{v_p, v_q\} \in E$:

U_i^ℓ:	$w \succ v_i \succ z \succ \{v_j\}_{j \neq i} \succ \{e_{r,s}\}_{\{r,s\}} \succ t$	for $\ell \in [k]$
F_{pq}^p:	$v_p \succ e_{pq} \succ z \succ w \succ \{v_j\}_{j \neq p} \succ \{e_{r,s}\}_{\{r,s\} \neq \{p,q\}} \succ t$	
F_{pq}^q:	$v_q \succ e_{pq} \succ z \succ w \succ \{v_j\}_{j \neq q} \succ \{e_{r,s}\}_{\{r,s\} \neq \{p,q\}} \succ t$	
D_i^ℓ:	$v_i \succ z \succ w \succ \{v_j\}_{j \neq i} \succ \{e_{r,s}\}_{\{r,s\}} \succ t$	for $\ell \in [K - 1 - \delta(v_i)]$
T^ℓ:	$t \succ z \succ w \succ \{v_j\}_j \succ \{e_{r,s}\}_{\{r,s\}}$	for $\ell \in [K]$
E_{pq}^ℓ:	$e_{pq} \succ z \succ w \succ \{v_j\}_j \succ \{e_{r,s}\}_{\{r,s\} \neq \{p,q\}} \succ t$	for $\ell \in [K-2]$
Z^ℓ:	$z \succ w \succ \{v_j\}_j \succ \{e_{r,s}\}_{\{r,s\}} \succ t$	for $\ell \in [K-1]$

Finally, the tie-breaking rule is as follows: $z \rhd t \rhd \cdots \rhd w$.

The winner of the election with the truthful ballot profile is candidate w. The details of the scores for this truthful ballot profile are given in the second column of Table 5.2.

We claim that G admits a clique of size k iff we can force the election of candidate t by announcing scores which differ from the truthful ones by at most $k^2 + 1$ vote changes.

$\Longrightarrow$ **:** Suppose first that there exists a subset of vertices $S \subseteq V$ such that S is a k-Clique of G, i.e., $|S| = k$ and $\{v_i, v_j\} \in E$ for every $v_i, v_j \in S$. Let us consider manipulated communicated scores which differ from the sincere ones by taking k^2 votes initially given to w in order to give one additional vote

Table 5.2: Candidates' scores in the complexity proof of Theorem 66

candidate	initial score	announced score	score after manipulation
v_i $(i \in [n])$	$K-1$	K if $v_i \in S$ $K-1$ otherwise	K if $v_i \in S$ $K-1$ otherwise
e_{ij} $(\{v_i, v_j\} \in E)$	$K-2$	K if $v_p, v_q \in S$ $K-2$ otherwise	K if $v_p, v_q \in S$ $K-2$ otherwise
w	kn	K	K
t	K	$K-1$	K
z	$K-1$	K	$K-1$
winner	w	z	t

to each $v_i \in S$ (there are k such candidates) and two additional votes to each e_{ij} such that $v_i, v_j \in S$ (there are $\frac{k(k-1)}{2}$ such candidates), and finally by taking one vote initially given to t in order to give it to z. In total, the communicated scores differ from the sincere ones by exactly k^2+1 vote changes.

In the manipulated scores, z is winning with K votes thanks to the tie-breaking, while only the k candidates v_i corresponding to the vertices of the clique are announced as potential winners with K votes, as well as the $k(k-1)/2$ candidates corresponding to the edges of the clique, and candidate w. These manipulated scores are summarized in the third column of Table 5.2.

It follows from these communicated scores that all voters U_i^ℓ such that $v_i \in S$ deviate from w to v_i. By these deviations, candidate w loses k^2 votes, and thus obtains in total K votes, while each candidate $v_i \in S$ gains k votes. However, by definition of the clique, for each $v_i \in S$, there are exactly $k-1$ voters F_{ij}^i (or F_{ji}^i) who will deviate from v_i to the potential winner e_{ij} (or e_{ji}) corresponding to an edge incident to v_i. Therefore, each $v_i \in S$ also loses $k-1$ votes, and thus obtains in total K votes. Note that, by these deviations, each candidate e_{ij} such that $v_i, v_j \in S$ gains two additional votes and thus obtains in total K votes. No other deviation is possible because all remaining voters prefer z to all potential winners that are not at top of their preferences. The scores after all deviations are summarized in the fourth column of Table 5.2. The maximum score is K, which is obtained by w, k candidates v_i, and $\frac{k(k-1)}{2}$ candidates e_{ij}, and t. Candidate t is favored by the tie-breaking among these candidates and thus wins the election.

$\Longleftarrow$: Suppose now that there exist communicated scores such that the target candidate t becomes the winner after deviations from the voters. The global idea of the proof is that the only possibility for communicated scores to lead to the victory of the target candidate t is to announce candidate z the winner and, as potential winners, k candidates v_i, as well as $k(k-1)/2$ candidates e_{pq}, such that for each potential winner v_i, there are $k-1$ potential winners e_{ij} (or e_{ji}) corresponding to edges incident to v_i.

We will first prove that z must be announced as the winner. Observe that no voter can deviate to t because every voter, except all voters T^ℓ who already vote for t, ranks it last. It follows that we need that at least k^2 voters U_i^ℓ, who currently vote for w, deviate to another candidate, and thus w cannot be announced as the winner.

Let us analyze the case where the announced winner would be a candidate v_i, e_{pq} or candidate t, by considering the candidates that can be announced potential winners:

- If candidate z or w is a potential winner, then at least all voters D_i^ℓ and all voters E_{rs}^ℓ (except voters E_{pq}^ℓ if e_{pq} is announced as the winner) would deviate to z if z is a potential winner or to w otherwise, and thus z or w would gain too many votes compared to t and t would never win. Therefore, none of them is a potential winner.
- Now, if a candidate $v_{i'}$ is a potential winner, for $i' < i$ or when e_{pq} or t is the winner, then all voters D_i^ℓ and all voters E_{rs}^ℓ (except voters E_{pq}^ℓ if e_{pq} is announced as the winner) would deviate to such candidate $v_{i'}$, that we call v^*, which is declared potential winner with the smallest index i'. Thus, such v^* would gain too many votes compared to t and t would never win. Therefore, such $v_{i'}$ cannot be a potential winner.
- Now, if a candidate $v_{i'}$ or e_{rs} is a potential winner, for $i' > i$ and v_i winner, then all voters $U_{i''}^\ell$, for $i'' \neq i'$, would keep their vote for w and thus w would have too many votes compared to t and t would never win. Therefore, such $v_{i'}$ or e_{rs} cannot be a potential winner.
- Now, if a candidate e_{rs} is a potential winner, for $\{r, s\} < \{p, q\}$ when e_{pq} winner or for t winner, then at least all voters D_i^ℓ and all voters E_{rs}^ℓ (except voters E_{pq}^ℓ if e_{pq} is announced as the winner) would deviate to such candidate e_{rs}, that we call e^*, which is declared potential winner with the smallest index $\{r, s\}$. Therefore, e^* would get too many votes compared to t and t would never win. Thus, such e_{rs} cannot be a potential winner.
- Now, finally, if a candidate e_{rs} is a potential winner, for $\{r, s\} > \{p, q\}$ and e_{pq} winner, then all voters $U_{i'}^\ell$ would keep their vote for w and thus w would have too many votes compared to t and t would never win. Therefore, such e_{rs} cannot be a potential winner.
- It follows that t is the only potential winner, and thus all voters $U_{i'}^\ell$ keep their vote for w. Thus, w has too many votes compared to t and t cannot win, a contradiction.

Hence the communicated scores must announce z as the winner.

Since z is ranked among the first two most preferred candidates by all voters D_i^ℓ, T^ℓ, E_{pq}^ℓ and Z^ℓ, none of these voters will deviate. Recall that we need at least k^2 voters U_i^ℓ (for $i \in [n]$ and $\ell \in [k]$) who deviate to another candidate, and the only candidate other than their top candidate that voters U_i^ℓ prefer to z is v_i, for all $\ell \in [k]$. Therefore, we need to announce at least k candidates v_i as potential winners. In such a way, each chosen candidate v_i gains k additional votes, whereas it initially had $K-1$ votes from voters D_i^ℓ, who cannot deviate, and from voters F_{ij}^i (or F_{ji}^i) for each edge $\{v_i, v_j\} \in E$. Since t will have at most K votes, we need at least $k-1$ voters F_{ij}^i (or F_{ji}^i) who deviate from ballot v_i. The only other candidate that such voters prefer to z is candidate e_{ij} (or e_{ji}). Therefore, for each chosen v_i potential winner, we also need to announce as potential winners at least $k-1$ candidates e_{ij} (or e_{ji}) which correspond to edges incident to v_i.

Recall that we can only announce scores which differ from the truthful ones by at most k^2+1 vote changes. If we announce z the winner with at most $K-1$ votes, then we need to remove at least k^2+1 votes for w and one vote for t, therefore we have already exceeded our budget. If we announce z the winner with at least $K+1$ votes, then we need to add two votes to at least k candidates v_i, three votes to at least $\frac{k(k-1)}{2}$ candidates e_{ij} and one vote to z, therefore we have already exceeded our budget. It follows that we need to announce z the winner with exactly K votes. In this case, we need to add one vote to z, one vote to at least k candidates v_i and two votes to at least $\frac{k(k-1)}{2}$ candidates e_{ij}. Therefore, to meet our budget, we need to declare exactly k candidates v_i and exactly $\frac{k(k-1)}{2}$ candidates e_{ij} as potential winners, in such a way that for potential winner v_i there exist $k-1$ potential winners e_{ij} corresponding to incident edges. Hence, the chosen candidates v_i correspond to a k-clique in G. $\square$

Nevertheless, we prove below that the restricted poll manipulation problem can be efficiently solved if the parameter k of the maximum distance to the truthful scores is a constant.

Proposition 67. *The restricted manipulation problem is in XP w.r.t. the maximum distance k to the truthful scores. More precisely, it can be solved by an algorithm which runs in time $\Theta(m^{2k+1} \cdot n)$.*

Proof. We give an upper bound to $|I_k|$. We denote that any move of voters is characterized by the origin and the destination candidate. Since our distance counts the number of swaps, one swap is defined by choosing two candidates, we then get $\binom{m}{2} = \frac{m(m-1)}{2}$ and $|I_1| \leqslant \frac{m(m-1)}{2}$. We start from s^T and iterate the upper bound argument and we get: $|I_k| \leqslant (\frac{m(m-1)}{2})^k \leqslant m^{2k}$. It is then enough to visit every score of I_k and add the winner determination in $\Theta(m \cdot n)$. At the end, we get $\Theta(m^{2k+1} \cdot n)$. $\square$

However, the previous result cannot be used if k is large and does not tell whether there actually exists a successful manipulation. We thus use a probabilistic approach to analyze the possibility of poll manipulation. Let S_k

denote the event of success for the restricted poll manipulation where k denotes the maximum allowed distance to the truthful scores. We first prove that when k is small compared to $\sqrt{n}$, the restricted poll manipulation tends to be impossible.

Theorem 68. *For any culture $C(n, \Pi^m)$, if the maximum distance k to the truthful scores is such that $k = o(\sqrt{n})$ and the target candidate x^* is not winning in the initial score, then the probability of existence of a successful poll manipulation to elect x^* tends toward zero, i.e., $\lim_{n\to\infty} \mathbb{P}_C(S_k) = 0$.*

Proof. Let us start by identifying the probability law of the truthful scores.

The truthful scores follow a multinomial law because there are n voters' preferences drawn independently at random with the same law, and we have m possibilities for the most preferred candidate of each voter, and these are the only necessary elements to compute scores s^T. We will use the following result on multinomial laws.

