



Compromise and elicitation in social choice : a study of egalitarianism and incomplete information in voting

Beatrice Napolitano

► To cite this version:

Beatrice Napolitano. Compromise and elicitation in social choice : a study of egalitarianism and incomplete information in voting. Artificial Intelligence [cs.AI]. Université Paris sciences et lettres, 2022. English. ⟨NNT : 2022UPSLD054⟩. ⟨tel-04123938⟩

HAL Id: tel-04123938

<https://theses.hal.science/tel-04123938v1>

Submitted on 9 Jun 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization



THÈSE DE DOCTORAT
DE L'UNIVERSITÉ PSL

Préparée à l'Université Paris-Dauphine

**Compromise and elicitation in social choice: a study of
egalitarianism and incomplete information in voting**

Soutenue par

Beatrice NAPOLITANO

Le 09 Décembre 2022

École doctorale n°543

École Doctorale SDOSE

Spécialité

Informatique

Composition du jury :

Antoinette BAUJARD
Professeure, Université Jean Monnet *Rapporteuse*

Olivier CAILLOUX
MdC, Université Paris-Dauphine *Examineur*

Ayça EBRU GIRITLIGIL
MdC, İstanbul Bilgi University *Présidente*

Vincent MOUSSEAU
Professeur, CentraleSupélec Université
Paris-Saclay *Rapporteur*

Remzi SANVER
Directeur de recherche, CNRS Univer-
sité Paris-Dauphine *Directeur de thèse*

Paolo VIAPPANI
Chargé de Recherche, CNRS Université
Paris-Dauphine *Examineur*

ACKNOWLEDGMENTS

While I am incredibly grateful to so many people who have supported me throughout my PhD journey, I have never been very good at expressing my gratitude in words. For this reason, as I write this section, I realize that one of the things I am grateful for is this very thesis for prompting me to take the time to express feelings that too often go unspoken. This process has allowed me to reflect on the incredible support and encouragement that I have received from so many people, and has reminded me of how fortunate I am to have such wonderful friends, family, colleagues, and mentors in my life.

The first people I would like to thank are the members of my jury for their valuable comments and insightful feedback during the revision and the defense of my thesis. Their constructive criticism and suggestions have been immensely helpful in improving the quality and clarity of my work. Without their guidance during the pre-defense of my thesis this dissertation would not be what it is today.

My deep gratitude goes to my supervisors, Remzi Sanver and Olivier Cailloux, for their unwavering support, advice, and encouragement throughout my research journey. Remzi's expertise, mentorship, and willingness to share knowledge have been instrumental in shaping the direction of my research. Olivier's dedication to work, guidance and attention to details have contributed in shaping who I am today as a researcher but also as a person. Olivier, even though I may have rolled my eyes at times, I will never forget your meticulous comments, jokes and our insightful discussions. Your words of encouragement, constructive feedback and your confidence in my abilities have been a source of motivation and inspiration, even during the most challenging times. Thank you both for being there even when I did not know I needed you.

I would also like to extend my appreciation to all the PhD students in our lab for creating such a positive and collaborative atmosphere. Most of these past few years have been marked by a pandemic that disrupted all of our lives and we have not been able to take advantage of the classical social events. However, this has not prevented the growth of important relationships and deep friendships that I hope will continue in the future. I also thank all the members of LAMSADE, including professors and administrative staff, for the warm welcome I received when I joined the lab. I feel grateful to have had the opportunity to work alongside such an amazing group of people.

I am especially grateful to Jérôme Lang, for helping me during a tough time when he himself was going through a very hard moment in his life. His generosity and kindness have been invaluable and I will never be able to fully express my gratitude.

Although I never told him, I would also like to thank Umberto Grandi, for responding to a stranger's email in the middle of the summer of some years ago asking for an internship for the following week. I consider this event as the sliding doors moment of my life. Umberto has been extremely kind and welcoming, taught me the fundamentals of a field that was completely new to me until then, and believed in me and supported me from the beginning. This opportunity not only allowed me to gain valuable experience and skills but also opened the door for my PhD.

Moreover, it allowed me to get to know the person who, years later, became my life partner.

Perhaps the greatest thanks goes to this very person, my partner in crime, my soulmate. Arthur, your presence in my life has been a constant source of comfort and strength, I am grateful for your love and patience during the most challenging times. You are my number one fan, my best supporter. You have stood by my side through the ups and downs, always offering a shoulder to lean on and you pushed me to keep going, even when I felt like giving up. I feel incredibly lucky to have you in my life. Thank you for being such an amazing partner and for always having my back. Thank you for being you.

Finally, I would like to express my deepest gratitude to my family and friends for their love, encouragement, and understanding during this challenging yet rewarding period of my life. Their support and belief in me have been a constant source of motivation and inspiration. Thanks to the old friends who have always stood by my side despite kilometers of distance, friendships made up of texts and video calls. Friendships where you don't hear from each other for weeks but a phone call is enough to shorten time and distance. Friendships made up of venting and support but also laughter and memes. Thanks to the new friends, who became my Parisian family since we met the first week I set foot in this new city. The friends who bring you a hot meal when you are sick and who come to help you move even when you don't ask them to. Thanks to my acquired family, Arthur's parents and siblings, who welcomed me into their family with open arms. Their kindness, warmth, and generosity have made me feel like a part of the family from the very beginning and they have shown me nothing but love and support.

And, last but not least, thanks to my family. Their sacrifices and support, have allowed me to pursue my dreams and achieve milestone after milestone in my academic career. Every departure has been heartbreaking, but my family has always been the first to encourage me to follow my ambitions. The packages that they periodically send from Cosenza to Paris are full of coffee, chocolate and sausages only to the outside eye, but in reality they are primarily full of love. I feel grateful and blessed to have such a wonderful family who have always been there for me, no matter the distance. Thank you for being my rock and my foundation, and for making me the person I am today.

CONTENTS

List of Figures	vii
List of Tables	ix
List of Abbreviations	xi
I Introduction	1
1 Introduction	3
1.1 Voting rules from practice to theory	4
1.2 Research questions and contributions	8
1.3 Organization of the Thesis	9
2 Preliminaries and Notation	11
2.1 Voting with ranked ballots	11
2.1.1 Axioms	12
2.1.2 Voting rules	15
2.2 Voting with rated ballots	21
3 History and Literature Review	25
3.1 Voting as a compromise	25
3.2 Median compromise: the Majority Judgment rule	32
3.3 Preference elicitation	35
II Contributions	41
4 Simultaneous Elicitation of PSR and Agent Preferences	43
4.1 Introduction	43
4.2 Social choice with partial information	45
4.3 Robust winner determination	46
4.4 Interactive Elicitation	49
4.5 Empirical Evaluation	52
4.6 Conclusions	56
5 Java Package to support Elicitation in Social Choice	57
5.1 UML	57
5.1.1 Preference Representation	58

CONTENTS

5.1.2	Knowledge Representation	58
5.1.3	Questions Representation	60
5.1.4	Preference Information	60
5.1.5	Regret Computation	63
5.1.6	Strategies	63
6	Majority Judgment winner determination under incomplete information	67
6.1	Introduction	67
6.1.1	Related work	68
6.2	Notation	69
6.2.1	Incomplete knowledge	70
6.3	Reasoning on incompleteness	71
6.4	Experimental results	75
6.4.1	Profile uniformly distributed	75
6.4.2	Profile real scenario	76
6.5	Conclusions	78
7	Compromising as an Equal-Loss Principle	81
7.1	Introduction	81
7.2	Basic notions and notation	84
7.3	Egalitarian versus Paretian compromises	84
7.3.1	Egalitarian compromises	84
7.3.2	Paretian compromises	85
7.4	Which Social Choice Rules (SCRs) are compromises?	86
7.4.1	Condorcet consistent rules	86
7.4.2	Scoring rules	87
7.4.3	BK compromises	88
7.4.4	Restrictions on sigma	89
7.5	Two-voters case	90
7.6	Concluding remarks	91
III	Conclusion	93
8	Conclusion	95
8.1	Summary of the contributions	95
8.2	Future work	97
	Bibliography	99
	Résumé long en français	111
8.1	Introduction	111
8.2	Simultaneous Elicitation of PSR and Agent Preferences	114
8.3	Majority Judgment winner determination under incomplete information	118

8.4	Compromising as an equal loss principle	121
8.5	Conclusions et Perspectives	123

LIST OF FIGURES

4.1	Profile representing a question to the chair in extended (a) and compact (b) form.	51
4.2	Average MMR in problems of size (5, 10) after k questions.	53
4.3	Average MMR (normalized by n) after k questions with Pessimistic strategy for different datasets.	55
5.1	UML class diagram of preferences representation.	59
5.2	UML class diagram of knowledge representation.	61
5.3	UML class diagram of questions representation.	62
5.4	UML class diagram of elicited knowledge obtained by asking questions.	62
5.5	UML class diagram of regret computation.	64
5.6	UML class diagram of strategies.	65
6.1	Average probability of missing the winner under uniform distribution of preferences, given $n = 100, m = 50$ and $k \in \llbracket 1, 25 \rrbracket$ for 1000 batches of 1000 runs.	77
6.2	Result obtained from UAV experiment by asking $n = 1147$ voters to judge $m = 12$ alternatives.	78
6.3	Average probability of missing the winner using a real case distribution of preferences, given $m = 12$ for $n \in \{10, 25, 50, 100, 250, 500\}$ and $k \in \llbracket 1, 5 \rrbracket$.	79
8.1	Profil représentant une question posée à la commission sous forme étendue (a) et compacte (b).	116
8.2	MMR moyen (normalisé par n) après k questions avec la stratégie pessimiste pour différents ensembles de données.	116
8.3	MMR moyen dans des problèmes de taille (5, 10) après k questions.	117
8.4	Probabilité moyenne de ne pas trouver le gagnant dans le cadre d'une distribution uniforme des préférences, étant donné $n = 100, m = 50$ pour 1000 lots de 1000 exécutions et $k \in \llbracket 1, 25 \rrbracket$.	120
8.5	Résultat obtenu à partir d'une expérience de vote en demandant à $n = 1147$ électeurs de juger $m = 12$ alternatives.	120
8.6	Probabilité moyenne de ne pas trouver le gagnant en utilisant une distribution des préférences de cas réel, étant donné $m = 12$ pour $n \in \{10, 25, 50, 100, 250, 500\}$ et $k \in \llbracket 1, 5 \rrbracket$.	120

LIST OF TABLES

2.1	Axioms satisfied by the voting rules described.	21
4.1	Average MMR in problems of size $(10, 20)$ after k questions assuming geometric weights.	53
4.2	Questions asked by Pessimistic strategy on several datasets to reach $\frac{n}{10}$ regret, columns 4 and 5, and zero regret, last two columns.	54
4.3	Average MMR in problems of size $(10, 20)$ after 500 questions, among which q_c to the chair.	56
6.1	Average probability of missing the winner $\mathcal{P}$, given $n = 100, m = 50$ considering 1000 complete profiles uniformly distributed and 1000 incomplete profiles for each of them.	76
6.2	Average probability of missing the winner using a real case distribution of preferences, given $m = 12$ for different values of k and n	79
8.1	MMR moyen dans des problèmes de taille $(10, 20)$ après 500 questions, parmi lesquelles q_c à la commission.	117

LIST OF ABBREVIATIONS

- AHP** Analytic Hierarchy Process. 38
- BK** Brams and Kilgour. iv, 21, 82, 88–92, 97, 122
- DM** Decision Maker. 38, 39
- FB** Fallback Bargaining. 20, 29, 31, 121
- MC** Majoritarian Compromise. 20, 29, 121
- MJ** Majority Judgment. 22–24, 30, 32–34, 67–69, 71, 77, 80, 96, 118
- PSR** Positional Scoring Rule. 16, 17
- PV** Pareto-and-Veto rules. 30, 31
- SCR** Social Choice Rule. iv, 12–15, 19, 20, 29, 81–92, 122
- SL** Shortlisting. 30, 31
- STV** Single Transferable Vote. 36
- UML** Unified Modeling Language. 57
- VR** Veto-Rank mechanism. 30, 31

PART I

INTRODUCTION

INTRODUCTION

Democracy is difficult, flawed and unstable. It involves barely distinguishable political parties taking part in lengthy, over-complicated and expensive decision-making processes. Trying to engage so many people with political issues seems to lead only to complexity and disagreement. So why bother?

— Michela Murgia, *Istruzioni per diventare fascisti*
(Trans. *How to be a Fascist: A Manual* by Alex Valente)

Collective decision-making is a process in which individual preferences are aggregated to form a single group choice. This broad definition allows many instances, that seem very different from each other, to fit into this category. Some examples include *fair allocation problems*—which address the problem of fairly dividing some resources among individuals who have different interests in such resources, e.g. houses allocation, creation of a working schedule etc.—*matching problems*—which deal with the problem of matching individuals from two separate groups considering their preferences, e.g. students to schools, tenants to houses etc.—and *judgment aggregation problems*—which try to group the beliefs of different individuals into a judgment that reflects the society as a single entity. In this thesis we will focus on another example of collective decision-making: *voting*. Given a set of individuals who express their preferences over a given set of candidates, how do we select the *best* candidate for the group? This is a very ancient dilemma that has been faced multiple times during centuries and that we are going to tackle throughout this manuscript. As we shall see, the answer depends on many factors, such as the information we have at our disposal —do we know the preferences of all voters with respect to all candidates? —and the properties we want the outcome to satisfy —what do we mean by the best candidate? Who decides what is best?

An important methodology in our work is to define desirable properties we want a rule to satisfy. This allows us, on one hand, to divide aggregation rules into classes of methods satisfying the same properties, and, on the other hand, to do the inverse process: starting from an already defined method to help the decision-maker understand its properties. The lack of information and the assistance to decision-makers are two of the themes explored in the course of this work.

This introductory chapter aims to introduce the contributions and organization of this dissertation, but also, and maybe more importantly, to introduce its motives. We have mentioned the analysis of aggregation processes through the definition of desiderata, but this axiomatic approach is very recent and started only after the publication of the Ph.D. thesis of Kenneth Arrow [1951]. The theory behind collective decision-making processes is, however, much older and, as

McLean [1990] wrote, “the theory of voting is known to have been discovered three times and lost twice”.

To understand what we are studying today, we think it is essential to take a look at the past. Starting from the first traces of decision-making processes in history, we consider early voting mechanisms that focused mainly on particular cases. From the analysis of specific instances, we move on to the start of *Social Choice Theory* as a formal scientific discipline. We then focus on how it has managed to evolve by incorporating computational aspects typical of the computer science field. Finally, we discuss how this led us to investigate the problems that will be described in the following chapters and we outline our contributions.

1.1 Voting rules from practice to theory

The first records of collective choice problems can be traced back to ancient Greece (ca. 500 BC), and the word democracy itself originated from the greek *dēmokratiā*: *dēmos* *people* and *kratos* *rule*. Many scholars, including Aristotle, have described in details the functioning of Athenian democratic institutions, however they never formally analyzed voting processes. This appears to be partly due to the fact that Athenian voters were asked to express their preferences over only two alternatives: the choices were often in the form of *yes/no*, *banish/not banish* [McLean and Urken, 1995]. Today we know that most of the problems in social choice arise when at least three alternatives are considered.

From ancient Rome to the Middle Ages: the precursors of social choice. A problem concerning the use of three alternative is exactly the reason that brought Pliny the Younger, a magistrate of Ancient Rome, to write a letter to his jurist friend Titius Aristo (105 AD). In the letter, that can be found translated in McLean and Urken [1995, Chapter 2], Pliny asks Titius advice about a problem that emerged during a vote in the Senate. A consul had been found dead and his servants were suspected of his death, the Senate had to decide whether to free them, to banish them or to put them to death. For the first time they were facing a problem with more than two choices and, with it, its related issues. Pliny explains that the first option was acclaimed by the largest number of senators but not by the majority of them. Thus, he asked the Senate to consider the alternatives as three separate ones on the principle of huge disparity in punishments, and tried to arrange a series of binary votes among the three verdicts. Pliny was aware of his own attempt to manipulate the elections in favor of his top choice, freeing the slaves, but so was one of his adversary, supporter of the death penalty. The latter knew that the most lenient sentence would have won in binary votes against banishment and death taken separately, so he withdrew his proposal that could not have defeated any of the others. The supporters of death penalty then moved their support to banishment that easily won, having now a strong majority. Note how by changing the way we aggregate the same set of preferences we can have completely different outcomes.

Pliny’s concern shows how some desiderata of electoral protocols are not only restricted to modern society. The desire for a transparent system that was difficult to manipulate was present also in medieval elections. However, after Pliny, the democratic process did not find much room in the public discussion. The only two exceptions of abstract works on voting rules in the Middle

Ages were the ones of Ramon Llull (ca. 1232-1316) and Nicolaus Cusanus (1401-1464). The lack of theoretic analysis led to the creation of overly complicated electoral processes in order to make manipulations so difficult as to be impossible in practice [Uckelman and Uckelman, 2010]. For example, Venice’s 1268 ordinance, which described the procedure to elect the Doge (the head of the state), involved nine stages from the elections of the first committee to the election of the Doge himself. The procedure, discussed in detail by Lines [1986], Coggins and Perali [1998], Mowbray and Gollmann [2007] can be summarized by this folk poem [Doglioni and Ziotti, 1666, p.79]:

<i>“Trenta elegge il consiglio.</i>	Thirty elect the board.
<i>Di quei nove hanno il meglio;</i>	Of those, nine prevail.
<i>Questi eleggon quaranta;</i>	These elect forty.
<i>Ma chi di lor si vanta</i>	But those of them who boast
<i>Son dodici che fanno</i>	are twelve, that elect
<i>Venticinque: ma stanno</i>	twenty-five: of these
<i>Di questi solo nove</i>	only nine remain;
<i>Che fan con le lor prove</i>	that make with their proofs
<i>Quarantacinque a ponto</i>	forty-five in total.
<i>De’ quali undici in conto,</i>	Of those, eleven
<i>Eleggon quarantuno,</i>	elect forty-one,
<i>Che chiusi tutti in uno,</i>	who closed all together,
<i>Con venticinque almeno</i>	with at least twenty-five
<i>Voti, fanno il sereno</i>	votes, make the serene
<i>Principe che corregge</i>	Prince who corrects
<i>Statuti, ordini e legge.”</i>	articles, orders and laws.

Wolfson [1899] and Keller [2014] describe more electoral processes in Italian communes.

At the end of the thirteenth century Ramon Llull designed and discussed two voting methods in his works: *Artifitium electionis personarum* [ca. 1274-83] (The method of the election of persons), *En qual manera Natana fo eleta a abadessa* [ca. 1283] (In which way that Natana was elected abbess which corresponds to the 24th chapter of his novel *Blaquerna*), and *De arte eleccionis* [1299] (On the method of election). Hägele and Pukelsheim [2001] present the original texts together with their English translation. Llull was a Catalan missionary and philosopher, and he was particularly interested in ecclesiastic elections. In the first two works we cited, Llull describes a precise schema of comparisons between each pair of candidates. The candidate winning the largest number of pairwise comparison was selected as a winner. The last method, described in *De arte eleccionis*, is also based on pairwise comparisons but with a system of successive eliminations. The candidate winning the last comparison was selected as a winner. While the first two methods can be viewed as equivalent to either a Condorcet’s method or a Borda count depending on the interpretation of Llull’s text, the third method is a Condorcet’s procedure [McLean, 1990]. As we will see in the course of this manuscript, a candidate is the Condorcet winner if she beats every other candidate in pairwise comparisons. A voting rule is a Condorcet procedure if it elects the Condorcet winner whenever it exists.

Unfortunately, his works have not received much attention, except from Nicolas Cusanus, a

German theologian and philosopher who studied Llull's *De arte electionis*. He rejected Llull's work and proposed his own voting procedure in a chapter of his work *De concordantia catholica* [ca. 1431-34] (Catholic concordance), which can be found in McLean and Urken [1995, Chapter 4]. Cusanus' main concern was to reduce manipulation and strategic voting and he recommended a Borda count method with secret voting. We must say that there is no evidence that Cusanus was aware of the other two methods proposed by Llull, but only of the Condorcet-like one. With this in mind, we can see that although both Llull and Cusanus wanted to reduce manipulation, the methods they proposed turned out to be quite different. This is probably due to the different contexts they were considering. Llull had in mind an election where members of an abbey had to select their own leader: voters were people who knew and trusted each other and would have lived together even after the vote. Cusanus, on the other hand, wrote *De concordantia catholica* while attending the Council of Basel: voters were people that would have met only once to elect the new head of the Catholic Church and had their own interests and strategies. This underlines a very important point about voting methods: context matters. If we were to define a voting procedure, the first thing to consider is "to elect what?"

Enlightenment and the revival of choice. To find new discussions on voting procedures we must skip the rest of the Middle Ages and move on to the Enlightenment. In particular, we must travel to the France of the late eighteenth century. This should come as no surprise since by that time, in France, sciences had been extremely well organized for at least a century already through the establishment of the "Académie royale des sciences" by Louis XIV in 1666 [Demeulenaere-Douyère, 1995]. Starting from 1753, membership to the Academy was decided through elections using plurality rule and in 1770 Jean-Charles de Borda argued that this method could lead to undesirable results. In particular, he showed how the candidate elected under plurality could lose in case of pairwise comparisons and he proposed, instead, a rank order count —what we now call Borda count which is also equivalent to the Cusanus method. Although Borda presented these observations to the Academy in 1770, elections continued to be conducted using plurality rule and his work was published many years later [de Borda, 1784]. One of the reasons for this delay has been attributed to one of his political opponents: Nicolas de Condorcet [Urken, 2004] —also known as Marquis of Condorcet. In fact, the latter had meanwhile become secretary of the Academy, presumably also thanks to the voting method in place, and was responsible for scheduling publications. The year after the publication of Borda's method, the Marquis of Condorcet also proposed a voting procedure, the one we now recognize as equivalent to Llull's method [de Condorcet, 1785]. After the revolution, the Academy was reformed into the so-called "Institut de France" and Napoleon appointed a commission to study the voting mechanisms to be used within it. The commission chose Borda's method but after a while Napoleon took the unilateral decision to switch back to the use of plurality.

Once again, all these works were ignored until the second half of the 20th century. One of the few noteworthy works is that of Charles Dodgson, better known as Lewis Carroll. Although not aware of the studies of the Marquis of Condorcet, in 1876 he proposed a voting procedure that solves the paradox arising in the Condorcet method when cycles occur. In his proposition the winner is the only candidate that wins every pairwise comparison (Condorcet winner). If such a candidate does not exist then for each candidate we count the number of changes in the

voter preferences required for them to be the Condorcet winner. The elected candidate is the one associated to the lowest number of changes.

The axiomatization of social choice theory. This concludes a brief introduction to the start of voting theory which has so far been concerned with analyzing procedures, and their related problems, on the basis of examples. From this point on, and especially since the publication of the Ph.D. thesis of Kenneth Arrow [1951], social choice theory emerges as a scientific discipline. Arrow, in fact, was the first to introduce axiomatic analysis of aggregation methods and to characterize classes of rules on the basis of the set of properties they satisfy. In his thesis, he also defines a small set of axioms that seem naturally desirable and then he proves that the only voting rule able to satisfy them all is the dictatorship. This opened the door to the research that followed over the next sixty years.

In the meantime, however, the aforementioned works were still unknown to the general public, until a few years later, when Duncan Black [1958] published the book that soon became the basis of Social Choice Theory: *The Theory of Committees and Elections*. From this point onward, new axioms are recurrently proposed and new electoral rules analyzed. In Chapter 2, we will describe the most relevant ones to this work.

For reasons of space, and complexity of the topic, many authors who have worked on electoral procedures —e.g. Pufendorf, Daunou, LaPlace, Lhulier, Nanson, Galton, Scitovsky, Bergson, Samuelson etc.—have not been mentioned in this chapter. More details on the history of Social Choice Theory can be found in Black [1958], McLean and Urken [1995], McLean [1990], Urken [2004]. Some other excellent sources for approaching the study of this discipline are: Arrow et al. [2002, 2011], Gaertner [2006], Fishburn [2015], Taylor [2005], Nitzan [2009].

The origin of computational social choice. Starting from the first traces of decision-making processes we have almost reached the present days. But we are still missing something to complete the picture. All rules mentioned so far never considered the computational cost of executing these procedures. Since Arrow, an aggregation method has been viewed as more or less reasonable on the basis of the axioms, the properties, it satisfies. However, the same rule may be proven impractical in its use [Bartholdi et al., 1989, Rothe et al., 2003, Hemaspaandra et al., 2005, Davenport and Kalagnanam, 2004]. Furthermore, most of the traditional rules require voters to express an order of preference over all alternatives. When the set of candidates is very large this comes at a huge cost, both from a cognitive and communicative point of view. If we were to decide where to have dinner, how would we expect all voters to order dozens, or even hundreds, of restaurants? What if some choices are not comparable [Pini et al., 2007]? Even assuming that voters are prepared to make such an effort, what is the cost of communicating all this information [Lang, 2004, Conitzer and Sandholm, 2005]? It is from all these questions (and many others such as the cost of manipulability for example) that many computer scientists have become interested in problems previously restricted to social theory. The first time the term *Computational Social Choice* was used to refer to this specific area of research was in 2006, during the first edition of the COMSOC workshop (Amsterdam). “*The Handbook of Computational Social Choice*” brings together the main notions of all major research sub-areas related to this new field [Brandt et al., 2016].