Lemma 69 ([Severini, 2005]). *If $(N_n)_{n\geqslant 0}$ is a multinomial law in $\mathbb{R}^m$ with parameters n and $q = (q_1, \ldots, q_m)$ and $\mathcal{N}(0; K)$ a multivariate normal distribution then $\frac{1}{\sqrt{n}}(N_n - nq) \underset{n\to+\infty}{\longrightarrow} \mathcal{N}(0; K)$, where $K_{i,j} = q_i\delta_{i,j} - q_iq_j$, for every $1 \leq i, j \leq m$, with $\delta_{i,j} = 1$ if $i = j$ and $\delta_{i,j} = 0$ otherwise.*

Let c^* be the truthful winner, i.e., $c^* := \mathcal{W}_P(b^T)$. Informally, a necessary condition for the existence of a successful manipulation with the two-candidate heuristic is that there is at least one candidate that is sufficiently close to the winner. The pair of candidates would then be this candidate and the current winner. Of course, this is not necessarily sufficient, as the pair may not be the right one. However, we will see that this necessary condition occurs with probability 0, and that's enough for us to conclude. We then write

$$S_k \subset \{ \bigcup_{z\neq c^*} \{|s^T_{c^*} - s^T_z + \mathbb{1}_{c^*\rhd z}| \leq k\}\}$$

We will analyze the probability of the second event to get an upper bound on the probability of success of the restricted poll manipulation problem. By using Observation 48 and Lemma 69 with $N_n = s^T$, we get: $\frac{1}{\sqrt{n}}(s^T - nq) \underset{n\to+\infty}{\longrightarrow} \mathcal{N}(0; K)$, where $K_{i,j} = q_i\delta_{i,j} - q_iq_j$, for every $1 \leq i, j \leq m$. We denote $\mathcal{N}(0; K) = (\mathcal{N}_1, \ldots, \mathcal{N}_m)$ and remark that each $\mathcal{N}_j$ follows a Gaussian law. For any $z \in M \setminus \{x^*\}$, we have

$$\lim_{n\to+\infty} \mathbb{P}_C(|s^T_{c^*} - s^T_z + \mathbb{1}_{c^*\rhd z}| \leq k)$$

$$= \lim_{n\to+\infty} \mathbb{P}_C(|\frac{1}{\sqrt{n}}s^T_{c^*} - nq_{c^*} - \frac{1}{\sqrt{n}}s_z + nq_z + \frac{\mathbb{1}_{c^*\rhd z}}{\sqrt{n}} + n(q_{c^*} - q_z)| \leq \frac{k}{\sqrt{n}})$$

Combining this equality with the previous convergence result using the test function

$$\Phi(s^T) = \mathbb{1}_{\{|\frac{1}{\sqrt{n}}s^T_{c^*} - nq_{c^*} - \frac{1}{\sqrt{n}}s^T_z + nq_z + \frac{\mathbb{1}_{c^* \rhd z}}{\sqrt{n}} + n(q_{c^*} - q_z)| \leq \frac{k}{\sqrt{n}}\}}$$

and

$$\lim_{n \to +\infty} \frac{k}{\sqrt{n}} = 0$$

by assumption, we deduce that

$$\lim_{n \to +\infty} \mathbb{P}_C(|\frac{1}{\sqrt{n}}s^T_{c^*} - nq_{c^*} - \frac{1}{\sqrt{n}}s^T_z + nq_z + \frac{\mathbb{1}_{c^* \rhd z}}{\sqrt{n}} + n(q_{c^*} - q_z)| \leq \frac{k}{\sqrt{n}}))$$

$$= \mathbb{P}_C(|\mathcal{N}_{c^*} - \mathcal{N}_z + \frac{\mathbb{1}_{c^* \rhd z}}{\sqrt{n}} + n(q_{c^*} - q_z)| \leq 0) = 0$$

Therefore,

$$\lim_{n \to +\infty} \mathbb{P}_C(|s^T_{c^*} - s^T_z + \mathbb{1}_{c^* \rhd z}| \leq k) = 0$$

It follows for the probability of the success event that

$$\lim_{n \to +\infty} \mathbb{P}_C(S_k) \leqslant \lim_{n \to +\infty} \mathbb{P}_C(\bigcup_{z \neq c^*} |s^T_{c^*} - s^T_z + \mathbb{1}_{c^* \rhd z}| \leq k)$$

$$\leq \lim_{n \to +\infty} \sum_{z \neq c^*} \mathbb{P}_C(|s^T_{c^*} - s^T_z + \mathbb{1}_{c^* \rhd z}| \leq k) = 0$$

We then get:

$$\lim_{n \to +\infty} \mathbb{P}_C(S_k) = 0$$

which concludes the proof. □

Then, we get immediately the following corollary if we include the case where x^* might win in the initial poll, because it is always possible to communicate scores that keep the same winner.

Corollary 70. *For any culture $C(n, \Pi^m)$, if the maximum distance k to the truthful scores is such that $k = o(\sqrt{n})$, then $\lim_{n \to \infty} \mathbb{P}_C(S_k) = \mathbb{P}_C(\{\mathcal{W}_P(s^T) = x^*\})$.*

We might note that, e.g., $\mathbb{P}_C(\{\mathcal{W}_P(s^T) = x^*\}) \approx \frac{1}{m}$ when considering the impartial culture.

We now focus on a case where poll manipulation can be successful, and prove that we can even efficiently compute it, thanks to an adaptation of the global heuristic where the sub-heuristic to call is Restricted 2PW-H(x^*, ℓ) which, starting from s^T, tries to announce ℓ as the winner and x^* as the only other potential winner, while taking into account the maximum allowed distance k.

Let $S_{\text{Restr-2PW-H}(x^*, \ell)}$ denote the event of success of this sub-heuristic.

Algorithm 3: Restricted 2PW-H(x^*, ℓ)

Input: $(N, M, \succ, \rhd, k)$, Target candidate x^*, Candidate ℓ

$s \leftarrow s^T$; $R \leftarrow 0$;
while $\exists c \in M \setminus \{\ell\}$ *s.t.* $s_c \geq s_\ell - \mathbb{1}_{c \neq x^* \rhd \ell} + \mathbb{1}_{\ell \rhd c = x^*}$ *and* $R < k$ **do**
  $y \leftarrow \arg\min_{x^*, \ell}\{s_{x^*}, s_\ell - 1\}$; $s_y \leftarrow s_y + 1$; $s_c \leftarrow s_c - 1$;
   $R \leftarrow R + 1$;
while $s_{x^*} < s_\ell - \mathbb{1}_{x^* \rhd \ell}$ *and* $R < k$ **do**
  $j^* \leftarrow \arg\max_{j \in M \setminus \{\ell\}} s_j$;
  if $s_\ell > \max_{j \in M \setminus \{\ell\}} s_j + 2$ **then** $j^* \leftarrow \ell$;
  $s_{x^*} \leftarrow s_{x^*} + 1$; $s_{j^*} \leftarrow s_{j^*} - 1$; $R \leftarrow R + 1$;
if $\mathcal{W}_P(b^s) = x^*$ **then return** $(True, s)$;
else return $(False, None)$;

Theorem 71. *For a balanced culture C, if the maximum distance k to the truthful scores is such that $n = o(k)$ where and $c^* := \mathcal{W}_P(b^T)$, then there exists $\ell \in B_C(x^*)$ such that the probability of success of Restricted 2PW-H(x^*, ℓ) tends toward 1, i.e., $\lim_{n\to\infty} \mathbb{P}_C(S_{\text{Restr-2PW-H}(x^*,\, \ell)}) = 1$ and thus $\lim_{n\to\infty} \mathbb{P}_C(S_k) = 1$.*

Proof. We can first observe that each score the polling institute may send can be summarized by its set of potential winners and its winner, since two announced scores with the same potential winners and winner produce the same voters' deviations. A type $\mathcal{T}(s)$ for a score vector s is thus defined as a pair $(PW, w) \in 2^M \times M$ where $w \in PW$, representing its potential winners and its winner. The set of all possible score types is denoted by $\mathcal{T}$. We will then show that:

$$\mathbb{P}_C(\{\bigcup_{s \in I_k} \mathcal{T}(s) = \mathcal{T}\}) = 1$$

Let c^* be the truthful winner, i.e., $c^* := \mathcal{W}_P(b^T)$. Informally, a sufficient condition for the existence of a strategy of each type is that all candidates are sufficiently close to the winner. More precisely, we would like them all to be closer than $\frac{k}{m}$, so that the cost of making potential winners any pair of candidates never exceeds k. We then get:

$$\{\bigcap_{z \neq c^*} \{|s^T_{c^*} - s^T_z + \mathbb{1}_{c^* \rhd z}| < \frac{k}{m}\}\} \subset \{\bigcup_{s \in I_k} \mathcal{T}(s) = \mathcal{T}\}$$

We again use the same technique adding and subtracting $n \cdot q_{c^*}$ and $n \cdot q_z$ and a central limit theorem on the truthful scores s^T following a multinomial law (Observation 48). However, we have this time a remaining term $\sqrt{n}(q_{c^*} - q_z)$ that is bounded by assumption ($n = o(k)$).

We then get:

$$\lim_{n \to +\infty} \mathbb{P}_C(|s^T_{c^*} - s^T_z + \mathbb{1}_{c^* \rhd z}| < \frac{k}{m}) = 1$$

Since a countable intersection of events of probability 1 is of probability 1, we have:

$$\lim_{n\to+\infty} \mathbb{P}_C(\bigcap_{z\neq c^*} \{|s^T_{c^*} - s^T_z + \mathbb{1}_{c^*\rhd z}| < \frac{k}{m}\}) = 1$$

Then, we have:

$$\mathbb{P}_C(\bigcup_{s\in I_k} \mathcal{T}(s) = \mathcal{T}) = 1$$

Using the fact that a successful strategy exists in the unrestricted case and since all strategies are accessible, we get:

$$\lim_{n\to\infty} \mathbb{P}_C(S_{\text{Restr-2PW-H}(x^*,\,\ell)}) = 1$$

Therefore, we have also:

$$\lim_{n\to\infty} \mathbb{P}_C(S_k) = 1$$

□

The case $k = \alpha \cdot n$ with $\alpha \in]0,1]$ works exactly as in Theorem 71 if $\alpha > \max_{c,c'\in M} |p_c - p_{c'}|$, where p_c denote the probability that candidate c is elected. This has a clear interpretation: if the polling institute is allowed to lie by a fraction α on scores then we will fall in the manipulation regime for a sufficiently large number of voters.

Like for the unrestricted problem, the general result of Theorem 71 holds for the concrete cultures under the condition mentioned in Section 5.3.

Corollary 72. *For a culture C and $n = o(k)$ and $c^* := \mathcal{W}_P(b^T)$, there exists $\ell \in B_C(x^*)$ such that the probability of success of Restricted 2PW-H(x^*, ℓ) tends toward 1, i.e., $\lim_{n\to\infty} \mathbb{P}_C(S_{\text{Restr-2PW-H}(x^*,\,\ell)}) = 1$ and thus $\lim_{n\to\infty} \mathbb{P}_C(S_k) = 1$, when:*

- *C corresponds to the impartial culture, or*
- *C is a single-peaked culture and x^* is not an extreme candidate or x^* is extreme but $\mathbb{P}_C(\succ_i|\ worst_{\succ_i} = x^*) \leqslant \frac{1}{2}$, which includes Walsh's and Conitzer's cultures, or*
- *C corresponds to a Mallows culture $\mathcal{M}^{\phi,\sigma}$ where $x^* \neq worst_\sigma$.*

5.6 . Toward a Generalization of Strategic Behavior

Until now, we only considered strategic moves from pivotal voters. However, one can argue that voters might want to deviate when they are close enough to be pivotal. Such a strategic behavior can be captured by considering pivotal thresholds $p_i \in \mathbb{N}$ for each voter i, as done by Wilczynski [2019] in an idea close to local-dominance [Meir et al., 2014]. This slightly modifies the definition of potential winners:

Definition 20 (General potential winners)**.** *A candidate y is a* general potential winner *for voter i w.r.t. score s if i believes that adding p_i votes to y will make candidate y the new winner, i.e., $s^{-i}_{\mathcal{W}_P(s^{-i})} - s^{-i}_y + \mathbb{1}_{\mathcal{W}_P(s^{-i}) \rhd y} \leqslant p_i$. We denote PW_i^{s,p_i} the set of general potential winners for i w.r.t. score s.*

The definition of best response naturally follows by considering general potential winners. Our initial setting corresponds to the case where $p_i = 1$.

Let us first analyze the impact of pivotal thresholds on strategic voting. For this purpose, we suppose that the polling institute is sincere and sends truthful scores $s = s^T$, and that all thresholds are equal and denoted by p, i.e., $p_i = p$, for every voter i. Let us define the expected proportion of strategic voters $\mathcal{P}_{SV}$ w.r.t. culture $C(n, \Pi^m)$, n, m, and p. Let U_i^p denote the event where the top candidate of voter i is not a general potential winner for i, i.e., $U_i^p = \{top_{\succ_i} \notin PW_i^{s,p}\}$, and D_i^p the event where voter i could favor a potential winner other than the current winner, that she prefers to it, i.e., $D_i^p = \{\exists w \in M \setminus \{top_{\succ_i}\} : w \succ_i \mathcal{W}_P(s)$ and $w \in PW_i^{s,p}\}$. By definition, the proportion of strategic voters counts the voters for who the two events are true, i.e., $\mathcal{P}_{SV}(C, n, m, p) = \mathbb{E}_C[\frac{1}{n}\sum_{i=1}^n \mathbb{1}_{U_i^p \cap D_i^p}]$.

The following proposition provides several insights on the proportion of strategic voters at the limits, by showing that the variations of the dependent events U_i and D_i are opposed with respect to p. The point 5 is in spirit quite close to the work of Xia [2012] since it is related to the margin of victory. Indeed, the lower bound in Theorem 14 from [Xia, 2012] would be enough to conclude.

Proposition 73.
1. *$U = \mathbb{E}_C[\frac{1}{n}\sum_{i=1}^n \mathbb{1}_{U_i^p}]$ is decreasing w.r.t. p.*
2. *$D = \mathbb{E}_C[\frac{1}{n}\sum_{i=1}^n \mathbb{1}_{D_i^p}]$ is increasing w.r.t. p.*
3. *$\mathcal{P}_{SV}(C, n, m, p) \leqslant \min(U, D)$.*
4. *$\lim_{p \to +\infty} \mathcal{P}_{SV}(C, n, m, p) = 0$ and $\mathcal{P}_{SV}(C, n, m, 0) = 0$.*
5. *$\lim_{n \to +\infty} \mathcal{P}_{SV}(C, n, m, p) = 0$ when p is fixed.*

Proof.

1-2. The statements follow from the inclusions $U_i^{p'} \subseteq U_i^p$ and $D_i^p \subseteq D_i^{p'}$, for each $p' > p$.