1.2 Research questions and contributions

Studying the history of aggregation methods, we realized that starting from a wide spectrum of procedures we can identify classes whose elements satisfy common properties. Given a method we can analyze the axioms it satisfies. The power of this is that we could also go the other way round: given some preferences about the properties that the rule must respect, we can help the decision-maker to define their desired aggregation method. One of the points we will focus on is precisely this:

- how can we help the decision-maker formally define an electoral rule on the basis of generic preferences about its properties?

In Chapter 4 we will study a situation in which a committee of non-experts has to decide how to aggregate the preferences of voters. Assume that the committee assigns a score to each position in the preference order of the voters. So, for example, the first ranked gets 10 points, the second 5 etc. Imagine now that the committee wishes the score of the best choice to be “much higher” than the one of the third best choice. How do we translate this “much higher” into a voting rule? What exactly does it mean? In this thesis we will focus on a particular class of rules, *Positional Scoring Rules* —which we will describe in detail in Chapter 2 —and develop a method for asking questions to the committee in terms of picking winners in an example profile. So, we transform a complex question involving differences of weights between consecutive positions into questions more easily answerable by a human committee such as “who should win in this profile”? From the answer to this question we deduce the answer to the original question.

Note how an important concept that we encountered when discussing Lull and Cusanus returns here. Both wanted to solve the same problem by proposing two different methods. This is because they were considering the problem within different contexts. Context is extremely important when defining a voting rule; there is no best rule that we can choose to apply in every case. That is why it makes sense to ask questions to the committee and really understand what it means in their context for an alternative to be ranked first rather than third. The answers could vary a lot depending on what kind of preferences are being aggregated and what is being elected. This idea will come back later when we will discuss Chapter 7.

Another important problem —which we mentioned introducing the field of Computational Social Choice—when deciding on a voting rule is its cost. This is a point that we also find in our first research question. In this case it is represented by the cognitive cost for non-experts to formalize a voting rule on the basis of some generic preferences. But, as we know, the cost is also related to voters when they have to order a large set of preferences, and computational in transmitting all this information. This awareness leads us to study strategies of voters preferences elicitation:

- how can we acquire the most relevant information at the lowest cost?

We study this question for different voting rules (Positional Scoring Rules and Majority Judgment) in Chapters 4 and 6. In particular, in Chapter 4, we combine this research question with the first one and we try to understand whether, in a context of zero information, interleaving the elicitation of voter preferences and the elicitation of the voting rule yields better result than

a more linear approach. In Chapter 6 we will study concrete scenarios taken from real voting situations, in which voter preferences are obtained by asking a given number of random questions to voters. Among other things, we will investigate how to identify the minimum number of questions to be asked so that the probability of eliciting the winner, using the Majority Judgment rule, is high.

Coming back to the idea “different contexts, different definitions” we began to wonder how different situations could justify the use of different voting rules that seek to achieve the same goal. In the literature, various notions of compromise have been proposed, and with them various voting rules that more or less attempt to implement these definitions. When we talk about compromise what we mean depends very much on the context, when we want to decide where to dine we are probably fine with voters trying to find a middle ground even if the outcome does not turn out to be a compromise. This is the case with some voting rules that we will discuss later on, such as Fallback Bargaining, where it may happen that the outcome is the best choice for some voters but highly disliked by many others. However, when we deal with contexts where egalitarianism—which we will interpret as everyone conceding equally—is a major concern, this notion of compromise is not acceptable anymore. On these bases, we develop the grounds for our concept of compromise:

- how can we define the notion of compromise where each voter concedes as equally as possible?

In Chapter 7 we address this question and analyze which of the existing rules, if any, satisfy this new definition. We will see that different definitions make sense in different contexts and that one is not necessarily better than the other.

1.3 Organization of the Thesis

In what follows we will describe how this manuscript is organized.

Chapter 2 collects the important notions that are the basis of the themes discussed in this thesis and that will return recurrently in subsequent chapters. We will describe the difference between voting with different types of ballots: ranked ballots, which we will use in Chapters 4 and 7, and rated ballots, which we will consider in Chapter 6. We will describe the voting rules we will use in our contributions and some of the axioms that characterize them.

In Chapter 3, using the same approach used in this introduction, we will dig into the past to position our work. In particular, in Section 3.1, starting from the meaning of compromise we will trace its use throughout history. We will analyze the definitions that have been given of compromising rules in social choice theory and point out how some of them do not represent certain ideas of compromise. We will show how our approach fits within this literature. Noting that the Majority Judgment rule is seen as a form of compromise in the literature, we will delve into this rule in Section 3.2. We will study its introduction, its uses and also its criticisms. As the last topic of this chapter, we cover in Section 3.3 another important aspect of this thesis: Preference Elicitation. In fact, if so far it has been assumed to know the preferences of voters and the voting procedure, this cannot be taken for granted. That of preference elicitation is a well studied problem and we will describe the different approaches by which it has been tackled

in the literature. In addition, we will mention how it has been addressed in different fields, particularly in the decision aiding and machine learning literature, and the similarities with our approach.

Part II includes our contributions. Specifically, Chapter 4 is an article published in an international conference [Napolitano et al., 2021] and Chapter 7 is an article published in a journal [Cailloux et al., 2022]. Chapter 5 accompanies Chapter 4 by describing the code provided in support of the contribution. Finally, Chapter 6 is an original contribution, not yet published, whose goal is to study the consequences of the elicitation process in a voting situation that uses Majority Judgment.

To conclude, Chapter 8 provides a summary of the contributions of this thesis and some perspectives for future directions.

PRELIMINARIES AND NOTATION

In the introductory chapter we have mentioned different concepts like voting, preference aggregation processes, voting rules. All these terms are rather vague if we do not provide precise definitions. In this chapter we will present the main notions we will use throughout the course of the dissertation and assign them precise meanings.

It is worth specifying that when we talk about voting we imply an aggregation of votes expressed through a ballot. The term ballot comes from the ancient Italian *ballotta*, a small ball used in medieval elections to cast a vote. We can already find references to golden or silver balls used by the members of the various committees in the long procedure of electing the Doge after Venice’s ordinance [Doglioni and Zioti, 1666]. Over time, the term has assumed the generic meaning of a method to cast a vote (often secret). In this dissertation we mainly use linear orders of preferences as ballots, i.e. each voter specifies which candidate is her best choice, her second choice and so on. However, this representation is often blamed as the reason for the impossibility results described by Arrow [1950]. Other types of ballots have been proposed, for example Balinski and Laraki [2011] consider voters’ judgments on candidates rather than preferences. To give an idea of the difference between these two representations we can imagine that in the case of ordered ballots a voter would say “I prefer alternative a to alternative b ” while using judgments he would say “I consider alternative a excellent and alternative b mediocre”. The latter notion will be explored in Section 2.2 and further on in Chapter 6.

A set of ballots constitutes the input of the aggregation function, since there are different representations of this object, this means that we can define different functions on different inputs. But what is the output? Here again, there is no single answer. One can consider *resolute* functions, i.e. functions that return a unique winner. In such cases, whenever there are ties between candidates, mechanisms are designed to break them. However, this is not the assumption we will make in the course of this manuscript. The functions we will describe return the set of winning candidates.

It is important to emphasize that since these functions can be defined on different inputs and outputs, they are different mathematical objects. If we define a desirable property of a function that uses ranked ballots it might not be possible to convert it into a property of a function that uses judgments, and vice-versa.

2.1 Voting with ranked ballots

A social choice setting consists of a set N of n individuals—or voters—who have preferences over a set A of m alternatives. A social choice problem deals with aggregating those preferences, through the use of a social choice rule, to obtain a common choice for the group. The preferences

of individuals are represented by linear orders over the set of alternatives A . In general, a linear order $\succ$ on a set X is a relation that is:

- irreflexive: $a \not\succ a, \forall a \in X$
- transitive: $a \succ b \wedge b \succ c \Rightarrow a \succ c, \forall a, b, c \in X$
- complete: $a \succ b \vee b \succ a, \forall a \neq b \in X$

We can also define a non-strict linear order $\succeq$ on a set X as a relation that is transitive, complete, reflexive ($a \succeq a, \forall a \in X$) and antisymmetric ($a \succeq b \wedge b \succeq a \Rightarrow a = b, \forall a, b \in X$). For each non-strict linear order $\succeq$ there exists an associated linear order $\succ$ obtained by removing all pairs $a \succeq a, \forall a \in X$. In other words, we have $a \succ b \iff a \succeq b \wedge a \neq b, \forall a, b \in X$.

$\mathcal{L}(A)$ denotes the set of all linear orders over A . A generic element $\succ_i \in \mathcal{L}(A)$ stands for the *preference ranking* of the individual $i \in N$. A *profile* $P : N \rightarrow \mathcal{L}(A)$ associates with each individual $i \in N$ a preference $P(i) = \succ_i$. We will use $\succ_i$ instead of $P(i)$ when the profile is clear from the context. $\mathcal{L}(A)^N$ denotes the set of all functions from N to $\mathcal{L}(A)$, thus all possible profiles given N .

Example 2.1. Given a set of alternatives $A = \{a, b, c\}$ and a set $N = \{v_1, v_2, v_3, v_4\}$ of voters, we represent a profile P as:

$$\begin{array}{cccc} v_1 & v_2 & v_3 & v_4 \\ a & b & c & c \\ b & c & a & a \\ c & a & b & b \end{array},$$

where each column is the preference ranking associated to a certain voter. In this example the voter v_1 prefers the alternative a over b over c , v_2 prefers b over c over a , and both v_3 and v_4 prefer c over a over b .

A *Social Choice Rule* (SCR) is a mapping $f : \mathcal{L}(A)^N \rightarrow 2^A \setminus \{\emptyset\}$ that associates to each profile a set of (tied) winners. As we saw in Chapter 1, scholars have realized that the most efficient way to define aggregation functions is through the formulation of desirable properties. These properties are called *axioms* and they are helpful in dividing rules into classes whose elements satisfy the same set of axioms. Or, in the same way, one can prove that a given class is empty, i.e. that it is impossible to satisfy a set of axioms simultaneously. As we shall see, axioms that independently seem to be reasonable requirements can result, when combined, in undesirable effects that are often difficult to detect without an systematical analysis.

In what follows we will describe the axioms and the voting rules that we will consider through the rest of this dissertation. In particular we will divide the ones based on preference rankings from the ones using judgments as ballots, and we will see which axioms each rule satisfies.

2.1.1 Axioms

As already mentioned, an axiom formalizes an appealing property of a voting rule. A first intuitive desideratum is that the rule should not favor specific voters or candidates. In many real cases, when a committee has to make a decision, in case of a tie a member, usually the chair, has

the power to break it at his discretion. Similarly, in many legislatures in case of a tie between approving a law or not, the status quo wins, highlighting an inequality between the alternatives. However, these behaviors are not always desired or accepted. In what follows we present two axioms that promote equality among voters, *anonymity*, and among candidates, *neutrality*.

Anonymity. An SCR f is anonymous iff for any $i, j \in N$ and $P, P' \in \mathcal{L}(A)^N$ such that $\succ'_k = \succ_k, \forall k \in N \setminus \{i, j\}$ and $\succ'_i = \succ_j, \succ'_j = \succ_i$, we have that $f(P) = f(P')$. In other words, if the identity of voters is irrelevant. $\square$

Neutrality. An SCR f is neutral iff for any $P, P' \in \mathcal{L}(A)^N$ such that P' is obtained from P by swapping the position of two alternatives in every preference rankings, we can obtain $f(P')$ from $f(P)$ via the same swap. In other words, if the identity of alternatives is irrelevant. $\square$

Example 2.2. With this in mind we can rewrite Example 2.1 by considering anonymity. In fact, if we are not interested in the identity of the voters, P can be expressed as:

$$\begin{array}{ccc} 1 & 1 & 2 \\ a & b & c \\ b & c & a \\ c & a & b \end{array},$$

where each column represents the preference ranking associated to a certain number of voters. Here, one voter prefers the alternative a over b over c , one voter prefers b over c over a , and two voters prefer c over a over b . This way of denoting a profile will be used often throughout the dissertation.

We could also consider a weaker form of anonymity were the voters are not necessarily equal but which does not allow for a dictator—a specific voter whose preferred choice is always the outcome of the voting rule.

Nondictatoriality. An SCR f is nondictatorial iff $\forall i \in N$ there exists a profile $P \in \mathcal{L}(A)^N$ such that $\succ_i(1) \notin f(P)$; where $\succ_i(1)$ is the first element of the linear order representing i preferences. In other words, if no voter acts as a dictator. $\square$

Another fairly obvious property that might come to mind is that if *enough* voters agree on an alternative then it must be selected as the winner. To better define the concept of *enough*, given a profile $P \in \mathcal{L}(A)^N$ and an alternative $a \in A$, we consider the number of preference rankings for which a is ranked first: $\eta(a) = |\{\succ_i, i \in N \mid \succ_i(1) = a\}|$.