3. Using the inclusions $U \cap D \subset U$ and $U \cap D \subset D$, we show that: $\mathcal{P}_{SV}(C, n, m, p) \leqslant U$ and $\mathcal{P}_{SV}(C, n, m, p) \leqslant D$.
4. If p is maximum, then all candidates are potential winners and thus each voter keeps her truthful vote, while when $p = 0$ there are no potential winners to deviate to.
5. Using Lemma 69, we know that the winner c^* and any other candidate z will be spread out at least of order $\sqrt{n}$ asymptotically. We then deduce

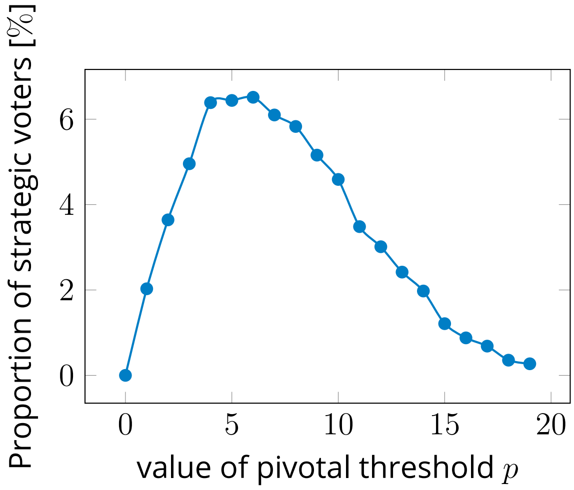

Figure 5.4: Proportion of strategic voters depending on the pivotal threshold p in an election with 100 voters and 4 candidates under impartial culture.

that there are no potential winners other than the winner in that case, since p is fixed. □

Although the previous proposition helps to better understand the proportion of strategic voters at the limits, it is still difficult to exactly determine the behavior for other values of p, in particular when $\mathcal{P}_{SV}$ is maximum, because of the dependency between U and D.

These results gives us some ideas of how this proportion evolves in some cases but we are still not able to derive some computations and give a precise view of the graph. To complement Theorem 73, we run simulations to find this maximal proportion w.r.t. p. where we compute the average proportion of strategic voters with impartial culture, $n = 100$, $m = 4$, for each $p \in [n]$ and 3,000 runs per p.

Experimental results in Figure 5.4 show a peak for manipulation around $p = 5$, which can be interpreted as voters who believe that a candidate which is at at distance less than 5% to win is a plausible winner, which seems reasonable.

Now that we have a better understanding of the impact of the pivotal threshold, let us now analyze the poll manipulation problems. Let S^G (resp., S_k^G) denote the associated event of success for the unrestricted (resp., restricted) problem with generalized strategic behavior.

Proposition 74. *For a balanced culture $C(n, \Pi^m)$, $p > 0$ and $p = o(n)$, we have $\mathbb{P}_C(S^G) \geqslant \mathbb{P}_C(S)$ and $\mathbb{P}_C(S_k^G) \geqslant \mathbb{P}_C(S_k)$.*

Proof. In an idea similar to the proof of Theorem 71, for each score of a given type in the initial setting, we can always choose a score of the same type which works for the generalized strategic behavior, since they would trigger the same deviations. Indeed, even if the polling institute is not sending the same score, it might construct a score with the same potential winners and the same winner since p is negligible against n. □

It follows from Theorem 74 that Theorems 62 and 64, for successful unrestricted poll manipulation, also hold under a generalized strategic behavior with the given weak hypotheses. Similarly, the convergence result toward 1 for the probability of a successful restricted poll manipulation (Theorem 71) and the generalization of Theorem 68 can also be extended to a generalized strategic behavior.

5.7 . Conclusion and Future Works

5.7.1 . Conclusion

In the context of political elections where voters are assumed to be strategic, we have studied the poll manipulation problem: *Can a polling institute lie about candidates' scores it communicates to voters in order to influence the outcome of the election?* Two variants are investigated: an unrestricted one where any scores can be sent, and a restricted one, more realistic, where the polling institute cannot announce scores too far from the reality. We show that both problems are computationally hard and answer an open question from Baumeister et al. [2020]. However, we go beyond this worst-case analysis by using probabilistic tools to balance computational hardness.

First, from a computational point of view, we solved the open problem of the complexity of the unrestricted version and added two new parameterized complexity results for the restricted version which is XP and W[1]-hard parameterized by the distance between the truthful poll and the one sent by the polling institute.

Second, we show relationships between cultures and prove that, under a very weak hypothesis, most of the cultures studied in the literature are included in a general condition.

Third, we use this preliminary work to study the probability of manipulation with an easily computable heuristic under this broad condition on culture, giving a balancing argument for NP-hardness. Under a broad condition on cultures, satisfied by many concrete preference distributions, we prove a lower bound on the probability of success of an easily computable heuristic for the unrestricted problem. This enables us to obtain a rapid convergence toward 1 of the manipulation probability, meaning that large elections are highly manipulable when the polling institute can freely manipulate without altering the trust of voters.

When it may not be the case, i.e., in a restricted context, our asymptotic results show that manipulability strongly relies on whether the allowed distance to truthful scores depends on the election size. Manipulation tends to fail when this distance is negligible w.r.t. the number of voters. However, when the distance is significant, e.g., is a given proportion of the election size, which appears as a very natural assumption, efficient and successful manipulation tends to be always possible, showing that political elections are highly susceptible to poll manipulation in practice.

5.7.2 . Future Works

This work provides different aspects of poll manipulation under plurality voting. Nevertheless, it opens the door to several promising directions for future research:

- One could consider other voting rules or other types of information communicated in the poll would be natural.
- Another avenue of work could be to examine different strategic voting behaviors, by, e.g., examining that abound in the literature, such as local dominance [Meir, 2018].
- Finally, a challenging future direction would be to adapt our analysis to dependent cultures such as the Pólya-Eggenberger urn.

6 - Modelling the Variability of Strategic Voting Outcomes under Uncertainty: A General Approach

Abstract

We present a new model of strategic voting where uncertainty is modeled by probability sets and the decisions are taken according to lower and upper expected utility gains. Focusing on the particular case of belief functions, we show that this generic model encompasses in one sweep previous models that considered either sets or probabilities to model uncertainty, and enables us to generalize some well-known convergence results from the literature. We also discuss the case of uncertain voters, and emphasize the challenges that come with considering richer and therefore more realistic models of uncertainty.

Résumé

Nous proposons un nouveau modèle de vote stratégique dans lequel l'incertitude est représentée par des ensembles de probabilités, et les décisions sont prises sur la base des gains espérés inférieurs et supérieurs. En nous concentrant sur le cas particulier des fonctions de croyance, nous montrons que ce modèle générique englobe, de manière unifiée, les approches antérieures qui utilisaient soit des ensembles, soit des probabilités pour modéliser l'incertitude. Il nous permet également de généraliser certains résultats classiques de convergence de la littérature. Nous abordons aussi le cas des électeurs incertains, et soulignons les défis soulevés par l'adoption de modèles d'incertitude plus riches et donc plus réalistes.

Most of the content of this chapter is based on a paper co-authored with Sébastien Destercke, which was accepted at the 16^{th} Multidisciplinary Workshop on Advances in Preference Handling (M-PREF 2025) [Destercke and Surugue, 2025]

6.1 . Introduction

In this final chapter, we end our study of voting outcomes by incorporating uncertainty into strategic voting, thereby generalizing the existing literature on the topic and allows for a more expressive model. As already emphasized in Chapter 2 and Chapter 4, the strategic voting model is motivated by the need to reflect the unavoidable incentives that voters face to manipulate outcomes. However, as already mentioned, all models considered so far assume complete preferences and perfect information on the voters' side, except for the model from [Meir, 2017] which remains purely qualitative. Indeed, this raises several practical challenges. First of all, building polls is creating some uncertainty by construction since you have to query a subset of voters and make some inferences about the remaining voters. This is the assumption of the two main models in the literature, where authors have considered both a probability around received polls [Myerson and Weber, 1993] for the first model, and the use of an uncertain neighborhood [Meir, 2017] in the second model. However, the first model makes all strategic moves in one-shot, without voters being able to change their votes afterwards, which is hard to justify. The second model is iterative but with a more qualitative behavior, meaning that it considers any score in the neighborhood with the same importance. In this chapter, our goal will be to generalize these models by introducing some new concepts that both include the existing literature and give a framework to capture new insights. Specifically, we will be interested in giving a generalization of Meir's models (see Chapter 2 and Chapter 4) by including some quantitative aspect, namely introducing the idea that some scores are less important than others, if they are far from the broadcast poll for example. In this context, we will give convergence results of our iterative voting process.

However, our model also allows us to conceptualize new ideas as an uncertainty about voter's themselves, similar to the work of Kreiss and Augustin but for iterative voting. For example, a voter may not know her full preference ordering or may be unable to compare certain candidates. Our model capture situations where voters have either uncertain ballots or incomplete preferences as in [Conitzer et al., 2011; Dey et al., 2018], which is often the case in practice. We will show how to model such behaviors and explain why it remains complicated to implement this in practice.

Thus, in both situations, the type of information and its reliability is key to understand how information impacts voter opinions is crucial. Previous work in this direction has considered specific types of information, such as identifying the winner only under certain voting rules (different from plurality since it is complete information in that case), or using the majority graph as a basis (see [Endriss et al., 2016; Reijngoud and Endriss, 2012]).

The rest of the chapter is organized as follows: Section 6.2 describes in detail the required background of uncertainty theory for our new model based on the one already introduced in Chapter 2, as it is a central contribution of the chapter. Section 6.3 then motivates it, first by showing how it includes and generalizes previous frameworks, second by providing practical examples where one would need to consider generalized representations of uncertainty in order to have an accurate model. Section 6.4 provides convergence results for our newly introduced model, showing that while being more general, it preserves some key properties discussed in the literature. Finally, Section 6.5 provides a discussion rather than a conclusion, as we think this work opens up many avenues of research that would bring strategic voting closer to real-life behavior of voters, hopefully offering a higher descriptive power for such models.

6.2 . A Strategic Voting Model with Uncertainty

We start by introducing the elements of the newly proposed model. As the model itself is one of the main contribution of this chapter, we will describe it in details, before illustrating some of its aspects.

6.2.1 . A voting Situation Equipped with an Uncertainty Model

We will follow all the same notations as in previous chapters Chapter 4 and Chapter 5 to described the strategic voting framework under plurality.

Before introducing our uncertainty model, we make the following remark to situate it within a broader theoretical framework.

Remark 75. *Probabilistic theory assumes that any state of uncertainty can, in practice, be modeled by a probability distribution. However, in many real-world situations, the available probabilistic information is only partial. Data may be imprecise, unreliable, or too scarce to justify the computation of frequencies. In such cases, the uniform distribution is often used as a default model of ignorance, though this choice is highly debatable.*

We equip our election with an uncertainty model on scores. We consider that the voter receives a convex probability set $\mathcal{P}$ over the set of possible scores $I_n^m = \{s : \sum_{x \in M} s_x(b) = n\}$. We omit n and m when the number of voters and candidates is not specified or not relevant to the current discussion. Two common ways to build $\mathcal{P}$ are (1) to consider an uncertainty model around an original score s and (2) by having uncertainty about each ballot, assumed independent, thus building a global $\mathcal{P}$ over M^n through each individual $\mathcal{P}_i$ (the uncertainty of voter i). Note that since a ballot profile maps to a single score vector, a set defined over M^n can be unambiguously mapped to a set $\mathcal{P}$ over I_n^m.

Possible scores are sufficient if we work under plurality rule, otherwise we would have to consider probability sets on the spaces of vote profiles.

From a probability set $\mathcal{P}$ on I_n^m, one can then define, for any function $f : I_n^m \to \mathbb{R}$ (typically a utility of some kind), a lower expectation

$$\underline{\mathbb{E}}_{\mathcal{P}}(f) = \inf_{P \in \mathcal{P}} \mathbb{E}_P(f) \tag{6.1}$$

where $\mathbb{E}_P$ is the standard, linear expectation operator. Upper expectation $\overline{\mathbb{E}}_{\mathcal{P}}(f)$ can be defined like-wise, taking a $\sup$. Of particular interest are lower and upper expectations over identity functions of events A, that corresponds to lower and upper probabilities over such events, and that we will denote $\underline{P}(A), \overline{P}(A)$. We will also use the following notations for sets of scores: an interval for a given candidate simply means that the score of this candidate can vary within these bounds. For instance, if $n = 5$ and $m = 3$, the set-valued score $E = ([2, 3], [1, 3], [0, 2])$ means that the score of b will be between 1 (one voter will vote for b for sure) and 3.

Probability sets $\mathcal{P}$ are very generic uncertainty models that include as special cases [Destercke and Dubois, 2014] quite a number of existing uncertainty representations. We detail some of those that will be of particular interest in this chapter.

The specific case of sets A first special case is the one of sets, where we know that $s \in E$ for some subset $E \subseteq I$. The corresponding probability set is defined as

$$\mathcal{P}_E = \{P : P(E) = 1\},$$

summarised by the constraint $\underline{P}(E) = 1$. In such a case, Equation (6.1) has the solution $\underline{\mathbb{E}}_{\mathcal{P}_E}(f) = \inf_{x \in E} f(x)$.

The specific case of probability Probability sets obviously generalize probabilities by construction. Therefore, if we take $\underline{P}(A) = \overline{P}(A)$ for all events A, we come back to the very well-known probability setting, in which $\mathcal{P}$ reduces to a singleton.