Majority. An SCR f is majoritarian iff for any $P \in \mathcal{L}(A)^N$ and $a \in A$ such that $\eta(a) > \frac{n}{2}$ then $f(P) = \{a\}$. In other words, if the majority of voters rank the same alternative first, then that alternative should be the sole winner. $\square$

This idea leads us to two more general considerations. First, if all voters prefer an alternative a to another one b then this preference must be translated into the result: b cannot be part of the winners. Second, if a majority of voters prefer an alternative a to any other, then a should be

the winner. To better define this last concept, given a profile $P \in \mathcal{L}(A)^N$ and two alternatives $a, b \in A$, we denote with $\mu(a, b) = |\{\succsim_i, i \in N \mid a \succsim_i b\}|$ the number of preference rankings for which a is preferred to b .

Pareto Property. *Given a profile $P \in \mathcal{L}(A)^N$ and any distinct $a, b \in A$, a Pareto dominates b at P (or equivalently b is Pareto dominated by a at P) iff $a \succ_i b, \forall i \in N$.*

We denote by $PO(P) = \{a \in A \mid \forall b \in A \setminus \{a\}, \exists i \in N, a \succ_i b\}$ the set of Pareto optimal alternatives at P .

An SCR f is Paretian iff $f(P) \subseteq PO(P) \forall P \in \mathcal{L}(A)^N$. In other words, f is Paretian if $f(P)$ never contains Pareto dominated alternatives. $\square$

Condorcet criterion. *Given a profile $P \in \mathcal{L}(A)^N$, an alternative $a \in A$ is the Condorcet winner iff $\mu(a, b) > \mu(b, a), \forall b \in A \setminus \{a\}$. In other words, if a beats every other alternatives in pairwise comparisons.*

An SCR f is a Condorcet procedure if it elects the Condorcet winner whenever it exists. $\square$

Note that for an SCR f being majoritarian implies that there is an alternative $a \in A$ that is preferred to any other by a majority of the voters: $\mu(a, b) > \frac{n}{2}, \forall b \in A \setminus \{a\}$. This also means that $\mu(a, b) > \mu(b, a), \forall b \in A \setminus \{a\}$ so that a is the Condorcet winner. The converse is not always true, if an alternative is the Condorcet winner then it is not necessarily the majority choice.

Example 2.3. *Consider the following profile P :*

1	1	1
a	c	d
b	b	b
c	a	a
d	d	c

$$\begin{aligned}
 \mu(a, b) &= 1, & \mu(a, c) &= 2 & \mu(a, d) &= 2 \\
 \mu(b, a) &= 2, & \mu(b, c) &= 2 & \mu(b, d) &= 2 \\
 \mu(c, a) &= 1, & \mu(c, b) &= 1 & \mu(c, d) &= 2 \\
 \mu(d, a) &= 1, & \mu(d, b) &= 1 & \mu(d, c) &= 1
 \end{aligned}$$

b is the Condorcet winner but it is not the majority choice.

We conclude that the Condorcet criterion is a stronger axiom and its satisfaction implies that of the majority criterion: if f is a Condorcet procedure then f is also majoritarian.

It is also worth to mention that the Condorcet winner does not always exist and that, with three or more alternatives, *majority cycles* may occur.

Example 2.4. *Consider the following profile P :*

$$\begin{array}{ccc} 1 & 1 & 1 \\ a & b & c \\ b & c & a \\ c & a & b \end{array},$$

$$\begin{aligned} \mu(a, b) &= 2, & \mu(b, a) &= 1 \\ \mu(b, c) &= 2, & \mu(c, b) &= 1 \\ \mu(c, a) &= 2, & \mu(a, c) &= 1 \end{aligned}.$$

Thus, a is preferred to b , b is preferred to c , but c is preferred to a , so there is no Condorcet winner. Such cycle of preferences is also known as **Condorcet's paradox**.

In all properties considered so far we have assumed that voters report their sincere preferences. But what if a voter can achieve better results by casting an insincere ballot rather than expressing her true preferences?

Manipulability. An SCR f is manipulable iff there exist two profiles $P, P' \in \mathcal{L}(A)^N$ and a voter $j \in N$ such that $\succ'_i = \succ_i, \forall i \in N \setminus \{j\}$ and $f(P') \succ_j f(P)$. In other words, if there exists a voter who by changing her preference ranking can change the outcome to something more appealing to her. $\square$

Strategy-Proofness. An SCR f is strategyproof iff it is not manipulable. $\square$

2.1.2 Voting rules

Having established desirable properties for aggregating preferences, we can now describe some methods designed for this purpose. Probably, one of the most intuitive approaches when a group of people get together and wish to find a common choice, is to elect the alternative supported by the majority of voters. The first thing we note is that this rule, as formulated, would not allow for the existence of multiple winners. Thus, while it might produce acceptable results with only two alternatives and an odd number of voters, in all other cases it would not always yield a winner.

Nevertheless, let us step back from these considerations for a moment and assume that there is always a majority that agrees on the choice of the winner. What would the decision of the majority imply? In his rules for a open society, Popper [1945, vol. 2, ch. 19] writes: “*Democracy cannot be the majority, although fully characterized as the rule of the institution of general elections is most important. For a majority might rule in a tyrannical way. (The majority of those who are less than 6 ft. high may decide that the minority of those over 6 ft. shall pay all taxes.)*”. Also thanks to his example, it is easy to imagine what Popper meant when he said that the majority can rule tyrannically. But one thing that is also very important to point out is that a sort of tyranny could also lie in the election itself, in the sense that the majority choice could be hated by the rest of the population. Imagine a situation where in a group of 100 voters, 51 of them have the following preference ranking $a \succ b \succ \dots \succ z$ and 49 instead $z \succ y \succ \dots \succ a$. A majority of 51 individuals would elect the candidate a who is detested by the remaining 49.

This makes us realize that rather intuitive methods, which seem fair at a first glance, can actually have unintended implications. To mitigate this problem and analyze the properties of voting rules, we can make use of the axioms just described.

The literature on voting rules is extensive and many have been proposed over the years, here we will present just a few examples of three particular categories: *Positional Scoring Rule*, that are methods assigning points to positions and electing the candidate with the greatest cumulative score; *Condorcet Procedures*, which are methods satisfying the Condorcet criterion, i.e. always electing the Condorcet winner whenever it exists; *BK Compromises*, which is the term we use to refer to the generalization of bargaining rules explored by Brams and Kilgour [2001], where, starting from the top choice of each voters, we fall back until one or more candidates have *enough* support to be considered winner.

Let us analyze these procedures in more detail.

Positional Scoring Rule (PSR)

The PSR f_w is the voting rule defined by the *scoring vector* $w = (w_1, \dots, w_m)$ which associate weights $w_r \in \mathbb{R}$ to positions, with $w_1 \geq w_2 \geq \dots \geq w_m$. Let α_r^a be the number of times that alternative a was ranked in the r -th position, then the score of a is defined by

$$s(a) = \sum_{i \in N} w_{\succ_i(x)} = \sum_{r=1}^m \alpha_r^a w_r. \quad (2.1)$$

The winners $f_w(P) = \{a \in A \mid s(a) \geq s(b), \forall b \in A \setminus \{a\}\}$ are the alternatives with highest score, we will write this as $f_w(P) = \operatorname{argmax}_{a \in A} s(a)$. Different scoring vector w define different rules, here we describe three of them.

Plurality.

The vector defining plurality rule is $w = (1, 0, \dots, 0)$. This means that the score of each candidate corresponds to the number of times she was ranked first $s(a) = \eta(a)$. The winners are the candidates who are ranked first by the largest number of voters.

Anti-Plurality.

The vector defining anti-plurality rule is $w = (1, \dots, 1, 0)$. The winners are the candidates who are ranked last by the smallest number of voters.

k-Approval.

Plurality and Anti-Plurality can be seen as particular cases of a more general approval voting where each voter assigns a score of 1 to, respectively, only one candidate and $m - 1$ candidates. For a generic $k \in \llbracket 1, m - 1 \rrbracket$, the vector defining k-approval rule is $w = (\underbrace{1, \dots, 1}_k, \underbrace{0, \dots, 0}_{m-k})$.

Borda.

The vector defining Borda rule is $w = (m - 1, m - 2, \dots, 0)$. This means that, for each preference ranking, each candidate gets points inversely proportional to her ranking. The winners are the candidates with the highest score. As we mentioned in Chapter 1, the Borda count is also known as Lull's method.

These are only few examples of scoring rules, but more can be defined by changing the vector. In some instances of real applications:

- the Eurovision Song Contest uses a PSR with scoring vector $w = (12, 10, 8, 7, 6, 5, 4, 3, 2, 1, 0, \dots, 0)$;
- the Formula One racing uses a PSR with scoring vector $w = (25, 18, 15, 12, 10, 8, 6, 4, 2, 1, 0, \dots, 0)$

Condorcet Procedures

Any method satisfying the Condorcet criterion is a Condorcet procedure. As we saw in the previous section, however, the Condorcet winner does not always exist. In these cases different ways to solve this ambiguity have been used, resulting in different procedures. These rules proposed over the years are many and go beyond the scope of this thesis, here we will describe only two of them but, for more details, Fishburn [1977] describes and compares nine famous Condorcet procedures.

We would like to warn the reader, that these two specific rules will not be mentioned again in the course of the thesis. The following material could be skipped (until the next set of BK rules) without compromising the understanding of the rest of the manuscript. However, it seems worthwhile to us to give an example of Condorcet procedures and, in particular, two methods which, as we shall see, do not satisfy the same axioms.

Given a profile $P \in \mathcal{L}(A)^N$ and two alternatives $a, b \in A$, recall that $\mu(a, b) = |\{ \succ_i, i \in N \mid a \succ_i b \}|$ represents the number of preference rankings for which a is preferred to b . We denote with $\triangleright$ a simple majority win, and we say that $a \triangleright b$ iff $\mu(a, b) > \mu(b, a)$.

Copeland.

Copeland's method—which was also proposed by Lull as we saw in Chapter 1—is often considered to be an extension of Condorcet's method. In fact, the idea of this procedure is based on the fact that if a candidate whose pairwise wins are greater than anyone else is deserving of winning, then one whose number of pairwise wins minus pairwise losses is greater than everyone else is even more deserving. Let us formalize this concept, we define with $\gamma(a) = |\{b \in A \setminus \{a\} \mid a \triangleright b\}| - |\{b \in A \setminus \{a\} \mid b \triangleright a\}|$ the number of simple majority wins of a minus the number of simple majority losses of a . The winners under the Copeland rule are the alternatives for which this value is the greatest $f_C(P) = \operatorname{argmax}_{a \in A} \gamma(a)$.

Schwartz.

The idea behind the procedure proposed by Schwartz [1972] is to identify the sets of alternatives such that every member in the set wins in simple majority all other alternatives outside the set. The union of all these sets represents the Schwartz set. We say that a Schwartz dominates b , $a \vdash b$, iff $\exists x_1, \dots, x_k \in A \mid a \triangleright x_1 \wedge x_1 \triangleright x_2 \wedge \dots \wedge x_k \triangleright b$ and $\nexists y_1, \dots, y_l \in A \mid b \triangleright y_1 \wedge y_1 \triangleright y_2 \wedge \dots \wedge y_l \triangleright a$. In other words, if there exists a path from a to b in $\triangleright$ but not vice-versa. The members of the Schwartz set are all candidates who are not Schwartz dominated. The winners are all the candidates in the Schwartz set, $f_S(P) = \{a \in A \mid \nexists b \in A \setminus \{a\} \text{ s.t. } b \vdash a\}$.

Theorem 2.1 (Fishburn [1977]). $f_S(P)$ is not *Paretian*.

Proof. Consider the following profile P :


$$\begin{array}{ccc} 1 & 1 & 1 \\ a & d & c \\ b & a & d \\ c & b & a \\ d & c & b \end{array},$$

Consider, for each pair of alternatives, the number of profiles that prefer one alternative to the other.

$$\begin{array}{lll} \mu(a, b) = 3, & \mu(a, c) = 2 & \mu(a, d) = 1 \\ \mu(b, a) = 0, & \mu(b, c) = 2 & \mu(b, d) = 1 \\ \mu(c, a) = 1, & \mu(c, b) = 1 & \mu(c, d) = 2 \\ \mu(d, a) = 2, & \mu(d, b) = 2 & \mu(d, c) = 1 \end{array}.$$

From this we deduce that a Pareto dominates b and, also, that the simple majority relations are $a \triangleright b$, $a \triangleright c$, $b \triangleright c$, $c \triangleright d$, $d \triangleright a$, $d \triangleright b$.