The specific case of belief functions A belief function over M^n (or I_n^m) consists in defining a positive mass function $\mathcal{M} : 2^{M^n} \to [0, 1]$ that sums up to one, i.e., $\sum_{E \subseteq M^n, E \neq \emptyset} \mu(E) = 1$. From such a mass $\mathcal{M}$ can then be defined two bounds over events $A \subseteq M^n$ that are defined as

$$\underline{P}(A) = \sum_{E \subseteq A} \mu(E) \tag{6.2}$$

$$\overline{P}(A) = \sum_{E \cap A \neq \emptyset} \mu(E) \tag{6.3}$$

and from which one can define a corresponding set of possible probabilities

$$\mathcal{P}_{\mathcal{M}} = \{P : \forall A, \underline{P}(A) \leq P(A) \leq \overline{P}(A)\} \tag{6.4}$$

Should we collect a mass function $\mathcal{M}_i$ per voter defined over M, $\mu_i(E)$ would then express the mass of evidence we have that the voter ballot will be in E. From those, one can easily build a joint mass function over M^n (hence over I_n^m): if E_i are some subsets of M receiving positive masses for voter i, then the joint mass given over $E_1 \times \ldots \times E_n$ is simply $\mathcal{M}(E_1 \times \ldots \times E_n) = \prod_{i=1}^{N} \mathcal{M}_i(E_i)$, which is well-suited to denote independence between agents [Smets and Kennes, 1992]. Each imprecise ballot can then be mapped to a corresponding set of possible score vectors. Belief functions, while strictly less expressive than generic lower probabilities $\underline{P}(A)$, are still a quite expressive model, as they include probabilities and sets as special cases: probabilities correspond to masses bearing only on singletons, while a set E correspond to the mass $\mu(E) = 1$.

An advantage of belief functions is that Equation (6.1) as well as the upper expectation can be easily solved by adopting the following formula using $\mathcal{M}$, i.e.,

$$\underline{\mathbb{E}}_{\mathcal{M}}(f) = \sum_{E \subseteq M^n} \mu(E) \inf_{x \in E} f(x),$$

$$\overline{\mathbb{E}}_{\mathcal{M}}(f) = \sum_{E \subseteq M^n} \mu(E) \sup_{x \in E} f(x).$$

Note that since M^n is finite the sup (resp. inf) is a max (resp. min).

Of interest within the literature devoted to belief functions is the notion of pignistic probability p^* of $\mathcal{M}$, that allows one to go from $\mathcal{P}_{\mathcal{M}}$ to a single probability within it. This pignistic probability corresponds to the Shapley value of the lower probability $\underline{P}(A)$ viewed as a game, and is given by:

$$p^*(s) = \sum_{E, s \in E} \frac{\mu(E)}{|E|}.$$

Therefore, we can also see that taking the pignistic probability also considering a Laplacian assumption within each set E receiving positive evidence. Note also that if the mass $\mathcal{M}$ is equivalent to a probability p, then $p^* = p$, meaning that when one starts with a probability as uncertainty model, the pignistic criterion is just the standard expected utility for this probability.

Finally and before stepping to the decision part of our model, we will mention two particular cases of belief functions that are of interest to us: the first is the case of necessity measures, where the masses are given to nested elements, that is $\mu(E_i) \neq 0$ and $\mu(E_j) \neq 0$ iff $E_i \subset E_j$ or $E_i \subset E_j$, and the case of inner measures [Denneberg, 2013, Ch. 2.], where masses are given on a partition of the space, that is $\mu(E_i) \neq 0$ and $\mu(E_j) \neq 0$ if $E_i \cap E_j = 0$ and $\cup_{E, \mu(E) > 0} E = I_n^m$ (or a subset of interest).

6.2.2 . Some Strategic Decisions

In this model, we want to encode the strategic behavior of voters under uncertainty. To do so, let us define a function $st(s, a_i, a_i') : I_n^m \times M^2 \to I_n^m$ that sends back an updated score if voter i moves his vote from a_i to a_i' in its votes. Then, let $u_i : M \to \mathbb{R}$ be for each voter i the associated utility that will often reflect how satisfied the voter is with the current winner $\mathcal{W}_P(s)$. This utility function could be more complex, e.g., be defined on M^2 to denote pairwise preferences over candidates and represent how much some candidates are preferred to other ones, however we will therefore stick to this definition depending only on M. Moreover, even though we have strict linear orders $\succ_i$ for each voter, we can always build an associated utility functions by Debreu et al. [1954]. Therefore, to understand how voters may change their vote, we need to evaluate the benefit of moving from a_i to a_i' that we denote

$$u_i(a_i'|a_i, s), \tag{6.5}$$

that is the utility of making the move from a_i to a_i', given the initial state of affairs s. In particular, if we are certain about our current state of affairs, if we know the score vector, it suffices for (6.5) to be positive for the voter to make a strategic move, that is

$$a_i' \succeq a_i \text{ iff } u_i(a_i'|a_i, s) \geq 0$$

However, in our model this is not the case since voters receive a probability set $\mathcal{P}$. Assuming that we have some probability set $\mathcal{P}$ defined over I, we will define four decision criteria (named DC) under uncertainty.

- The first criteria that we will consider is a pessimistic criteria to decide whether an action is better than another as follows:

  $$a_i' \succ_{pess} a_i \text{ iff } \underline{\mathbb{E}}_{\mathcal{P}}(u_i(a_i'|a_i, s)) \geq 0$$

  $$\text{and } \overline{\mathbb{E}}_{\mathcal{P}}(u_i(a_i'|a_i, s)) > 0$$

  where S denotes uncertain scores whose knowledge is modeled by $\mathcal{P}$. In essence, this corresponds to the maximmin criterion put forward by Gilboa and Schmeidler [1989].

- A pignistic criteria to decide whether an action is better than another as follows:

  $$a_i' \succeq_{pig} a_i \text{ iff } \mathbb{E}_{p^*}(u_i(a_i'|a_i, s)) \geq 0,$$

  where p^* is the pignistic probability. Note that it includes standard probabilistic decision as a special case.

- A mixture criteria to decide whether an action is better than another as follows:

$$a'_i \succeq_{\alpha,p^*} a_i \text{ iff}$$

$$\text{iff } \alpha \cdot \underline{\mathbb{E}}_{\mathcal{M}}(u_i(a'_i|a_i,s))$$

$$+(1-\alpha) \cdot \mathbb{E}_{p^*}(u_i(a'_i|a_i,s)) \geq 0$$

  with $\alpha \in [0,1]$. Note that α can be seen as a level of completeness of the obtained, as it allows one to go from a very partial order ($\alpha = 0$) to a complete one ($\alpha = 1$).

- Another mixture of criteria to decide whether an action is better than another as follows:

$$a'_i \succeq_{H(\alpha)} a_i \text{ iff}$$

$$\text{iff } \alpha \cdot \underline{\mathbb{E}}_{\mathcal{M}}(u_i(a'_i|a_i,s))$$

$$+(1-\alpha) \cdot \overline{\mathbb{E}}_{\mathcal{M}}(u_i(a'_i|a_i,s)) \geq 0$$

  with $\alpha \in [0,1]$. In essence, this corresponds to the well-known Hurwicz criterion [Hurwicz, 1951] that intends to balance between optimism ($\overline{\mathbb{E}}$) and pessimism ($\underline{\mathbb{E}}$), and that has been justified within a belief setting by Denoeux and Shenoy [2020].

Note that for the three last rules, the strict relation $\succ$ corresponds to having strict positive values on the right side. Those decision rules mirror some of the commonly used decision rules within the imprecise probabilistic setting [Troffaes, 2007], and we hope to leverage other commonly used rules in future works.

An equilibrium is a situation where no voter has an incentive to deviate from its ballot, i.e., $\forall i \in N, \nexists\ a'_i$ such that $a'_i \succ a_i$ (we assume voters will only make a strategic move if their preferences are strict). To sum up, we can fully describe an election under plurality with a lexicographic tie-breaking and strategic voting under uncertainties with the following tuple $(N, M, \mathscr{P}, \rhd, \mathcal{P}, (u_i)_{i \in N}, DC)$.

6.3 . Motivations of the Model

Our model is enriched because it includes the existing models from the literature but also provides a deeper understanding of the uncertainty in iterative voting in Plurality elections.

6.3.1 . Relations with Existing Models

One of the strengths of our model is its power to unify two theories that were first considered as different, namely the one from Meir [2017]; Meir et al. [2010] and the one from Myerson and Weber [1993].

For the first one, we need to define a neighborhood as Meir [2017] with respect to a distance d : $S_r(s) = \{s' \in I \mid d(s, s') \leqslant r\}$, d can be ℓ_1, ℓ_∞ or another one. For example, the ℓ_1 distance corresponds to the number of voters that need to be added or removed in order to convert one score vector into another. The ℓ_∞ distance, on the other hand, measures the largest difference in the number of votes received by any single candidate between the two score vectors. We could also use an Earth Mover's distance, which is a variant of the ℓ_1 distance that assumes a fixed total number of voters. By abuse of notation, we will use S_r when it is clear from the context. The seminal work of Meir et al. [2010] is working on the special case $r = 0$, meaning the poll information is complete for plurality and certain. This corresponds to considering $\forall i \in N, \mu_i(S_0) = 1$, particular utilities functions:

$$\forall i \in N, u_i(a'_i|a_i, s)$$

$$= \begin{cases} 1 & \text{if } st(s, a_i, a'_i) \succ_i s \text{ and } \mathcal{W}_P(st(s, a_i, a'_i)) = a'_i, \\ 0 & \text{if } st(s, a_i, a'_i) \succ_i s \text{ and } \mathcal{W}_P(st(s, a_i, a'_i)) \neq a'_i, \\ 0 & \text{if } st(s, a_i, a'_i) \sim_i s, \\ -1 & \text{else,} \end{cases}$$

and the pessimistic decision rule $\succeq_{pess}$ to recover the model. Recall that in Meir [2017], the convergence is guaranteed and we call these strategic moves “direct best response" since all strategic moves where a voter deviates is to the new winner.

Meir et al. [2014] then extended this framework to add uncertainty by considering some strictly positive value $r_i > 0$ for voter i. This comes down in our model to take, $\forall i \in N, \mu_i(S_{r_i}) = 1$, particular utilities functions

$$\forall i \in N, u_i(a'_i|a_i, s) = \begin{cases} 1 & \text{if } st(s, a_i, a'_i) \succ_i s, \\ 0 & \text{if } st(s, a_i, a'_i) \sim_i s, \\ -1 & \text{else.} \end{cases}$$

and again the pessimistic decision rule $\succeq_{pess}$ to recover the model.

Let us give a short example.

Example 32. *Consider an election with twelve voters and three candidates where one voter i has the following preference and votes initially truthful: $c \succ_i a \succ_i b$. The initial poll result is as follows: $(4, 6, 2)$ (see the picture).*

The voter would not change her vote with the initial model (i.e., $r = 0$). However, if we choose a level of uncertainty $r = 1$, using the EMD distance for instance,

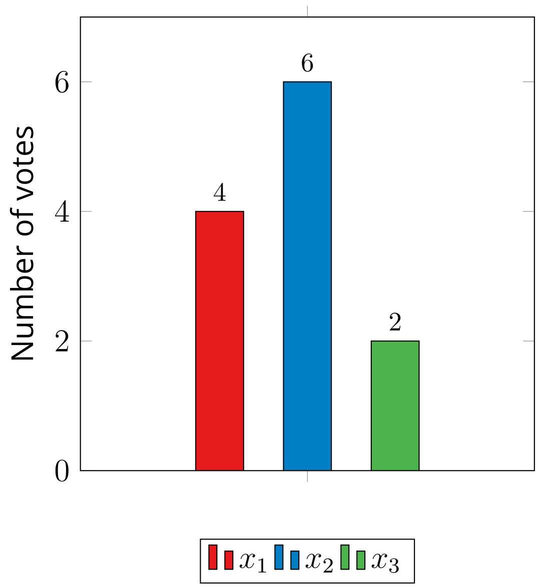

Figure 6.1: Vote distribution for the score profile $(4, 6, 2)$

a strategic move occurs, indicating that this model captures new behaviors. Indeed, the neighborhood of all possible score will be

$$S_1 = \{(4,6,2), (3,7,2), (3,6,3), (5,5,2), (4,5,3), (5,6,1), (4,7,1)\}$$

Therefore, in almost all score profiles, the voter cannot influence the outcome by changing her vote from c to a, except in the cases $(4, 5, 3)$ and $(5, 6, 1)$. Let us detail this point: with score $(4, 5, 3)$, voter i can change her vote from c to a in order to make a win instead of b, which is more preferable to her. A similar situation occurs with score $(5, 6, 1)$. The incorporation of uncertainty in the model thus justifies a strategic move from c to a.

For the probabilistic model of Myerson and Weber [1993], we define $\mathcal{P}$ as a probability $\mathbb{P}(s) \sim M(s, n)$ where s is the observed score and M a multinomial distribution. The utilities can be chosen arbitrarily provided and verify the voter preferences:

$$u_i(a_i'|a_i, s) \geq 0 \text{ iff } st(s, a_i, a_i') \succ_i s$$

Then, we need to remark that in Myerson and Weber [1993] all voters are acting simultaneously and the decision is that each voter maximizes its expected utilities. The computation can be greatly simplified by only looking at the pivotal state and because we look only at plurality:

$$\mathbb{E}_{\mathbb{P}}(u_i(a_i'|a_i, s))$$
$$= \sum_{a_i \in M} \mathbb{P}(a_i, a_i') u_i(a_i'|a_i, s)$$

where $\mathbb{P}(a_i, a'_i)$ is the probability to be pivotal by doing the action from a_i to a'_i.

We suggest two ways to build richer model of uncertainty about the poll, either from the polling institute or from the voters themselves. Building polls, even when done honestly, comes with uncertainty because you cannot ask everyone's preferences (see the case of dishonest polls [Mousseau et al., 2024]). We can also interpret uncertainty from voters as they may be uncertain or have incomplete knowledge of their preferences. The next section presents some examples to illustrate these ideas.

6.3.2 . Examples

Our model not only includes other models from the literature but is also able to capture other situations. First of all, one of the weakness of the first model from Meir [2017] is that any score of the neighborhood is considered identically. Ideally, we would like to affect weights relative to the distance to the broadcast score. Let us detail that idea in the two following examples.

Example 33. *Let $m = 3$, $n = 30$ and the broadcast score being $s = (10, 9, 11)$. We will define information belief available directly on I^3_{30}, $\mathcal{M}(S_1) = 0.5$, $\mathcal{M}(S_2) = 0.3$ and $\mathcal{M}(S_3) = 0.2$.*

This model describes exactly what we expect with a belief function that models a certainty decreasing with respect to the distance to the broadcast poll s. This corresponds to building possibility distributions, that among imprecise probabilistic models are the most direct extensions of sets. They are therefore the ideal candidate to extend Meir [2017] approach.

However, we may want to have a model with a more probabilistic flavour. In this case, one can consider the partition generated by the S_i sets, and associate a probability mass to each member of this partition, generating a so-called inner measure [Denneberg, 2013, Ch. 2.] on any event of the initial space I. An intrinsic interest of this model is that it remains a probabilistic one, even if defined on an algebra where the atoms are sets of scores.

Example 34. *Let $m = 3$, $n = 30$ and the broadcast score being $s = (10, 9, 11)$. We will define uncertain information directly on I^3_{30}, $\mathcal{M}(S_1) = 0.5$, $\mathcal{M}(S_2 \setminus S_1) = 0.3$ and $\mathcal{M}(S_3 \setminus S_2) = 0.2$.*

Moreover, we can also collect the lack of knowledge or incomplete information about voter's preferences. In fact, in most elections we usually have undecided voters especially far from the election day. Let us develop that idea in the two followings examples. In the following example, voter 1 is undecided between candidate a and candidate b, voter 2 and 3 are certain.

Example 35. *Let $m = 3$ and $n = 3$, $\mu_1(\{a, b\}) = 1$, $\mu_2(\{b\}) = 1$ and $\mu_3(\{c\}) = 1$. The joint belief function derived from this information over I^3_3 is $\mathcal{M}(\{(1, 1, 1), (0, 2, 1)\}) = \mathcal{M}(([0, 1], [1, 2], 1)\}) = 1$.*

Now, we move to a last example that shows our model is able to capture more complex uncertainty. Indeed, we might have a voter (i.e., voter 1) that is able to tell she is not voting for c but it still hesitating between a and b, with a slight preference for a, and other voters being certain. The next example proposes a belief functions accounting for that.

Example 36. *Let $m = 3$ and $n = 3$, we could have $\mu_1(\{a\}) = 0.5$, $\mu_1(\{a, b\}) = 0.5$, so that $\underline{P}_1(\{a\}) = 0.5$, $\overline{P}_1(\{a\}) = 1$, and $\overline{P}_1(\{b\}) = 0.5$, meaning that voting for a is definitely more probable, yet voting for b remains plausible, albeit less. Provided we have $\mu_2(\{b\}) = 1$ and $\mu_3(\{c\}) = 1$, the joint mass over scores would be $\mathcal{M}(\{(1, 1, 1)\}) = 0.5, \mathcal{M}(([0, 1], [1, 2], 1)\}) = 0.5$.*

Let us now illustrate one of our decision rules on this same example, to give the reader an idea of how this plays out.

Example 37. *Consider for instance that the preference of voter 2 are $b \succ_2 c \succ_2 a$, which is coherent with its ballot and a Hurwicz criterion with $\alpha = \frac{1}{3}$, meaning that the weight of optimism is larger. We will consider the following utilities:*

$$\forall i \in N, u_i(a'_i|a_i, s) = \begin{cases} 1 & \textit{if } st(s, a_i, a'_i) \succ_i s, \\ 0 & \textit{if } st(s, a_i, a'_i) \sim_i s, \\ -1 & \textit{else}. \end{cases}$$

We want to evaluate the deviation from b to c. Then, with the linearity of the mass functions, we compute

$$\underline{\mathbb{E}}_{\mathcal{M}}(u_i(a'_i|a_i, s)) = \frac{1}{2} \cdot 1 + \frac{1}{2} \cdot (-1) = 0$$

and

$$\overline{\mathbb{E}}_{\mathcal{M}}(u_i(a'_i|a_i, s)) = \frac{1}{2} \cdot 1 + \frac{1}{2} \cdot 1 = 1$$

Indeed, note that for the set $\{(1, 1, 1)\}$ having mass 0.5, the move makes c elected instead of a, and the utility is a constant 1. For the set $([0, 1], [1, 2], 1)$ that also has mass 0.5, the move is still beneficial for the possible score $(1, 1, 1)$, but would bring negative utility for $(0, 2, 1)$, as c would be elected instead of b, a clear downside for voter 2. Thus, we get:

$$\begin{aligned} &\frac{1}{3} \cdot \underline{\mathbb{E}}_{\mathcal{M}}(u_i(a'_i|a_i, s)) \\ &+(1 - \frac{1}{3}) \cdot \overline{\mathbb{E}}_{\mathcal{M}}(u_i(a'_i|a_i, s)) \\ &= \frac{1}{3} \cdot 0 + \frac{2}{3} \cdot 1 = \frac{2}{3} \geq 0. \end{aligned}$$

The conclusion is that voter 2 will do the deviation from b to c with such a decision criterion.

Note that we restrict our uncertainty models here to belief functions, yet some natural assessments would need to opt for the richer language of generic convex sets of probabilities, such as providing partial order between the probabilities of voting for some candidates [Miranda and Destercke, 2015]. Considering such cases is however beyond the scope of the current chapter.

6.4 . Convergence with Quantitative Uncertainty on Scores

In this section we generalize existing results on convergence from Meir [2017] by adding quantitative uncertainty about the score as in Example 33 and Example 34. Let us recall that a neighborhood S_r with respect to a distance d is defined as in Meir [2017]: $S_r(s) = \{s' \in I \mid d(s,s') \leqslant r\}$, where d can be ℓ_1, ℓ_∞, earth moving distance (EMD) or another one. Let us consider from now on the ℓ_1 distance that is easy to interpret in terms of transfer of votes. For each voter i, we will let r_i be the support of its uncertainty, so that S_{r_i} is the biggest set of voter i's uncertainty.

6.4.1 . Extending the Meir Framework

Let us consider any belief functions of voter i, $\mathcal{M}_i : 2^{S_{r_i}} \to [0,1]$ that sums up to one, i.e., $\sum_{E \subseteq S_{r_i}, E \neq \emptyset} \mu_i(E) = 1$.

We will consider two particular cases of interest. The first is the case of nested sets, i.e., $\forall i \in N, \forall k \geqslant 1, \mu_i(S_k) = \beta^i_k$, with $\sum_{k=1}^{r_i} \beta^i_k = 1$. Indeed, we might want to describe the fact that the belief mass is depending of the distance to the true score. Second, we look at the case of partitioned belief function, i.e., $\forall i \in N, \forall k \geqslant 2, \mu_i(S_k \setminus S_{k-1}) = \beta^i_k$, with $\sum_{k=1}^{r_i} \beta^i_k = 1$ and $\forall i \in N, \mu_i(S_1) = \beta^i_1$. In both cases, we will assume decreasing $(\beta_k)_{1 \leqslant k \leqslant r_i}$, meaning that our evidence decreases as we get further from the observed s.

For the rest of this section, we will consider the following utilities that make sense in terms of improved or deteriorated outputs:

$$\forall i \in N, u_i(a'_i|a_i,s) = \begin{cases} 1 & \text{if } st(s,a_i,a'_i) \succ_i s, \\ 0 & \text{if } st(s,a_i,a'_i) \sim_i s, \\ -1 & \text{else.} \end{cases}$$

Theorem 76. *Voters considering uncertainty given by a nested or partitioned belief function, and making strategic decisions according to either pessimistic ($\succeq_{pess}$) or mixed ($\succeq_{\alpha,p^*}, \succeq_{h(\alpha)}$) decision rules with α large enough, will converge to an equilibrium.*

Proof. At first, we will take $\forall i \in N, \forall k \geqslant 1, \mu_i(S_k) = \beta^i_k$, with $\sum_{k=1}^{r_i} \beta^i_k = 1$. We consider a feasible move $a_i \to a_{i'}$. We remark that the hypothesis that $\underline{\mathbb{E}}_{\mathcal{M}}(u_i(a'_i|a_i,s)) \geqslant 0$ implies that either:

- $\forall s \in S_{r_i}, u_i(a'_i|a_i,s) \neq -1$,

- $\exists \tilde{r} \in [0, r_i], \forall s \in S_{\tilde{r}}, u_i(a_i'|a_i, s) = 1$

since, if some neighborhoods have a negative utility, this must be compensated by a positive contribution, which happens only if all moves within a neighborhood the results. However, the second case is impossible, for the reason that if we allow a voter to transfer its vote to another candidate ($r = 1$), then there is a situation for which the voter is not pivotal, and therefore $\exists s \in S_{\tilde{1}}, u_i(a_i'|a_i, s) = 0$, showing that the second case never happens. Therefore, if

$$\underline{\mathbb{E}}_{\mathcal{M}}(u_i(a_i'|a_i, s)) \geqslant 0$$

this means that $\forall s \in S_{r_i}, u_i(a_i'|a_i, s) \in \{0, 1\}$, with at least one s giving the null value.

Second, note that we cannot have $\overline{\mathbb{E}}_{\mathcal{M}}(u_i(a_i'|a_i, s)) > 0$ if $\forall s \in S_{r_i}, u_i(a_i'|a_i, s) = 0$, and there must be a situation for which making this strategic move is a local dominance move, meaning it verifies Theorem 4 from Meir [2017].

We do exactly the same reasoning for the second type of belief functions, i.e., $\forall i \in N, \forall k \geqslant 2, \mu_i(S_k \setminus S_{k-1}) = \beta_k^i$, with $\sum_{k=1}^{r_i} \beta_k^i = 1$ and $\forall i \in N, \mu_i(S_1) = \beta_1^i$.

For the mixed decision, one can prove that the decision behaves as the pessimistic one if α is large enough. □

We now want to go a step further by showing the convergence can hold even with a less, but still pessimistic behavior, namely the Hurwicz criterion with $\alpha > \frac{1}{2}$. This allows us to consider negative outcomes in the uncertainty neighborhood. Let us consider some belief functions as follows: $\forall i \in N, \forall k \geqslant 1, \mu_i(S_k) = \beta_k^i$, with $\sum_{k=1}^{r_i} \beta_k^i = 1$

Theorem 77. *Voters considering a belief function around the true score and making strategic votes according to a sufficiently pessimistic Hurwicz criterion (i.e., $\alpha > \frac{1}{2}$) will converge to an equilibrium.*

Proof. We assume that the weights β_k^i of our belief functions, defined as $\forall i \in N, \forall k \geqslant 1, \mu_i(S_k) = \beta_k^i$.

We consider a feasible move $a_i \to a_{i'}$. Using the fact that the lower (an upper) expectation is positive homogeneous, i.e., $\alpha\underline{\mathbb{E}}(f) = \underline{\mathbb{E}}(\alpha f)$, we get

$$\begin{aligned}
&\alpha \cdot \underline{\mathbb{E}}_{\mathcal{M}}(u_i(a_i'|a_i, s)) \\
&+ (1 - \alpha) \cdot \overline{\mathbb{E}}_{\mathcal{M}}((u_i(a_i'|a_i, s)) \\
&= \sum_{i=1}^{r_i} \mu(S_i)[\alpha \inf_{s \in S}(u_i(a_i'|a_i, s)) \\
&+ (1 - \alpha) \cdot \sup_{s \in S}(u_i(a_i'|a_i, s))]
\end{aligned} \tag{6.6}$$

Let us denote by

$$\tilde{u_i}(S, a_i, a_i') = \alpha \inf_{s \in S}(u_i(a_i'|a_i, s))$$

$$+(1-\alpha) \cdot \sup_{s \in S}(u_i(a_i'|a_i, s))$$

the term associated to subset S, meaning that Equation (6.6) can be rewritten $\sum_{i=1}^{r_i} \mu(S_i)\tilde{u_i}(S, a_i, a_i')$, and that the strategic move is decided by a weighted average of $\tilde{u_i}$ values.