For an alternative a to be in the Schwartz set, there must not be an alternative b for which there is a path from b to a in $\triangleright$ but no path from a to b . If we represent the preferences with a graph in which each node is an alternative and the arcs correspond to the $\triangleright$ relations, it is easy to see that from whichever node we start we can always reach any other node.



$$\begin{array}{ll} a \triangleright b & \wedge \quad b \triangleright c \wedge c \triangleright d \wedge d \triangleright a \\ a \triangleright c & \wedge \quad c \triangleright d \wedge d \triangleright a \\ a \triangleright c \wedge c \triangleright d & \wedge \quad d \triangleright a \\ b \triangleright c & \wedge \quad c \triangleright d \wedge d \triangleright b \\ b \triangleright c \wedge c \triangleright d & \wedge \quad d \triangleright b \\ c \triangleright d & \wedge \quad d \triangleright a \wedge a \triangleright c \end{array}$$

This means that the Schwartz set is equal to the set of all alternatives, thus that $f_S(P) = \{a, b, c, d\}$. But b is a Pareto dominated alternative, thus the Schwartz rule does not satisfy the Pareto criterion. $\square$

This example shows us that even "basic" properties like Pareto are worth verifying, because rules that seem reasonable may not satisfy them. And, to be fair, deviating a little, this may lead us to another consideration that will be useful later on: is Pareto really that basic and essential? Sen [1970, 2004] thinks that this is not always the case. In fact, he introduced the concept of *Minimal Liberty*, according to which there should be at least a pair of alternatives a, b such that if an individual prefers a to b then society prefers $a \succ b$. Considering all preference orders to

be possible (*Unrestricted Domain*), he showed that there is no way to aggregate preferences into a single common choice that simultaneously satisfies the Pareto criterion and Minimal Liberty. Sen suggests that one way to escape this paradox is to give up Paretian efficiency: “While the Pareto criterion has been thought to be an expression of individual liberty, it appears that in choices involving more than two alternatives it can have consequences that are, in fact, deeply illiberal.” [Sen, 1970] In our quest to identify rules that could reflect a true compromise between voters (Chapter 7), we will come across a class of rules that we find reasonable in some contexts but which do not satisfy the Pareto criterion.

BK compromises

In order to describe the next set of rules it is necessary to make a premise. We have decided to group under the umbrella name of Brams and Kilgour (BK) compromises, all those rules which, starting from everyone’s ideal alternative, fall back to the voters’ second, third and more generally k -th best, until a certain quota of voters q support some of the alternatives considered so far.

This method, in its main idea, has been rediscovered several times over the years. Its proposal can be traced back to de Condorcet [1789, 1793], whose procedure, translated by McLean and Hewitt [1994, pp. 249-250], reads: “If one candidate has the absolute majority of first votes, he will be elected. If one candidate has the absolute majority of first votes and second votes together, he will be elected. If several candidates obtain this majority, the one with the most votes will be preferred. If one candidate has the absolute majority of the three votes together, he will be elected, and if several candidates obtain this majority, the one with the most votes will be preferred”. Here the quota q of supporters an alternative must have to get elected is $\lfloor \frac{n}{2} \rfloor + 1$, i.e. the majority of voters. It is important to note, however, that Condorcet specifies a tie-breaking procedure to apply when several alternatives are supported by at least q voters. The rules belonging to the BK Compromise family do not use a tie-breaking procedure, in such instances they return the set of alternatives respecting the quota.

As Camps et al. [2014] point out, the same method was later analyzed by Lhuillier [1793] and adopted in Geneva. Sometime after that, in the early 20th century, this procedure was promoted by the politician James W. Bucklin and adopted first in the municipality of Grand Junction (Colorado, USA) and then in many other cities of the United States [Barber, 2000, p. 167]. In 1986 it was proposed again by Sertel [1986] and Sertel and Yılmaz [1999] under the name of Majoritarian Compromise. Later, Brams and Kilgour [2001] generalized this concept and introduced a class of SCR called *q-approval fallback bargaining*, where q is the required quantity of support that can vary from a single voter up to unanimity. Different choices of q lead to different SCRs. Considering n voters, the choice of $q = 1$ corresponds to the plurality rule and $q = n$ represents a bargaining procedure called *Fallback bargaining*, which has been further analyzed by Kibris and Sertel [2007] and Congar and Merlin [2012]. Again, Brams and Kilgour do not use a tie-breaking mechanism, thus the BK rule with $q = \lceil \frac{n}{2} \rceil$ is similar but not equal to Majoritarian Compromise, which, instead, is a subset of the former.

To describe these BK compromises, recall that we are considering a finite set N of n individuals and a finite set A of m alternatives. A *profile* $P : N \rightarrow \mathcal{L}(A)$ associates with each individual $i \in N$ a preference $P(i) = \succ_i$. We write $r_{\succ_i}(a) = \#\{b \in A \mid b \succ_i a\} + 1$ for the

rank of $a \in A$ at $\succ_i \in \mathcal{L}(A)$.

Given any $1 \leq q \leq n$ and $1 \leq d \leq m$, consider the (q, d) -compromise set, i.e. the set of alternatives that are ranked among the top d by at least q voters $CS_d^q(P) = \{a \in A \mid \#\{r_{\succ_i}(a) \leq d, \forall i \in N\} \geq q\}$. The q -approval fallback bargaining depth, d_q^* , is the smallest value that makes this set nonempty: $d_q^* = \min\{d \mid CS_d^q(P) \neq \emptyset\}$.

The **SCR q -approval fallback bargaining** is defined as the set of all alternatives that are ranked d_q^* for at least q bargainers: $f_{BK}^q = CS_{d_q^*}^q(P)$ where $d = d_q^*$.

Consider the maximum support at depth d as the biggest value that makes the compromise set nonempty: $q_d^* = \max\{q \mid CS_d^q(P) \neq \emptyset\}$. We can define a tie breaking version of the BK rule, in a Condorcet fashion, by adding an extra step. If multiple alternatives are ranked d_q^* for at least q bargainers, the **tie breaking q -approval fallback bargaining** returns the ones with the greatest support: $f_{tbBK}^{q_d^*} = CS_{d_q^*}^{q_d^*}(P)$ where $d = d_q^*$. Note that when CS is a singleton, when $q = n$ and when all the alternatives in CS are supported by the same amount of voters then $f_{BK}^q = f_{tbBK}^q$. Example 2.5 shows how different quotas lead to different results.

Example 2.5 ([Brams and Kilgour, 2001]). Consider the following profile P :

1	1	1	1
a	a	a	d
b	c	d	b
c	b	b	c
d	d	c	a

when $q = 1$ then $d_1^* = 1$ and $f_{BK}^1 = \{a, d\}$ and $f_{tbBK}^1 = \{a\}$;

when $q = 2$ then $d_2^* = 1$ and $f_{BK}^2 = \{a\}$ and $f_{tbBK}^2 = \{a\}$;

when $q = 3$ then $d_3^* = 1$ and $f_{BK}^3 = \{a\}$ and $f_{tbBK}^3 = \{a\}$;

when $q = 4$ then $d_4^* = 3$ and $f_{BK}^4 = \{b\}$ and $f_{tbBK}^4 = \{b\}$.

When $q = n$ the BK rule is referred to as **Fallback Bargaining** (FB): $f_{BK}^n = f_{FB}$.

When $q = \lceil \frac{n}{2} \rceil$ the revised BK rule is denoted as **Majoritarian Compromise** (MC): $f_{tbBK}^{\lceil \frac{n}{2} \rceil} = f_{MC}$. Example 2.6 shows a profile where the two sets of outcomes f_{MC} and $f_{BK}^{\lceil \frac{n}{2} \rceil}$ do not coincide.

Example 2.6 (Sertel and Yilmaz [1999] and Brams and Kilgour [2001]). Consider the following profile P :

1	1	1
a	c	b
c	a	c
b	b	a

and $q = \lceil \frac{n}{2} \rceil = 2$. Therefore, $d_2^* = 2$ and the two SCRs return $f_{BK}^2 = \{a, c\}$ and $f_{MC} = \{c\}$.

To conclude this section we show in Table 2.1 an overview of which of the axioms described is satisfied by the rules discussed so far. Note that we say that an axiom is satisfied by the class BK Compromises, which we informally use to design the class of *q-approval fallback bargaining*, if it is satisfied by every element in the class, and the converse if there is at least one element in the class for which it is not satisfied.

	Anon.	Neut.	Nondict.	Pareto	Condorcet	Majority	Manipul.	SP
Plurality	✓	✓	✓	✓	×	✓	✓	×
Anti-Plurality	✓	✓	✓	×	×	×	✓	×
Borda	✓	✓	✓	✓	×	×	✓	×
Copeland	✓	✓	✓	✓	✓	✓	✓	×
Schwartz	✓	✓	✓	×	✓	✓	✓	×
BK Compromises	✓	✓	✓	✓	×	✓	✓	×

Table 2.1: Axioms satisfied by the voting rules described.

2.2 Voting with rated ballots

In the classical Arrovian framework for social choice a set of voters is assumed to give a complete preferences order over a set of alternatives [Arrow, 1950]. When we think about it, this is not so obvious to achieve. For example, there may be too many alternatives and the cost for the voters to provide a complete order is simply too high. In this situation, one could opt for an approval system, in which voters approve only the desired alternatives [Brams and Fishburn, 2007]. Yet, this does not resolve other criticisms raised against the system of ranked ballots: that of preference intensity. One may argue that preference rankings provide already some level of intensity, but they do not capture the full story. Different voters may have the same order but prefer one alternative to another with different intensity levels. Suppose that a slight majority of voters slightly prefer alternative a to alternative b , these would appear in the preference rankings as $a \succ b$. Now consider that the remaining slight minority despises a , and strongly prefers b to a . There is no way for any rule in the framework of ranked preferences to distinguish the difference in intensity between $a \succ b$ in the first group and $b \succ a$ in the second. How to decide what is the most desirable alternative for society?

A possibility is to evaluate candidates according to a common language shared by society. Note that from the evaluation of candidates we can deduce a ranking, but the converse is not true. Thus, an evaluation is more informative than a ranking. This describes the framework introduced by Smith [2000], and later analyzed by Pivato [2014], Gaertner and Xu [2012], Zahid and de Swart [2015], when proposing his *Range Voting* rule. Here the common language consists of a range of numerical scores, for example from 0 to 100, and the rule works as follows: each voter assigns a score to all or some alternatives, the one with the highest average score is elected. If more than one alternative has the same highest score then the tie is broken randomly.

But a whole bunch of new questions arise. What is the right size for this scale? If it is too large it may create manipulation problems because voters may exaggerate their preferences; if it is not large enough it reduces expressiveness. Is it better to use a numerical scale or a more immediate language? Are there really benefits to using ratings instead of rankings? Luckily for

us, there is a vast literature of economists and psychologists who have addressed these questions, even sometimes arriving at very different answers. Unfortunately, to discuss this would go beyond the scope of this thesis, but for a (non-exhaustive) list of work on the subject one can read Cox [1980], Sparling and Sen [2011], Churchill and Peter [1984], Preston and Colman [2000], Maio et al. [1996].

In this manuscript we are going to focus on a specific scale of seven expressions: *Excellent*, *Very Good*, *Good*, *Passable*, *Inadequate*, *Mediocre*, *Bad*. This scale was used by Balinski and Laraki [2007] when proposing their Majority Judgment (MJ) rule. Each voter judges all the candidates and the one with the highest median is elected. We describe this rule in more detail in the next section. For now, we would like to conclude our introduction with a small anecdote. During a 2012 interview Arrow was asked how he thought voting systems should be evaluated in the future, he replied: “I’m a little inclined to think that score systems where you categorize in maybe three or four classes probably (in spite of what I said about manipulation) is probably the best” [Hamlin, 2012].

Majority judgment

MJ is a voting method proposed by Balinski and Laraki [2007, 2011] to elect one out of m candidates based on the judgments of n voters. The latter express their preferences by assigning to each candidate one of the following adjectives: Excellent, Very good, Good, Average, Mediocre, Inadequate, To be rejected. Those adjectives represent a common language whose semantic is assumed to be a shared knowledge among the voters carrying thus an absolute meaning. For each candidate the median of the grades she received is computed, this is called *majority-grade*. The candidate with the highest majority-grade is elected. Ties are broken by considering the majority-grade of first order: one vote associated with the majority-grade of each tied candidates is removed and their medians are recomputed. The candidate with the highest new median is elected. If there is still a tie the process is repeated until a unique winner is found.

In order to formalize this description, consider a finite set N of voters (or judges) with $\#N = n$ and a finite set A of alternatives (or competitors) with $\#A = m$. A *common language* $\Delta = \{\delta_1, \delta_2, \dots, \delta_7\}$ is a set of strictly ordered grades. The notation $\delta_1 \geq \delta_2$ indicates that δ_1 is a better or equivalent grade than δ_2 . Here δ_1 corresponds to Excellent, δ_2 to Very good etc. A profile $P : A \times N \rightarrow \Delta$ is a m by n matrix of grades.

The operator $\rho : \Delta^N \rightarrow \Delta^{|N|}$ defines an ordering function that given a vector of grades P_i returns the vector ordered by decreasing grades.

The *majority-grade*, f_{maj} , is the function that associates to a vector of grades $v \in \Delta^N$ its median grade value: $f_{\text{maj}}(v) = \rho(v)_{\lfloor \frac{|v|}{2} \rfloor + 1}$. Note that in case $|v|$ is even two medians could be used, but in Balinski and Laraki [2011] definition the lower grade is picked.

The Majority Judgment function $F_{\text{maj}} : \Delta^{A \times N} \rightarrow A$ is a function that selects the alternative with the highest median grade as winner. We can define it as $F_{\text{maj}}(P) = \arg\max_{i \in A} f_{\text{maj}}(P_i)$ assuming it is a singleton. To avoid adding further complexity to the notation, we describe only informally what happens when this is not the case, in other words when two or more alternatives are associated with the same highest median grade $h = \max_{i \in A} f_{\text{maj}}(P_i)$. In this case, ties are broken by removing one h grade from the vectors of grades of each tied alternative, recomputing

the new median grade and repeating the process until one unique winner is found or there are no more grades to remove. When n is odd, this is equivalent to take the next element after the median, i.e. the one at index $(\lfloor \frac{n}{2} \rfloor + 1) + 1$, if there is still a tie we then look at the previous element before the median, i.e. the one at index $(\lfloor \frac{n}{2} \rfloor + 1) - 1$, and keep alternating until the tie is broken or there are no more elements in the vector. When n is even the process is similar but we alternate starting from the element before the median, $\lfloor \frac{n}{2} \rfloor$, then the one after, $\lfloor \frac{n}{2} \rfloor + 2$, and so on. If after applying the mechanism there are still ties we break them using an arbitrary ordering defined on all alternatives, e.g. lexicographical order.

Criticisms of this rule include that the median does not respect the majority principle, meaning that if a majority of voters prefer an alternative a to an alternative b , the latter could still win with an higher median. Consider the following example.

Example 2.7 ([Laslier, 2018]). *Consider the following profile formed by three voters i_1, i_2, i_3 and two candidates a, b . Suppose the voters assign to candidates a a grade between 0 and 20.*

	i_1	i_2	i_3	
a	20	9	9	.
b	11	0	10	

The median of a is 9 and the median of b is 10, thus b is elected although only one voter slightly prefers b to a . This example can be reproduced with a larger number of voters:

	50	50	1	
a	20	9	9	.
b	11	0	10	

The alternative b is elected even if almost the unanimity of the voters strongly prefers a to b and only one voter slightly prefers b to a .

Is it worth mentioning that the same example is also studied by Balinski and Laraki [2011, p. 281], where the authors reproduce the example with MJ grades.

Example 2.8 ([Balinski and Laraki, 2011]). *For the effect to be reproduced in the MJ setting the example should be translated as:*

	50	50	1	
a	Excellent	Mediocre	Mediocre	.
b	Good	To be rejected	Average	

However, the authors stress the fact that is very unlikely that the last voter would associate a different grade to 9 and 10. In a large electorate the distinction would be too fine to make a significant difference between 20 and 19 or 10 and 9, so they affirm that the example would more realistically translate to:

	50	50	1	
a	Excellent	Mediocre	Mediocre	.
b	Average	To be rejected	Mediocre	

And a would be elected.

Another criticism regards manipulability. MJ is manipulable to the extent that voters can exaggerate by excess the judgments of their top choices and by defect those of their worst choices. It must be mentioned that Bassett and Persky [1999] proved that the high breakdown property of the median—which is the sensitivity to outlying observations —makes difficult for a minority to manipulate the ranking. However, Gehrlein and Lepelley [2003] studied that the probability of being subject to manipulation is just slightly smaller than other methods like Borda and Copeland.

HISTORY AND LITERATURE REVIEW

In this chapter we will analyze in detail the topics covered in this manuscript. We will describe the different definitions and approaches to the issues considered that have been proposed in the literature. In the first part of this chapter, Section 3.1, we will focus on the meaning of compromise. All voting rules can be considered a compromise between voters preferences, but some more than others. We will consider several notions of compromise and analyze different aggregation methods that satisfy them. In the second part, Section 3.2, we will focus on a specific rule that has been considered, by some authors, a compromising rule: Majority Judgment. We will describe its introduction, its uses in real life voting procedures and also its criticisms. In the final part, Section 3.3, highlighting how some assumptions made so far, notably the complete knowledge of preferences, are too stringent, we will analyze the literature on preference elicitation. We will consider different fields and different approaches within the same field. In all sections we will present the definitions introduced in our contributions and position them in the existing literature. We will show how our approach differ from existing works and lay the foundation for the content of the next chapters.

3.1 Voting as a compromise

After having described voting rules as procedures designed to aggregate a set of preferences into a common choice, we could say that each one of these mechanism reflects some sort of compromise. In fact, reaching a collective choice by unanimous accord is a fairly rare situation, especially in elections with a large number of voters. In all other cases, someone has to give up their top preference in favor of a second, or third etc., choice. Yet, can we call this a compromise? What do we mean when we talk about this notion?

How do people compromise? One way of approaching the analysis of a concept is to start from the etymology of the word itself. The word compromise comes from the Latin *compromissus*, past participle of *compromittere*: *com together* and *promittere to promise*. The idea of compromise can indeed be found in texts dating back to Roman times, where two disputing parties who wished to submit their contention to arbitration, would appoint a third party (an arbitrator) and make the mutual promise to abide by the latter's unquestionable judgment. If either party broke this promise, it would have had to pay a penalty. In a fragment of a letter written by Proculus, which is reported by Zimmermann [1996, p.529], it is clear that the two disputants agreed to give the arbitrator unlimited powers in his decision. No appeal was possible against the final award, which was binding no matter how unjust and unequal it was. This idea of compromise as a mere arbitration process seems very different from the notion we have of

it today. Let us consider the definitions of the two verbs in a modern dictionary: to arbitrate "to settle a dispute between two people or groups after hearing the arguments and opinions of both" [Merriam-Webster, 2022a]; to compromise "to come to agreement by mutual concession" [Merriam-Webster, 2022b]. The mutual concession is the element that immediately makes us think of a compromise and it is precisely the missing factor in the description of the Roman compromise and modern arbitration. Quoting Braybrooke [1982] "It is simply a bad joke to speak of someone's being a party to a compromise when she has got nothing out of it."

However, Roman society is not the first one to use mechanisms of compromise; we find references to use of arbitration and selection of arbitrators in Aristotle's Constitution of the Athenians [53.2–4]. Moreover, Aristotle is the first to associate the concept of fairness with that of arbitration:

*"Fairness, for example, seems to be just; but fairness is justice that goes beyond the written law."*Aristotle trd. Kennedy [2006][1.13.13] *"And it is fair to want to go into arbitration rather than to court; for the arbitrator sees what is fair, but the jury looks to the law, and for this reason arbitrators have been invented, that fairness may prevail."*Aristotle trd. Kennedy [2006][1.13.19]

Throughout history, however, compromise has been intended in its Roman meaning: the resolution of a dispute by a third party whose judgment individuals agree to abide by. The use of arbitration as a mean of compromise continued until the end of the Middle Ages, only to be overshadowed by the advent of absolute monarchies, like the one of Louis XIV, that tended to centralize all powers in a single strong figure[Zappalá, 2018]. However, this did not last long. In France, for example, the first constitution approved after the revolution, in 1791, defended the right of arbitration:

"Chapitre V. Article 5. - Le droit des citoyens, de terminer définitivement leurs contestations par la voie de l'arbitrage, ne peut recevoir aucune atteinte par les actes du Pouvoir législatif."[Constitution de 1791, 1791]

It would seem at this point that we are straying a little from the topics treated in this thesis. However, this excursus on how the notion of compromise has been understood for centuries can give us insights on how compromise has been approached in the history of social choice theory. In fact, as Examples 3.2 and 3.3 in Section 2.1.2 show, the best-known voting rules grouped under the name of compromise rules start with the intention of seeking a compromise among voters, however they do not ensure that the outcome of the procedure represents a mutual concession.

Consider an example used by Luban [1985]. Rich and Poor are offered 1000 dollars which they can obtain provided they reach a compromise on how to divide this sum. Rich who knows to have leverage over Poor offers him 100 dollars threatening to walk from the deal if he does not accept. Poor, who is in desperate need of money, accepts. Luban points out that this is considered a compromise because 1) it produces an agreement 2) the criterion of "equal satisfaction" is met if we imagine that 100 dollars give Poor the same level of satisfaction that 900 dollars give Rich. Most of us, however, would consider this outcome as a ridiculous extortion rather than a compromise. What is then the best compromise possible between Rich and Poor? There is no

clear answer. If we were to value equity as a fundamental concept of wealth redistribution, then an equitable solution would be for Rich to get 100 dollars and Poor 900 dollars. However, if equality is an important social value, we would intuitively think that the only acceptable result is a 500/500 split. That is, the result where, in the interpretation of compromise as "mutual concession", both parties concede equally. This consideration is the basis for our version of compromise rules presented in Chapter 7.

On equality, fairness and the principle of equal losses. It is important to point out, however, that the concepts of equality and fairness have very specific meanings depending on the considered literature. In this manuscript we consider equality with respect to losses from one's ideal point. We denote a compromise as a situation in which all parties are not only willing to concede, but that at the end of the process they all give up as equally as possible. One of the first author to refer to the equal-loss principle was Mill [1849, p. 396] who, when discussing fairness of taxation, wrote: *"As a government ought to make no distinction of persons or classes in the strength of their claims on it, whatever sacrifices it requires from them should be made to bear as nearly as possible with the same pressure upon all, which, it must be observed, is the mode by which least sacrifice is occasioned on the whole. [...] Equality of taxation, therefore, as a maxim of politics, means equality of sacrifice."* Since then, the equal-loss principle has received a lot of attention, especially in the problems of taxation [Edgeworth, 1897, Young, 1987] and bankruptcy [Herrero and Villar, 2001, Aumann and Maschler, 1985], and there has been much discussion about the interpretation of sacrifice, which is ultimately a very subjective concept. Chun [1988] was the first to apply the equal loss principle to bargaining theory. This new bargaining solution, called the equal-loss solution, aims at equalizing across agents the losses from their ideal point.

Following Nash [1950] description, a bargaining problem is a pair (S, d) where S is the set of feasible agreements and d is the disagreement point or status quo, i.e. what each agent gets in case an agreement is not reached. A solution, defined on a class of problems, is a function that maps each problem (S, d) to a point in S representing the agreement reached. Nash proposed to analyze solutions based on the properties, axioms, they satisfy. His solution consists of each agent getting at least what they would get if they disagreed, plus a share of the benefit the group would get if they cooperated. The outcome of this procedure satisfies certain desirable properties including Pareto optimality. It is important to mention that the Nash model is based on a two individuals bargain, but an extension to n individuals can be formulated. Nash's equilibrium, however, has been criticized for several reasons. Sen [2017, pp. 177-180] points out that his solution relies too much on the bargaining power of individuals. Even assuming that the outcome of a bargaining procedure is better than what each agent would have gotten in case of disagreement, it does not make such outcome fair or desirable. To support this thesis, Sen makes an example where some unemployed workers may accept extremely low wages and unfair employment treatments just because in the absence of agreement they may starve. He argues that although this outcome is a Nash equilibrium, the workers were exploited due to their poor bargaining power. Other solutions have been proposed, for example Kalai and Smorodinsky [1975] focused on equalizing the gains of each player relative to their maximum possible gain, and Kalai [1977] introduced the *egalitarian bargaining solution* which equalizes the gains of

the agents from the disagreement points. The latter solution is similar in its idea to the concept of *compromise* advanced by Yu [1973] which, instead, equalizes the losses of the agents from the ideal point. Both of those approaches are in accordance with Rawls' idea of fairness [Rawls, 1958, 1967], whose approach to inequalities is to increase the welfare level of the worst-off individual as much as possible. In other words, a maximin procedure.

Equal-loss vs Pareto. To return to the equal-loss solution, Chun [1988] himself mentions the analogies with Kalay's solution and he points out that it is a variant of Yu's solution, but he offers a characterization of it. Given a problem (S, d) , the ideal point, $a(S, d)$, corresponds to the maximal utility level each agent could obtain provided that all agents obtain at least what they would get in case of disagreement. That is, for each agent i , $a_i(S, d) = \max\{x_i | x \in S, x \geq d\}$. The equal-loss solution is defined as the maximal point $x \in S$ with $a_i(S, d) - x_i = a_j(S, d) - x_j$ for all agents i, j . As we already mention, the idea is to equalize across agents the losses from the ideal point. Although this sounds very appealing, one of the main criticisms of Chun's solution is that it does not satisfy the Pareto criterion.