We distinguish six possible cases:

- Case A:

$$\exists s \in S, u_i(a_i'|a_i, s) = 1 \text{ and } \forall s \in S, u_i(a_i'|a_i, s) \geq 0$$

- Case B:

$$\exists s \in S, u_i(a_i'|a_i, s) = -1 \text{ and } \forall s \in S, u_i(a_i'|a_i, s) \leqslant 0$$

- Case C:

$$\exists s \in S, u_i(a_i'|a_i, s) = -1 \text{ and } \exists s \in S, u_i(a_i'|a_i, s) = 1$$

- Case D:

$$\forall s \in S, u_i(a_i'|a_i, s) = 0$$

- Case E:

$$\forall s \in S, u_i(a_i'|a_i, s) = 1$$

- Case F:

$$\forall s \in S, u_i(a_i'|a_i, s) = -1$$

However, cases E and F can never happen for a non-singleton S, for the same reasons as the ones advocated in the proof of Theorem 76. We then get

$$\tilde{u_i}(S, a_i, a_i') = \begin{cases} 1-\alpha & \text{Case A,} \\ -\alpha & \text{Case B,} \\ 1-2\cdot\alpha & \text{Case C,} \\ 0 & \text{Case D} \end{cases}$$

It is clear that $\alpha > \frac{1}{2}$ implies that $\tilde{u_i}(S, a_i, a_i') \geq 0$ in cases A only.

Therefore, if

$$\alpha \cdot \underline{\mathbb{E}}_{\mathcal{M}}(u_i(a_i'|a_i, s))$$

$$+(1-\alpha) \cdot \overline{\mathbb{E}}_{\mathcal{M}}(u_i(a_i'|a_i, s)) \geq 0$$

then

$$\exists r_i' \leqslant r_i \quad \text{such that}$$
$$\exists s \in S_{r_i'}, u_i(a_i'|a_i, s) = 1$$
$$\text{and } \forall s \in S_{r_i'}, u_i(a_i'|a_i, s) \geq 0 \quad \text{(Case A)}$$

The last equivalence comes from the fact that only case A can lead to a positive $\tilde{u_i}(S, a_i, a_i')$, and that the criterion is an average of such $\tilde{u_i}$ values. Therefore, it has to exist $r_i' \leqslant r_i$ such that $\tilde{u_i}(S_{r_i'}, a_i, a_i') = 1 - \alpha$. If $R_i^+ = \{j : |\tilde{u_i}(S_j, a_i, a_i') = 1 - \alpha\}$ is the set of neighborhood indices in case A, we need for $\sum_{j \in R_i^+} \beta_j^i$ to be large enough for the decision criterion to be positive. In other words, we can accept a move with some case C only if there exists a local dominance move in a smaller neighborhood that receives enough evidence. Using the result from section VIII in Meir [2015] which tells us that the convergence holds for any level of uncertainty r_i and any starting point (even non-truthful states) for local dominance strategic behaviors, then we can do a bijection of our strategic moves and get the convergence also. □

6.4.2 . The Case of Pignistic Probability

Let us now consider the case of a pignistic probability on a neighborhood S_{r_i}, which is having a uniform distribution $\mathcal{U}$ over all single scores within S_{r_i} for any voter i. Note that for the rest of this chapter the notation card is used for cardinality. If S_{r_i} is our single set with positive mass, the corresponding strategic behavior is equivalent to have the following criterion:

$$a_i' \succeq_{pig} a_i$$
$$\text{iff} \quad \mathbb{E}_{\mathcal{U}}(u_i(a_i'|a_i, s)) \geq 0$$
$$\text{iff} \quad \mathrm{card}(s \in S \mid u_i(a_i'|a_i, s) = 1)$$
$$- \mathrm{card}(s \in S \mid u_i(a_i'|a_i, s) = -1) \geq 0$$

When this criterion is not strict, it is obvious that we have a cycle because the same voter could move back and forth between two indistinguishable states. Therefore, we consider the strict version of it, namely

$$a_i' \succ_{\mathcal{U}} a_i$$
$$\text{iff} \quad \mathrm{card}(s \in S \mid u_i(a_i'|a_i, s) = 1)$$
$$- \mathrm{card}(s \in S \mid u_i(a_i'|a_i, s) = -1) > 0$$

Proposition 78. *With a uniform pignistic criterion on a neighborhood of size 1 (with respect to the ℓ_1 distance), we are not guaranteed to reach convergence.*

Proof. Here is the counter example with a neighborhood of size 1. Let us take the following profile:

$$\begin{array}{ccccccc}
a & \succ_1 & b & \succ_1 & c & \succ_1 & d \\
c & \succ_2 & a & \succ_2 & b & \succ_2 & d \\
d & \succ_3 & c & \succ_3 & a & \succ_3 & b \\
c & \succ_4 & d & \succ_4 & a & \succ_4 & b \\
a & \succ_5 & c & \succ_5 & d & \succ_5 & b \\
d & \succ_6 & b & \succ_6 & a & \succ_6 & c \\
c & \succ_7 & b & \succ_7 & d & \succ_7 & a \\
d & \succ_8 & b & \succ_8 & a & \succ_8 & c \\
b & \succ_9 & d & \succ_9 & c & \succ_9 & a \\
b & \succ_{10} & d & \succ_{10} & c & \succ_{10} & a
\end{array}$$

Voter 8 will move from d to a with a cardinal difference of 1. For clarity, let us detail the computation of this first move: the original score is $s = (2, 2, 3, 3)$, so

$$\begin{aligned}
S_1 = \{&(2, 2, 3, 3), (1, 2, 3, 3), (2, 1, 3, 3), \\
&(2, 2, 2, 3), (2, 2, 3, 2), (3, 2, 3, 3), \\
&(2, 3, 3, 3), (2, 2, 4, 3), (2, 2, 3, 4)\}
\end{aligned}$$

that corresponds to add/remove one vote. When moving from d to a, Voter 8 is improved in the first state $(2, 2, 3, 3)$ as it becomes $(3, 2, 3, 2)$, and a is elected instead of c, which is better from voter 8's perspective. Voter 8 is also improved for states $(2, 1, 3, 3)$, $(2, 2, 3, 2)$ and $(2, 3, 3, 3)$, and deteriorated in states $(2, 2, 2, 3)$, $(2, 3, 3, 3)$ and $(2, 2, 3, 4)$. Other states are not impacted by this move. Then voter 2 will move from c to a with a cardinal difference of 2, voter 8 will move from a to d with a cardinal difference of 1 and finally voter 2 will move from a to c with a cardinal difference of 3. This creates a cycle, which prevents convergence. □

Of course, if the size of the neighborhood is larger, there is not much hope for convergence either. We think this example is quite interesting, in particular when put in perspective with our result of Theorem 76 about partitioned belief functions. Indeed, this latter result indicates that it is possible to consider a probability measure and a corresponding decision rule such that convergence holds, in the case where the probability measure is defined on an algebra coarser than the one induced by single score vectors. This indicate that the counter-example is not so much about having a probabilistic model itself than about the voters being perhaps too optimistic in their movement, hinting also at the fact that decision rules such as "maximax" ones, where one relies on the upper expectation alone are unlikely to lead to convergence of voting behaviors.

6.5 . Conclusion and Future Works

6.5.1 . Conclusion

In this chapter, we provide new tools to model uncertainty in strategic voting from a quantitative perspective. We think that our work substantially modifies our view on uncertainty models around broadcast polls. In particular, we have shown that they capture standard settings in a single framework, and allow to extend classical results to more complex situations and new strategic decision rules. Moreover, we have provided some convergence results in our model, generalizing results from the iterative voting literature.

In addition, we introduce a novel view on uncertainty coming from the voter's perspective in Examples 35 and 36. The uncertainty could come from the voters themselves because they may not be able to fully express their preferences.

This chapter concludes our exploration of the various aspects of voting outcome variability. We now outline some perspectives for future research related to this final chapter.

6.5.2 . Future Works

We think this opens up vast avenues of research, as to our knowledge no strategic voting frameworks considered voter induced uncertainty, perhaps because identifying a central, broadcast score is then an ill-defined problem.

- First, it is often the case that voters can have, at some stage, incomplete preferences or uncertain ideas about who they are going to vote for, and are perhaps expecting more information to be more decisive. We could however reframe the question of convergence of strategic results within this framework. Yet, in contrast with previous results where voters always receive as information a precise score issued from polls, and announce an updated precise ballot, they would start here from uncertain information about the voting result. Then, they may announce in the next round an updated uncertainty model in regard of the new information. Let us be a bit more concrete and illustrate this point with the following example.

  Example 38. *Let $m = 4$ and $n = 10$, $\mu_1(\{a, b\}) = 1$, $\mu_2(\{b, c\}) = 1$, $\forall i \in \{3, 4\}, \mu_i(\{b\}) = 1$, $\forall j \in \{5, 6, 7\}, \mu_j(\{c\}) = 1$, $\forall l \in \{8, 9, 10\}, \mu_l(\{d\}) = 1$. The joint belief function derived from this information over I^4_{10} is $\mathcal{M}(([0, 1], [2, 4], [3, 4], 3)\}) = 1$.*

  The question is how the belief of voters 1 and 2 are going to be affected by the broadcast of the poll, namely the joint belief $\mathcal{M}$. For voter 1, we might want to say that this belief is going to become $\mu_1(\{b\}) = 1$, as a

has no chances of being elected: voter 1, that was hesitating, may therefore report its vote to b. The situation is much less clear for voter 2, that may not want to make further commitments which is already an issue. Of course, this situation complexifies if we now consider that voters may provide more complex belief functions, such as the one in Example 36. One possible way to address this issue could be for instance to consider the recent framework introduced by Pomeret-Coquot et al. [2022], that considers game and strategic moves in which voters information is given in the form of belief functions, and where conditioning or updating of information is done in various ways. We however leave this endeavour for future work, as it would probably require to completely rethink and revisit the framework of strategic voting. Another added difficulty with respect to the standard framework is that the current decision/act of the voter (a_i) would have to be (re)defined, as well as what it means to move from it to consider another action (a_i'). One solution could be to consider that the voter is evaluating a set of potentially optimal actions, mirroring the fact that we are using decision rules such as $\succ_{pess}$ that results in partial orders among actions.

- Second, this work can be extended in several directions, notably by considering other voting rules. For instance, extensions to scoring rules such as Borda are likely feasible. However, it is worth noting that for many other voting rules, the notion of strategic behavior is less well defined, and convergence results remain scarce, even in classical settings.

- More generally, we believe that addressing social choice problems through the lens of uncertainty is particularly fruitful, as it allows one to account for both incompleteness and uncertainty in preferences. We therefore expect that such approaches could be applied to a wide range of settings in social choice theory, from voting to fairness problems and beyond.

Conclusion

In this thesis, we have studied various sources of variability in election outcomes. Each was examined through a consistent framework designed to provide concrete and quantitative insights. Indeed, throughout the thesis, probabilistic tools were employed to rigorously measure how and to what extent election outcomes can vary under different conditions.

In Chapter 3, we examined the probability of agreement between different positional scoring rules. We identified several structured preference distributions commonly used in social choice experiments in which the probability of agreement is significantly higher than under the impartial cultures previously studied in the literature. Our main results for Walsh's and Conitzer's distributions characterize the sets of positional scoring rules that converge to the same expected winner as the number of voters grows large. Importantly, we also demonstrate that this phenomenon is not only asymptotic since high levels of agreement already emerge in elections with a small number of voters. Furthermore, we show that such agreement is compatible with the winner of any Condorcet-consistent rule. We extended this analysis to other structured distributions, notably the Mallows distribution, under which we observed an even higher probability of agreement, and the Pólya-Eggenberger urn model, where the probability of agreement, while not perfect, can still be high in some cases. From an orthogonal perspective, this work also provided an opportunity to introduce new insights into how to sample preferences within the single-peaked domain. Then, we confronted all these theoretical results on agreement with real-world data. Overall, analyzing agreement under different preference distributions provides valuable insights into how the choice of voting rule contributes to the variability of election outcomes.

In Chapter 4, we adopted a different perspective on the variability of election outcomes by taking into consideration strategic voting in plurality elections. Specifically, we examined the diversity of winners that may arise under the plurality rule when voters behave strategically. To this end, we introduced the notions of possible and necessary winners in the context of iterative voting. Our experiments on the classical iterative voting model revealed that winner diversity is relatively rare. We provided a theoretical explanation for this phenomenon, showing that it is largely due to the frequent existence of a necessary winner, as the process often remains at equilibrium. From a computational point of view, we proved that determining whether a possible or a necessary winner exists belongs to two different complexity classes. Beyond these quantitative and computational insights, we also characterized the winner and showed that the probability of electing a Condorcet winner increases significantly compared to standard plurality rule without strategic behavior.

Altogether, these results help us understand how strategic voting can shape election outcomes.

We continued to study strategic voting in Chapter 5, focusing this time on the role of information provided to voters in plurality elections. Specifically, we examined whether a polling institute can influence the outcome of an election by manipulating the information it broadcasts. From a computational perspective, we resolved an open problem by establishing the NP-hardness of the unrestricted manipulation case, where the polling institute is allowed to send any score. We then turned to the restricted case, where only reasonable polls, that is, those close to the truthful one, are allowed. In this setting, we complemented the existing hardness results by providing parameterized complexity results, using as a parameter the distance between the truthful poll and the manipulated one. Building on probabilistic techniques from previous chapters, we further provided a quantitative analysis of the problem by evaluating how frequently manipulation is possible. Notably, we proved that the probability of successful manipulation converges to one as the number of voters increases, under a broad family of cultures. Additionally, we derived lower bounds that show this phenomenon already arises in elections of moderate size. Our results highlight a high probability of successful manipulation in the unrestricted case and a more nuanced behavior in the restricted case, where the feasibility of manipulation depends heavily on the allowed distance from the truthful poll. Nevertheless, when this distance is proportional to the number of voters, manipulation becomes highly likely. This study demonstrates that the variability of election outcomes can be significantly affected by an external agent, in this case the polling institute, through the strategic dissemination of information.