Example 3.1 (Adapted from Sen [2017] pp.194-195). *Consider two individuals i, j and three states a, b and c with the following welfare levels:*

	i	j	
a	20	1	,
b	10	1	
c	0	0	

where c is the disagreement point. Despite the state a being Pareto optimal, Chun's procedure would consider the states a and b indifferent. In fact, the worst off individual j is not better off under a than under b (nor the converse).

Sen uses this example to criticize the failure of maximin rules to satisfy the Pareto criterion. We added the state c of disagreement in order for it to apply in the context studied by Chun.

To "solve this problem", a few years later, Chun and Peters [1991] defined a lexicographic version of the equal-loss solution. First, a lexicographic order is defined on the domain of the feasible outcomes, then the equal loss solution is applied: the maximal outcome in which the individuals suffer equal losses from the ideal point is chosen. If this outcome is not Pareto optimal then the next one in lexicographic order is taken until a Pareto optimal outcome is found. Other characterizations of lexicographic maximin solutions have been proposed by Sen [2017] and Chang and Hwang [1999].

We have put "solve this problem" in quotes because we do not believe the failure to satisfy the Pareto criterion to be a problem. Consider again Example 3.1, agent j is indifferent to the two states a and b as she would get the same utility in both cases, but agent i clearly prefers a to b . The state a is therefore the Pareto optimal alternative. Nevertheless, let us imagine that a and b represent different distributions of money within a society. If we wanted to pursue an ideal of social justice and reduce inequalities immediately the alternative b seems the most appealing. Sen also makes a criticism of the Pareto criterion [Sen, 2017, Chapter 6]. While he acknowledges that it is not always satisfied in literature, he notes that the Pareto criterion

is always treated as a missed opportunity. Sen, however, observes that it is not a criterion that is so obviously desirable. In fact, he shows that even a weak Pareto condition conflicts with a weak condition of individual freedom [Sen, 2017, Theorem 6*3]. Example 3.1 can represent a variation of an illustration given by Sen himself to propose the rejection of the Pareto principle. Imagine that agents i and j have to decide who reads a given somewhat scandalous book. Agent i is what is considered a prude, she would not want to read that book, but she would rather be the one to read it than have the impressionable agent j read it. For agent j it is indifferent because she would like to read it, but she would be just as happy with agent i reading it. Consider a to be the state in which the book is read by i , b to be the state in which it is read by j and c to be the state in which none reads the book. We can use the utilities of the Example 3.1 to represent this situation. Again, a is the Pareto optimal alternative but only because i has the presumption of wanting to decide for others. Sen argues that it is not only important to know the preferences of individuals but also why individuals have that preference: “Preferences based on excessive nosiness about what is good for others, should be, it could be argued, ignored.” [Sen, 2017, Chapter 6.5]

In Chapter 7 we will introduce two different versions of our vision of compromise, in particular one that prioritizes equality at the expense of Pareto efficiency. Moreover, we consider the concept of equal-loss as the basis of any successful compromise in all situations where egalitarianism, in the sense of conceding equally, is a major concern. This is a novelty in the literature of social choice rules, which so far have only imposed willingness to compromise without actually ensuring that all parties concede something. The philosopher Day [1989] attempts an explanation for this phenomenon, thinking that it stems from the negation of the adjective “uncompromising”: “An uncompromising person is one who is not disposed to make any concessions, so (it is erroneously inferred) a compromising person is one who is pliant and disposed to make concessions—regardless of whether he receives any concession in return. However this may be, one must insist, as it is generally agreed, that compromise necessarily involves mutual concessions.”. This might explain why all voting rules that try to compromise actually settle for the will of the agents to compromise.

Compromising rules. Merlin et al. [2019] discuss the most famous of these procedures proposing to gather them in the same class of *compromise rules*. In Section 2.1.2 we already defined the Majoritarian Compromise, introduced by Sertel [1986] and further analyzed by Sertel and Yılmaz [1999]. Based on a revision of the Condorcet-Bucklin rule, this procedure starts from everyone’s ideal choice to find an alternative supported by the majority of voters. If no such alternative exists, then it falls back to the voters’ second, third and more generally k -th best choice, until at least one of the alternatives considered appears among the first k best for a majority. If instead of considering an agreement for the majority of voters we wish to select as winners the alternatives supported by the unanimity of voters, then the procedure corresponds to the Fallback Bargaining rule. More generally, if we consider the support of a certain quota q of voters we refer to the q -approval fallback bargaining rule [Brams and Kilgour, 2001].

All these SCRs impose to voters a willingness to compromise, but we argue that they do not effectively ensure an outcome where the agents have, indeed, compromised. Example 3.2 motivates our view.

Example 3.2. Consider the following preference profile with $n = 100$:

49	51	
c	a	
b	b	$\cdot$
a	c	

When $q \in \llbracket 1, \frac{n}{2} + 1 \rrbracket$, all BK compromises pick a and c while the tie breaking BK compromise selects a . These outcomes do not appear as a compromise as almost half of the voters obtain their best choice while the remaining half have to be contented with their worst one. Observe that b receives unanimous support when each voter falls back one step from her ideal point.

Although we will not study it in this context, it is interesting to mention that Merlin et al. [2019] also include the Majority Judgment in the class of compromise rules. Introduced by Balinski and Laraki [2007, 2011], we defined this procedure in Section 2.2 and we discuss it in more detail in Section 3.2. In particular, the authors consider MJ to be a compromise by transposing the possible grades into ranks of the preference orders, with the condition that each voters use the same grade only once. In this context the alternative with the median rank is the compromise. If more of such alternatives exist, the tie breaking mechanism proposed by Balinski and Laraki [2011] is used: one of the median rank is removed to each candidate winner until one of those will get the new highest median.

Other mechanisms studied in different contexts may still fit into the classification of compromise rules. For example, if we consider only two voters, de Clippel et al. [2014] analyze two methods for the selection of arbitrators: the *Veto-Rank mechanism (VR)*—which is the most common procedure for assigning arbitrators [de Clippel et al., 2014, Online Appendix]—and *Shortlisting (SL)*. Consider a list of m (odd) alternatives (that are candidates to be arbitrators), and two voters (that are the two parties that must agree on an arbitrator). In the VR procedure, both voters simultaneously veto their worst $\frac{m-1}{2}$ alternatives. The selected alternatives are the ones with the highest Borda score among the non-vetoed alternatives. In the SL procedure, one of the two parties starts by vetoing her worst $\frac{m-1}{2}$ alternatives, and then the second party chooses her best alternative out of the remaining ones. As the outcome of the procedure depends on the party that starts, symmetry among players is ensured by defining the solution as the union of the two outcomes where one and the other party starts. Another class of two voters rules based on the veto power is the one denoted as *Pareto-and-Veto rules (PV)* by Laslier et al. [2020]. The same class was also studied by Moulin [1983], who refers to it with the name of *veto-core*, and by Abreu and Sen [1991]. This procedure distribute a veto power of v_1 and v_2 alternatives to voters 1 and 2, respectively, with $v_1 + v_2 = m - 1$. Both voters, $i = 1, 2$, simultaneously veto their worst v_i alternatives. The selected alternatives are all the non-vetoed and Pareto optimal ones. Similar to the case $n \geq 3$, Example 3.3 motivates our argument in the case of two voters.

Example 3.3. Consider the following preference profile with $n = 2$:

i_1	i_2
a	y_1
x_1	y_2
x_2	b
x_3	a
b	x_1
y_1	x_2
y_2	x_3

All the rules aforementioned, except for PV when $v_1 < v_2$, selects the alternative a as winner:

FB: a is the first alternative to reach unanimous consent;

VR: both voters veto their $\frac{7-1}{2} = 3$ worst alternatives so i_1 vetoes $\{b, y_1, y_2\}$ and i_2 vetoes $\{x_1, x_2, x_3\}$, the only alternative left is a ;

SL: when i_1 starts she vetoes her $\frac{7-1}{2} = 3$ worst alternatives thus $\{b, y_1, y_2\}$, then i_2 picks her best alternative among the remaining ones thus a . When i_1 starts she vetoes $\{x_1, x_2, x_3\}$ and i_2 picks a . The outcome of the SL procedure is then a ;

PV: when $v_1 \geq v_2$:

- $v_1 = 5, v_2 = 1$, i_1 vetoes $\{x_2, x_3, b, y_1, y_2\}$ and i_2 vetoes $\{x_3\}$ PV selects a ;
- $v_1 = 4, v_2 = 2$, i_1 vetoes $\{x_3, b, y_1, y_2\}$ and i_2 vetoes $\{x_3\}$ PV selects a ;
- $v_1 = 3, v_2 = 3$, i_1 vetoes $\{b, y_1, y_2\}$ and i_2 vetoes $\{x_2, x_3\}$ PV selects a .

It is worth mentioning that PV does not select a as the winner in this profile when the veto powers are $v_1 < v_2$. However, this profile is extended as a desirable property of our notion of compromise that we will describe in more detail in Chapter 7. Generalizing this example we will show that PV, whatever the distribution of veto power is, fails to select what would be for us the compromise outcome: b .

3.2 Median compromise: the Majority Judgment rule

As mentioned in Section 3.1, an example of compromise can be represented by Majority Judgment, a voting system where voters assign grades to candidates using an ordinal scale. The evaluation of candidates instead of their ranking yields more information as a result of a higher expressiveness. In MJ each voter evaluates, or judges, each candidate and the winner is the candidate with the highest median of the grades received. This method was introduced by Balinski and Laraki [2007] in the most recent period of social choice history. Yet, it has attracted increasing attention of french associations and political parties which have started to use MJ for internal decisions or local elections.

An overview of the use of the median as aggregator. The idea of using the median in voting is not new, one of the first use can be traced back to Sir Francis Galton's *middlemost* method [Galton, 1907a,b]. He considered a situation where a group of people has to assess a damage. Galton argued that the aggregation metric that poses the least problems in this context is the median because it is not affected by over- or under-exaggerations. In his words "According to the democratic principle of one vote one value, the middlemost estimate expresses the *vox populi*, every other estimate being condemned as too low or too high by a majority of the voters". In other words, the median for Galton represents the most accurate estimate of *people's voice*. He also precised that in the occurrence of even votes two medians are possible, in this case the average of those is taken.

Numerous observers described the median grade as the highest level at which a candidate obtains the support of the majority of the voters. In other words, starting for the highest grade d we check if the majority of the voters assigned at least d to some alternative a . If this is not the case, we descend in the grading scale until such level d^* is found where a candidate a^* satisfies half population. The grade d^* is then the median of a^* grades, and, since it is the first level we stopped at, it corresponds to the best possible median. As we know from Section 2.1.2, this method was rediscovered several times and proposed under the name of Bucklin's rule [Hoag and Hallett, 1926], Majoritarian Compromise [Sertel, 1986, Sertel and Yılmaz, 1999] and q-approval fallback bargaining [Brams and Kilgour, 2001]. Moreover, note that when the number of grades is equal to two (approve, disapprove) then this method is reduced to Approval Voting.

More recently Bassett and Persky [1999] proposed the use of the median as voting rule for elections advocating for its robustness. In fact, they argue that the high breakdown property of the median minimizes the risk of outcome manipulation by a minority of voters. The authors also insist on the fact that the median is the only method with a high breakdown that also satisfies the monotonicity criteria [Bassett and Persky, 1994]. However, Gehrlein and Lepelley [2003], considering only three candidates elections, proved that plurality rule has only a slightly greater probability of manipulability than the median voting rule proposed by Bassett and Persky [1999] while having a significant greater probability of producing a decisive result. They also study Borda rule and Copeland rule finding similar results: Borda has a higher probability of manipulation but also of decisiveness, Copeland has even lower probability of manipulation and higher probability of decisiveness.

Other studies on the use of the median have been conducted. In particular, Barthélemy and

Monjardet [1981] survey mathematical problems and properties related to the notion of median in the context of cluster analysis and social choice theory. Nehring and Pivato [2022] analyze the median rule in judgment aggregation and they define a weighted median rule that is equivalent to MJ except for the treatment of ties.

Real-life use cases of the MJ procedure. In France, MJ has being adopted by a progressively larger number of associations and political parties including: Le Parti Pirate, Génération(s), LaPrimaire.org, France Insoumise and La République en Marche. "Mieux Voter" [MieuxVoter, 2022a] is a french association that promotes the use of MJ as voting method whenever a collective choice has to be selected: public administration, associations, companies. On their website it is possible to find all the citizens lists –party lists that are not affiliated to any national political party —that used MJ to rank their candidates during the local elections of 2020. In two cases, Bordeaux et Annecy, the candidate selected using MJ was then elected as a mayor.