Finally, in Chapter 6, we relax the assumption that voters' preferences are complete and certain, and propose a model that incorporates uncertainty, the final source of variability in election outcomes that we consider. This more general framework allows us to account for voters with incomplete or uncertain preferences, as well as to model uncertainty in poll information in a quantitative way. The main contribution of this chapter is to establish a connection between two well-known and widely accepted models of strategic voting from the literature. It also extends the current framework in which iterative voting is guaranteed to converge, and provides a basis for modeling new strategic voting scenarios. This enriched model helps explain how uncertainty in voters' preferences can lead to variability in election outcomes.

To sum up, adopting a probabilistic approach to study the sources of variability in election outcomes allows us to gain concrete and quantitative insights into their impact. These tools offer a sense of proportion regarding the results, which is particularly meaningful when assessing voting procedures or the practical relevance of social choice problems.

While each chapter already provides specific and technical directions for future works, the aim here is not to repeat them all, but rather to explain how some of them illustrate the relevance of advocating for the inclusion of probabilistic approaches to outcome variability in computational social choice problems. We have already shown throughout this thesis that adopting such an approach can shed new light on several problems. Of course, empirical experiments can reveal certain tendencies, even if their conclusions often depend on specific parameters, such as the number of voters or candidates. By contrast, a probabilistic approach allows us to derive formal results that hold for any election size. To illustrate this point, we present two future works.

First, we believe that comparing the outcome of strategic voting with that of a classical voting rule can help us better understand the consequences of strategic behaviors. For instance, one could compare the outcome of plurality voting with strategic voters with the outcome of the Borda rule, and examine how often the two agree. Such an analysis could be used to test the hypothesis that strategic voting tends to favor more consensual candidates as the Borda rule does.

Second, we recall that we have proved an increase in Condorcet efficiency under impartial cultures for the classical iterative voting model. The gold standard for this line of research would be to establish this result under any reasonable culture and voter behaviors. If such a result could be achieved, it would illustrate a situation where a classical requirement of voting rules fails, namely the absence of strategic voting. However, a probabilistic approach makes it possible to uncover a positive insight, showing that strategic behavior can actually increase the probability of electing the Condorcet winner.

Another set of future directions that further illustrates the value of a probabilistic approach includes the quantification of strategic behavior in the mean and the probabilistic evaluation of voting axioms.

In the iterative voting literature, a standard requirement is the convergence of the strategic voting model, primarily to ensure that the procedure is computationally feasible. However, ensuring convergence often necessitates restricting the range of voter behaviors, which can reduce strategic deviations to specific and sometimes unrealistic cases. Yet in practice, cycles are rarely observed in experiments and do not align with what we typically expect in real-world settings. This observation suggests that developing a theory of iterative voting in the mean could be particularly impactful, as it would allow for more relaxed assumptions while more accurately reflecting observed behaviors. Such a perspective could pave the way for more general and realistic models for analyzing iterative voting.

The probabilistic evaluation of voting axioms is also a promising avenue, as it aims to reinterpret classical impossibility theorems from a more quantitative perspective. In the classical axiomatic framework, an axiom is considered violated as soon as there exists a single profile that fails to satisfy it. The probabilistic approach, by contrast, measures how frequently such violations occur, offering a new lens through which to understand these impossibility results. The key question then becomes whether these theorems represent strong impossibilities, that is, whether the proportion of profiles that violate the axioms is large, or whether such violations are rare and thus perhaps less concerning in practice.

Bibliography

Measuring nearly single-peakedness of an electorate: Some new insights. In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT 2021)*, pages 19–34. Springer, 2021.

Aditya Anand and Palash Dey. Distance restricted manipulation in voting. *Theoretical Computer Science*, 891:149–165, 2021.

Kenneth J. Arrow. A difficulty in the concept of social welfare. *Journal of Political Economy*, 58(4):328–346, 1950.

Kenneth J. Arrow, Amartya Sen, and Kotaro Suzumura. *Handbook of Social Choice and Welfare*. Elsevier, 2010.

Krishna B. Athreya. On a characteristic property of Pólya's urn. *Studia Scientiarum Mathematicarum Hungarica*, 4:31–35, 1969.

Hossein Azari Soufiani, David C Parkes, and Lirong Xia. A statistical decision-theoretic framework for social choice. *Advances in Neural Information Processing Systems*, 27, 2014.

John J. Bartholdi, Craig A. Tovey, and Michael A. Trick. The computational difficulty of manipulating an election. *Social choice and welfare*, 6:227–241, 1989.

Dorothea Baumeister and Jörg Rothe. Taking the final step to a full dichotomy of the possible winner problem in pure scoring rules. *Information Processing Letters*, 112(5):186–190, 2012.

Dorothea Baumeister, Ann-Kathrin Selker, and Anaëlle Wilczynski. Manipulation of opinion polls to influence iterative elections. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020)*, pages 132–140, 2020.

Sven Berg. Paradox of voting under an urn model: The effect of homogeneity. *Public Choice*, 47(2):377–387, 1985.

Andrew C Berry. The accuracy of the gaussian approximation to the sum of independent variates. *Transactions of the american mathematical society*, 49 (1):122–136, 1941.

Nadja Betzler and Britta Dorn. Towards a dichotomy for the possible winner problem in elections based on scoring rules. *Journal of Computer and System Sciences*, 76(8):812–836, 2010.

Duncan Black. On the rationale of group decision-making. *Journal of Political Economy*, 56(1):23–34, 1948.

Duncan Black. *The theory of committees and elections*. Cambridge University Press, 1958.

Niclas Boehmer, Piotr Faliszewski, Łukasz Janeczko, Andrzej Kaczmarczyk, Grzegorz Lisowski, Grzegorz Pierczyński, Simon Rey, Dariusz Stolicki, Stanisław Szufa, and Tomasz Wąs. Guide to numerical experiments on elections in computational social choice. In *Proceedings of the 33rd International Joint Conference on Artificial Intelligence (IJCAI 2024)*, pages 7962–7970, 2024.

Nicolas Boehmer, Robert Bredereck, Edith Elkind, Piotr Faliszewski, and Stanisław Szufa. Expected frequency matrices of elections: Computation, geometry, and preference learning. *Advances in Neural Information Processing Systems*, pages 31347–31359, 2022.

Carlo Bonferroni. Teoria statistica delle classi e calcolo delle probabilita. *Pubblicazioni del R Istituto Superiore di Scienze Economiche e Commericiali di Firenze*, 8:3–62, 1936.

Sylvain Bouveret, Renaud Blanch, Antoinette Baujard, François Durand, Herrade Igersheim, Jérôme Lang, Annick Laruelle, Jean-François Laslier, Isabelle Lebon, and Vincent Merlin. Voter autrement 2017 - online experiment. Dataset and companion article on Zenodo, July 2018. URL https://doi.org/10.5281/zenodo.1199545.

Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D Procaccia. *Handbook of computational social choice*. Cambridge University Press, 2016.

Simina Brânzei, Ioannis Caragiannis, Jamie Morgenstern, and Ariel Procaccia. How bad is selfish voting? In *Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI 2013)*, volume 27, pages 138–144, 2013.

Francesco Paolo Cantelli. Considérations sur la convergence dans le calcul des probabilités. In *Annales de l'institut Henri Poincaré*, volume 5, pages 3–50, 1935.

Ioannis Caragiannis, Ariel Procaccia, and Nisarg Shah. Modal ranking: A uniquely robust voting rule. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI 2014)*, volume 28, pages 616–622, 2014.

Swarnendu Chatterjee and Ton Storcken. Frequency based analysis of collective aggregation rules. *Journal of Mathematical Economics*, 87:56–66, 2020.

Vincent Conitzer. Eliciting single-peaked preferences using comparison queries. In *Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2007)*, pages 420–427, 2007.

Vincent Conitzer, Toby Walsh, and Lirong Xia. Dominating manipulations in voting with partial information. In *Proceedings of the 25th AAAI Conference on Artificial Intelligence (AAAI 2011)*, volume 25, pages 638–643, 2011.

Andreas Darmann, Julia Grundner, and Christian Klamler. Evaluative voting or classical voting rules: Does it make a difference? empirical evidence for consensus among voting rules. *European Journal of Political Economy*, 59: 345–353, 2019.

Marie Jean Antoine Nicolas de Caritat marquis de Condorcet. *Essai sur l'Application de l'Analyse à la Probabilité des Décisions Rendues à la Pluralité des Voix*. L'imprimerie royale, 1785.

Gerard Debreu et al. Representation of a preference ordering by a numerical function. *Decision processes*, 3:159–165, 1954.

Théo Delemazure and Sylvain Bouveret. Voter autrement 2022-the online experiment (" un autre vote"). *Dataset and companion article on Zenodo*, 2024.

Dieter Denneberg. *Non-additive measure and integral*, volume 27. Springer Science & Business Media, 2013.

Thierry Denoeux and Prakash P Shenoy. An interval-valued utility theory for decision making with dempster-shafer belief functions. *International Journal of Approximate Reasoning*, 124:194–216, 2020.

Lihi Dery, Svetlana Obraztsova, Zinovi Rabinovich, and Meir Kalech. Lie on the fly: Strategic voting in an iterative preference elicitation process. *Group Decision and Negotiation*, 28:1077–1107, 2019.

Sébastien Destercke and Didier Dubois. Special cases. *Introduction to Imprecise Probabilities*, (chapter 4):79–91, 2014.

Sébastien Destercke and Henri Surugue. An enriched model of strategic voting under uncertainty. In *16th Multidisciplinary Workshop on Advances in Preference Handling (M-PREF 2025)*, 2025.

Palash Dey, Neeldhara Misra, and Yadati Narahari. Complexity of manipulation with partial information in voting. *Theoretical Computer Science*, 726: 78–99, 2018.

Aryeh Dvoretzky, Jack Kiefer, and Jacob Wolfowitz. Asymptotic minimax character of the sample distribution function and of the classical multinomial estimator. *The Annals of Mathematical Statistics*, pages 642–669, 1956.

Florian Eggenberger and George Pólya. Über die statistik verketteter vorgänge. *ZAMM-Journal of Applied Mathematics and Mechanics/Zeitschrift für Angewandte Mathematik und Mechanik*, 3(4):279–289, 1923.

Edith Elkind and Martin Lackner. On detecting nearly structured preference profiles. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI 2014)*, volume 28, 2014.

Edith Elkind, Piotr Faliszewski, and Arkadii M Slinko. On the role of distances in defining voting rules. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, volume 10, pages 375–382, 2010.

Ulle Endriss, Svetlana Obraztsova, Maria Polukarov, and Jeffrey S Rosenschein. Strategic voting with incomplete information. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI 2016)*, 2016.

Gábor Erdélyi, Martin Lackner, and Andreas Pfandler. Computational aspects of nearly single-peaked electorates. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI 2013)*, volume 27, pages 283–289, 2013.

Carl-Gustav Esseen. On the liapunov limit error in the theory of probability. *Ark. Mat. Astr. Fys.*, 28:1–19, 1942.

Piotr Faliszewski and Ariel D Procaccia. AI's war on manipulation: Are we winning? *AI Magazine*, 31(4):53–64, 2010.

Piotr Faliszewski and Jörg Rothe. Control and bribery in voting. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 7, pages 146–168. Cambridge University Press, 2016.

Piotr Faliszewski, Edith Hemaspaandra, and Lane A Hemaspaandra. The complexity of manipulative attacks in nearly single-peaked electorates. In *Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge (TARK 2011)*, pages 228–237, 2011.

Rupert Freeman, Markus Brill, and Vincent Conitzer. General tiebreaking schemes for computational social choice. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015)*, pages 1401–1409, 2015.

Ehud Friedgut, Gil Kalai, and Noam Nisan. Elections can be manipulated often. In *Proceedings of the 49th Annual IEEE symposium on Foundations of Computer Science (FOCS 2008)*, pages 243–249, 2008.

Michael R. Garey and David S. Johnson. *Computers and intractability : A guide to the theory of NP-completeness*. W.H. Freeman, 1979.

Serge Gaspers, Victor Naroditskiy, Nina Narodytska, and Toby Walsh. Possible and necessary winner problem in social polls. In *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014)*, pages 613–620, 2014.

William V Gehrlein. Weighted scoring rules, the impartial culture condition, and homogeneity. *Quality and Quantity*, 20(1):85–107, 1986.

William V. Gehrlein. Obtaining representations for probabilities of voting outcomes with effectively unlimited precision integer arithmetic. *Social Choice and Welfare*, 19(3):503–512, 2002.

William V. Gehrlein. Condorcet efficiency and proximity to single-peaked preferences. In *Proceedings of the 3rd International Conference on Logic, Game Theory and Social Choice (LGS 2003)*, pages 282–283, 2003.

William V. Gehrlein and Peter C. Fishburn. Robustness of positional scoring over subsets of alternatives. *Applied Mathematics and Optimization*, 6(1):241–255, 1980.

William V. Gehrlein and Peter C. Fishburn. Scoring rule sensitivity to weight selection. *Public Choice*, 40(3):249–261, 1983.

William V. Gehrlein and Dominique Lepelley. The unexpected behavior of plurality rule. *Theory and decision*, 67:267–293, 2009.

William V Gehrlein and Dominique Lepelley. *Voting paradoxes and group coherence: the Condorcet efficiency of voting rules*. Springer Science & Business Media, 2010.

Allan Gibbard. Manipulation of voting schemes: a general result. *Econometrica*, 41(4):587–601, 1973.

Itzhak Gilboa and David Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153, 1989.

Judy Goldsmith, Jérôme Lang, Nicholas Mattei, and Patrice Perny. Voting with rank dependent scoring rules. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI 2014)*, volume 28, pages 698–704, 2014.

Umberto Grandi. Social choice and social networks. In Ulle Endriss, editor, *Trends in Computational Social Choice*, chapter 9, pages 169–184. AI Access, 2017.

Umberto Grandi, Andrea Loreggia, Francesca Rossi, Kristen Brent Venable, and Toby Walsh. Restricted manipulation in iterative voting: Condorcet efficiency and Borda score. In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT 2013)*, pages 181–192, 2013.

Benjamin Hadjibeyli and Mark C. Wilson. Distance rationalization of anonymous and homogeneous voting rules. *Social Choice and Welfare*, 52(3):559–583, 2019.

Noam Hazon, Yonatan Aumann, Sarit Kraus, and Michael Wooldridge. On the evaluation of election outcomes under uncertainty. *Artificial Intelligence*, 189: 1–18, 2012.

Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *The collected works of Wassily Hoeffding*, pages 409–426, 1994.

Leonid Hurwicz. Optimality criteria for decision making under ignorance. Technical report, Cowles Commission discussion paper, statistics, 1951.

Marcus Isaksson, Guy Kindler, and Elchanan Mossel. The geometry of manipulation—a quantitative proof of the Gibbard-Satterthwaite theorem. *Combinatorica*, 32(2):221–250, 2012.

Eric Kamwa and Vincent Merlin. The likelihood of the consistency of collective rankings under preferences aggregation with four alternatives using scoring rules: a general formula and the optimal decision rule. *Computational Economics*, 53(4):1377–1395, 2019.

Joshua Kavner and Lirong Xia. Strategic behavior is bliss: iterative voting improves social welfare. *In Proceedings of the 35th Conference on Neural Information Processing Systems (NeurIPS 2021)*, 34:19021–19032, 2021.

Joshua Kavner and Lirong Xia. Average-case analysis of iterative voting. *arXiv preprint arXiv:2402.08144*, 2024.

Kathrin Konczak and Jérôme Lang. Voting procedures with incomplete preferences. In *Proceedings of the 18th International Joint Conference on Artifical Intelligence (IJCAI 2005) in the multidisciplinary Workshop on Advances in Preference Handling*, volume 20, 2005.

Dominik Kreiss and Thomas Augustin. Undecided voters as set-valued information–towards forecasts under epistemic imprecision. In *Proceedings of the 14th International Conference of Scalable Uncertainty Management (SUM 2020)*, pages 242–250. Springer, 2020.

Dominique Lepelley. Condorcet efficiency of positional voting rules with single-peaked preferences. *Economic Design*, 1:289–299, 1994.

Dominique Lepelley and Vincent Merlin. Choix social positionnel et principe majoritaire. *Annales d'Economie et de Statistique*, pages 29–48, 1998.

Dominique Lepelley and Laurent Vidu. Règles positionnelles itératives, principe majoritaire et préférences unimodales. *RAIRO-Operations Research*, 34(3):347–362, 2000.

Dominique Lepelley, Ahmed Louichi, and Fabrice Valognes. Computer simulations of voting systems. *Advances in Complex Systems*, 3(01n04):181–194, 2000a.

Dominique Lepelley, Patrick Pierron, and Fabrice Valognes. Scoring rules, condorcet efficiency and social homogeneity. *Theory and Decision*, 49(2):175–196, 2000b.

Colin L Mallows. Non-null ranking models. I. *Biometrika*, 44(1/2):114–130, 1957.

Nicholas Mattei and Toby Walsh. Preflib: A library for preferences http://www. preflib. org. In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT 2013)*, pages 259–270, 2013.

Kenneth O. May. A set of independent necessary and sufficient conditions for simple majority decision. *Econometrica: Journal of the Econometric Society*, pages 680–684, 1952.

Reshef Meir. Plurality voting under uncertainty. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI 2015)*, volume 29, 2015.

Reshef Meir. Iterative voting. In Ulle Endriss, editor, *Trends in Computational Social Choice*, chapter 4, pages 69–86. AI Access, 2017.

Reshef Meir. Strategic voting. *Synthesis lectures on artificial intelligence and machine learning*, 13(1):1–167, 2018.

Reshef Meir. *Strategic voting*. Springer Nature, 2022.

Reshef Meir, Maria Polukarov, Jeffrey Rosenschein, and Nicholas Jennings. Convergence to equilibria in plurality voting. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI 2010)*, volume 24, pages 823–828, 2010.

Reshef Meir, Omer Lev, and Jeffrey S Rosenschein. A local-dominance theory of voting equilibria. In *Proceedings of the 15th ACM conference on Economics and computation (ACM 2014)*, pages 313–330, 2014.

Vincent Merlin, Maria Tataru, and Fabrice Valognes. On the probability that all decision rules select the same winner. *Journal of Mathematical Economics*, 33(2):183–207, 2000.

Enrique Miranda and Sébastien Destercke. Extreme points of the credal sets generated by comparative probabilities. *Journal of Mathematical Psychology*, 64:44–57, 2015.

Hervi Moulin. Axioms of cooperative decision making. *Cambridge Books*, 1989.

Vincent Mousseau, Henri Surugue, and Anaëlle Wilczynski. Do we care about poll manipulation in political elections? In *Proceedings of the 27th European Conference on Artificial Intelligence (ECAI 2024)*, pages 3541–3548, 2024.

Vincent Mousseau, Henri Surugue, Magdalena Tydrichová, and Anaëlle Wilczynski. On iterative voting outcomes in plurality elections. In *10th Workshop on Computational Social Choice Conference (COMSOC 2025)*, 2025a.

Vincent Mousseau, Henri Surugue, and Anaëlle Wilczynski. Agreement among voting rules under single-peaked preference distributions. In *Proceedings of the 28th European Conference on Artificial Intelligence (ECAI 2025) [To appear]*, 2025b.

Roger B Myerson and Robert J Weber. A theory of voting equilibria. *American Political science review*, 87(1):102–114, 1993.

Hannu Nurmi. *Comparing voting systems*, volume 3. Springer Science & Business Media, 2012.

Pierre Pomeret-Coquot, Hélène Fargier, and Érik Martin-Dorel. Games of incomplete information: A framework based on belief functions. *International Journal of Approximate Reasoning*, 151:182–204, 2022.

Ariel D. Procaccia and Jeffrey S. Rosenschein. Average-case tractability of manipulation in voting via the fraction of manipulators. In *Proceedings of the 6th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2007)*, pages 718–720, 2007.

Zinovi Rabinovich, Svetlana Obraztsova, Omer Lev, Evangelos Markakis, and Jeffrey Rosenschein. Analysis of equilibria in iterative voting schemes. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI 2015)*, volume 29, pages 1007–1013, 2015.

Michel Regenwetter, Aeri Kim, Arthur Kantor, and Moon-Ho R Ho. The unexpected empirical consensus among consensus methods. *Psychological Science*, 18(7):629–635, 2007.

Annemieke Reijngoud and Ulle Endriss. Voter response to iterated poll information. In *Proceedings of the 4th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, volume 4, page 8, 2012.

Andra Maria Roescu. Preferences, voting rules, behaviour and outcomes. a field experiment on the local elections in romania. *Romanian Journal of Society and Politics*, 9(1):9–24, 2014.

Donald G. Saari. *Basic geometry of voting*. Springer Science & Business Media, 2012.

Donald G Saari and Fabrice Valognes. The geometry of black's single peakedness and related conditions. *Journal of Mathematical Economics*, 32(4):429–456, 1999.

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.

Thomas A Severini. *Elements of distribution theory*. Cambridge University Press, 2005.

Lloyd S Shapley and Martin Shubik. A method for evaluating the distribution of power in a committee system. *American political science review*, 48(3): 787–792, 1954.

Philippe Smets and R Kennes. The concept of distinct evidence. In *Proceedings of the 4th International Conference on Processing and Management of Uncertainty in Knowledge-Based Systems (IPMU 1992)*, pages 789–794, 1992.

Xin Sui, Alex Francois-Nienaber, and Craig Boutilier. Multi-dimensional single-peaked consistency and its approximations. In *Proceedings of the 23th International Joint Conference on Artificial Intelligence (IJCAI 2013)*, volume 13, pages 375–382, 2013.

Stanisław Szufa, Piotr Faliszewski, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. Drawing a map of elections in the space of statistical cultures. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020)*, pages 1341–1349, 2020.

David RM Thompson, Omer Lev, Kevin Leyton-Brown, and Jeffrey Rosenschein. Empirical analysis of plurality election equilibria. In *Proceedings of* *the 12th international Conference on Autonomous agents and Multiagent Systems (AAMAS 2013)*, pages 391–398, 2013.

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

Matthias CM Troffaes. Decision making under uncertainty using imprecise probabilities. *International Journal of Approximate Reasoning*, 45(1):17–29, 2007.

Magdaléna Tydrichová. *Structural and algorithmic aspects of preference domain restrictions in collective decision making: contributions to the study of single-peaked and Euclidean preferences*. PhD thesis, Sorbonne Université, 2023.

Toby Walsh. Generating single peaked votes. *arXiv preprint arXiv:1503.02766*, 2015.

Anaëlle Wilczynski. Poll-confident voters in iterative voting. In *Proceedings of the 33th AAAI Conference on Artificial Intelligence (AAAI 2019)*, volume 33, pages 2205–2212, 2019.

Lirong Xia. Computing the margin of victory for various voting rules. In *Proceedings of the 13th ACM Conference on Electronic Commerce (ACM 2012)*, pages 982–999, 2012.

Lirong Xia. Statistical properties of social choice mechanisms. In *Proceedings of the 5th International Workshop on Computational Social Choice*, 2014.

Lirong Xia. The smoothed possibility of social choice. *Advances in Neural Information Processing Systems*, 33:11044–11055, 2020.

Lirong Xia. How likely are large elections tied? In *Proceedings of the 22nd ACM Conference on Economics and Computation (ACM 2021)*, pages 884–885, 2021.

Lirong Xia. The impact of a coalition: Assessing the likelihood of voter influence in large elections. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC 2023)*, pages 1156–1156, 2023.

Lirong Xia and Vincent Conitzer. Generalized scoring rules and the frequency of coalitional manipulability. In *Proceedings of the 9th ACM Conference on Electronic Commerce (EC 2008)*, pages 109–118, 2008.

Lirong Xia and Vincent Conitzer. Determining possible and necessary winners given partial orders. *Journal of Artificial Intelligence Research*, 41:25–67, 2011.

Table of Notations

This table summarizes the main notations used throughout the thesis.

Symbol	Meaning
Sets & basic parameters	
$[k] = \{1, \dots, k\}$	Set of positive integers from 1 to k.
n	Number of voters.
$N = [n] = \{1, \dots, n\}$	Set of voters.
m	Number of candidates.
$M = \{x_1, \dots, x_m\}$	Set of candidates.
$\vert \cdot \vert$	Cardinality of a set.
Preferences & profiles	
$x \succ_i y$	Voter i strictly prefers x to y.
$x \succeq_i y$	Voter i weakly prefers x to y.
$\succ_i$	Strict linear order of voter i over M.
$\mathscr{P} = (\succ_1, \dots, \succ_n)$	Preference profile of all voters.
$\Pi^m_{y \succ z}$	Set of orders on m candidates where y is ranked before z.
$[\succ]^\tau$	Preference order obtained from $\succ$ by renaming candidates with permutation τ.
$>$	Single-peaked axis (a total order) over M.
Voting rules & scores	
$\mathcal{F}(\mathscr{P})$	Voting rule mapping a profile to a set of winners.
$\mathcal{W}_\mathcal{F}(\cdot)$	Winner set under rule $\mathcal{F}$ (when used in the text).
$\mathcal{W}_P(b)$	Winner set under Plurality for ballot profile b.
$b_i \in M$	Ballot submitted by voter i (plurality).
$b \in M^n$	Ballot profile $(b_1, \dots, b_n)$.
$s_x(b)$	Plurality score of candidate x for ballot b
$\alpha = (\alpha_1, \cdots, \alpha_m)$	Scoring vector of a positional rule.
$\rhd$	Fixed lexicographic tie-breaking order on M.
Probability	
$\pi \in \Pi$	Distribution over preference orders of one voter.
$\mathbb{P}_C$	Probability distribution on profiles with culture C
$\mathbb{P}_{IC}$	Impartial culture probability distribution.
$\mathbb{P}_{IAC}$	Impartial anonymous culture probability distribution.
$\mathcal{W}_\pi(\mathcal{F})$	Expected winners of $\mathcal{F}$ under π

(continued on the next page)

Symbol	Meaning
Strategic voting	
PW_i^t	Potential winners for voter i at step t (in an iterative process).
PW^t	All potential winners at step t (union over voters).
$PW(s)$	Set of potential winners compatible with score vector s.
w^0	Initial truthful plurality winner.