LaPrimaire.org [LaPrimaire.org, 2016] is a french political initiative whose goal is to select an independent candidate for the french presidential election using MJ as voting rule. All french citizens over 18 with rights to vote can participate as candidates or voters. The association Democratech implemented the platform for the first time in 2016 in view of the 2017 presidential elections. The number of voters who participated in the election was 10676 during the first round and 32685 during the second round. Any eligible citizen can submit her nomination to the platform and those who gets at least 500 supports are the candidates of the election. The vote is conducted in two rounds. In the first round each voter is asked to express her judgment, using MJ, on five random candidates. At the end of this phase the five candidates with the highest medians are considered the finalists who qualify for the second round. In the second round each voter is asked to express her judgment, using MJ, on all the five finalists. The candidate with the best median at the end of this phase is selected as representative for the presidential election. However, the participation of this candidate to the actual presidential election in France is not granted. In fact, by the french law a candidate must collect at least 500 signatures of elected officials in order to participate to the presidential election. The candidate selected by the voters of LaPrimaire.org in 2016 collected only 135 signatures and did not participate in the 2017 presidential elections.

Criticism and concerns. Despite his public popularity, MJ has received several criticisms from part of the scientific community. In particular, Felsenthal and Machover [2008], Zahid [2009] and Laslier [2018] have shown that this procedure can lead to undesirable and counter-intuitive results. The authors of MJ highlight the various strengths of their model insisting on:

- the maximization of voters expressiveness —as we saw in Section 2.2;
- the satisfaction of anonymity, neutrality, unanimity, monotonicity and independence from irrelevant alternatives —the axiom for which if the winner of an election is a , and a new candidate b is added, then the winner will be either a or b ;
- the immunity to candidate cloning —the axiom for which if a new candidate b similar to a candidate already existing a is added, then the chances of a to win do not change;

- the fact that truth is a dominant strategy, i.e. a single voter has no incentives to lie.

However, Felsenthal and Machover [2008], Laslier [2018] show that the monotonicity axiom holds only for a fixed population. This leads to the no show paradox [Fishburn and Brams, 1983]: a voter can obtain a more desirable outcome if she does not participate in the election than the one she obtains by participating and voting sincerely. Felsenthal and Machover [2008], Zahid [2009] prove that the two mechanisms to resolve ties proposed by Balinski and Laraki [2011] as equivalents, do not always yield the same result. Other criticisms involve the *consistency* axiom, also known as *reinforcement* [Young, 1974, 1975]. The axiom says that if a candidate a wins in two separate elections, then this candidate must also win when combining the two electorates. Felsenthal and Machover [2008] show that MJ violates consistency. It is worth mentioning that the authors of MJ defend their model against criticisms [Balinski and Laraki, 2011, Balinski, 2019]. They believe that many of the problems mentioned are not major concerns in practice but, instead, arise from particular toy cases. Moreover, the authors claim that no voting model is perfect and are convinced that the pros of MJ outweigh the cons.

In this manuscript we refrain from personal judgments on the voting rule. We believe that it is worth to investigate any rule that is used in practical applications. In Chapter 6 we will analyze MJ in more detail, in particular in a situation that has not been studied yet: incomplete information.

3.3 Preference elicitation

All the methods described so far assume complete information about the preference orders of all voters over the entire set of alternatives, and about the voting mechanism itself. But how reasonable is this assumption? If the set of alternatives is very large, it is unreasonable to expect voters to provide a complete order of preferences. First from a cognitive point of view: experimental psychological studies describe how in many situations our preferences are not predefined, but rather we construct them only at the moment when we have to express them [Lichtenstein, 2006]. Asking voters to express preferences when the set of choices is very wide can lead to various problems, such as confusion and low participation. Second, but no less important, this represents a huge communication cost. Conitzer and Sandholm [2005] studied the communication complexity of some of the most common voting rules and, for each, provided an upper and lower bound on the number of bits of information that voters are required to communicate before the rule can select a winner. Other metrics exist, such as the number of questions voters must answer before the rule can determine a winner. This approach was used by Procaccia [2008] to determine the minimum number of questions to select the Condorcet winner. In presidential elections, for example, we tend to prefer accuracy over cost, but this is not true in all applications. When it comes to choosing a restaurant to have dinner at with friends, we are likely to settle for an approximate winner rather than having to spend the entire evening ranking all the restaurants in town.

Possible and necessary winners. On these basis, Konczak and Lang [2005] introduced the concepts of possible and necessary winners. Given a partial profile there are many ways to add preferences in order to build a complete profile. If an alternative a is a winner in some of these completion then a is a possible winner in the partial profile. If a is a winner in *all* completion then a is a necessary winner. This information can be very important during an elicitation process because if we reach a necessary winner then we may decide to stop the elicitation process. Numerous studies have been published on the computational complexity of finding possible and necessary winners [Konczak and Lang, 2005, Pini et al., 2007, Walsh, 2007, Xia and Conitzer, 2008, 2011, Baumeister and Rothe, 2012]. See Lang [2020] for a survey on possible and necessary winner not only in voting but also in other social choice scenarios like fair division, partial tournaments and hedonic games.

Although being an important concept from a theoretical point of view, the problem still persists. It is often necessary to know all, or almost all, preference rankings in order to compute a necessary winner. As for possible winners, yes, they can be used to reduce the set of alternatives, but how to choose the right approximation? Boutilier et al. [2006], and later Lu and Boutilier [2011a], proposed the use of minimax regret: the alternative a that minimizes the regret of choosing a instead of the winner in the worst case—which is the winner in the worst possible completion of the preference profile from a 's perspective—is the best "current" approximation. Our elicitation procedure presented in Chapter 4 is based on this approach, but others have been suggested. For example, Bachrach et al. [2010] proposed a probabilistic method in which they considered for each possible winner the number of completions for which it is a winner. Lu and Boutilier [2011b] suggested a mix of these two strategies to determine a minimal value k

such that by asking the voters their best k alternatives is it possible to determine an approximate winner with a low regret.

Building an elicitation procedure. This brings us to another point, established that in many cases it is convenient to ask for only partial information rather than to demand the complete preferences from the beginning, how to conduct this elicitation process? And when to stop it? Conitzer and Sandholm [2002], and Walsh [2008] later on, studied the complexity of determining when to stop the elicitation process and concluded that it depends on the choice of elicitation procedure and voting rule. The authors proved that for some rules, like Single Transferable Vote (STV), the problem can be resolved in a polynomial time if voters are asked to provide their full preference rankings but it becomes NP-hard when considering pairwise comparison queries that ask voters to compare two specific alternatives. However, those two problems are both polynomials for some other rules like Borda and plurality.

Although there are some studies investigating the "one-shot" preference elicitation, as the work of Lu and Boutilier [2011b] cited earlier, much attention has focused on incremental preference elicitation. Kalech et al. [2011] proposed both approaches. In their first algorithm, *Iterative Voting*, at each round all voters are asked, at the same time, to provide their (next) best alternative. Possible and necessary winners are computed and once found one (or all) necessary winner(s) the elicitation stops. Their second algorithm, *Greedy Voting*, instead, ask all voters their best k alternatives and then stops the elicitation procedure. However, this method does not guarantee to find a necessary winner. Lu and Boutilier [2011a] also proposed an iterative procedure that use the minimax regret to find the next question to ask, as already mentioned, but also to decide when to stop the elicitation process. In fact, they define a threshold such that if the regret is lower than this threshold the current approximate winner is returned. Benabbou et al. [2016] extended this work to multi-attribute domains—settings in which voters express their preferences over a set of interrelated alternatives, e.g. a multiple referendum. The authors provided an iterative elicitation procedure that stops as soon as a Borda winner can be determined and also returns an aggregated ranking of the best k alternatives. Naamani-Dery et al. [2015] proposed two incremental procedures that differ in their heuristic for choosing the next question. Their first approach selects the query, among the possible ones, that maximizes the probability of information gain. The second mechanism selects first the candidate most likely to win and then identifies the query that maximizes its possibility of winning.

There are other aspects of preference elicitation that we will not cover, for example the complexity of queries, or situations where the set of alternatives is uncertain. For more on these topics, we recommend reading Brandt et al. [2016, Ch. 10].

Uncertainty of the voting rule. In this manuscript, we will instead deal with an aspect that has not been treated much in the literature: the uncertainty of the voting rule. It is often assumed that the uncertainty of the voting rule may be due to an attempt of manipulation by the body responsible for defining the procedure. In the concluding remarks of his work, Walsh [2007], mentioned as a future direction:

“Other forms of uncertainty have yet to be studied in detail. For instance, there may be uncertainty in the weight vector used by a scoring rule. The Chair may want to manipulate the

election by collecting all the votes and then deciding on the weights in such a way that ensures that their preferred candidate wins. As a second example, the Chair may want to hinder strategic voting by not announcing the weights in advance or by choosing weights randomly.”

A similar perspective has been adopted by subsequent works dealing with the uncertainty of the voting rule. These, just like the second part of Walsh consideration, focused on the impact of this uncertainty on voters’ manipulation and strategic behavior [Baumeister et al., 2011, Elkind and Erdélyi, 2012, Holliday and Pacuit, 2019].

However, we do not believe that this is the only interesting scenario. Similarly to the case of preference profile uncertainty where it is difficult to obtain voter preferences from a cognitive and computational point of view, it may be difficult for a committee of non-experts to define an aggregation procedure. Consider the example suggested by Walsh where a committee has to formalize a scoring rule. The members of the committee may have an idea about the importance of weights, but it may be very difficult for them to properly define the scoring vector. Consider, for example, the board of a company that must decide how to aggregate experts opinions about the candidate to hire. Maybe they consider that being ranked first is "much more" valuable than being ranked second and so on, but how to define this "much more"? Should the first position be worth twice the second position? On which basis can non-experts choose this parameter? It is important to mention that the problem of how to choose the weights associated with a scoring rule has been extensively studied and it has been shown how a winner with given weights might not be one when changing the scoring vector [Cook and Kress, 1990, Llamazares and Peña, 2013, Llamazares, 2016].

To the best of our knowledge very few works tackle the problem of eliciting a voting rule. We mention the works of Llamazares and Peña [2013], Viappiani [2018], that considers scoring rules with uncertain weights, and that of Cailloux and Endriss [2014] that proposes a framework for explaining the consequences of choosing a specific set of axioms to non-expert users so that they can take an informed decision on which rule to choose. We use this methodology of reasoning through examples when building questions to the committee to elicit their preference.

In Chapter 4, we propose an elicitation procedure that considers both preference profile and scoring vector as sources of uncertainty at the same time. This represents a novelty in the literature of voting under incomplete preference. It is important to mention that the elicitation process is conducted by a hypothetical third party external to the voting process. The committee cannot be influenced by voters preferences and vice versa. We focus on positional scoring rules with convex weights and develop elicitation strategies based on minimax regret.

Types of questions and preferences inconsistency. We consider two kind of questions: the questions to the agents are pairwise comparisons between two alternatives; the questions to the chair ask to select a winner out of an example profile. Empirical psychological studies show how alternative phrasing of the same question give rise to different preferences [Tversky and Kahneman, 1986]. For example, in an experiment about this phenomenon conducted by Irwin et al. [1993], several participants were asked to evaluate problems such as:

- 1 Improving the air quality in Denver.
- 2 Upgrading their TV.

When asked to choose an alternative in a pairwise comparisons between those two options, an

overwhelming majority choose to improve the air quality. However, when participants were presented these options separately, most of them were willing to pay more for upgrading their TV.

There is no best method of asking questions because all of them, in one way or another, are bound to produce inconsistencies. However, we want to provide some considerations on the two chosen formulations.

The Analytic Hierarchy Process (AHP) is a mathematical model, widely used in decision making theory, to make complex decision. It was proposed by Saaty [1986] and it is based on the law of comparative judgment conceived by Thurstone [1927]. The latter is a model to retrieve information through the use of pairwise comparisons. Thus, AHP is a model that, to arrive to a complex decision, uses pairwise comparisons between the various alternatives. Although it has some criticisms, such as the inconsistency of preferences that we mentioned earlier, it is one of the most used decision-making methods. Dividing a complex problem into a hierarchy of smaller, more addressable sub-problems allows users to analyze different aspects at the time, making the choice less complicated.

Regarding the uses of example profile to elicit preferences about the voting rule, other than the work of Cailloux and Endriss [2014], we want to mention a real life experiment performed by Giritligil et al. [2005]. The authors of the latter paper, in an attempt to understand how non-experts perceived the majoritarian approval principle, created a survey in which they presented different profiles to 288 participants who had to choose a winner in each one of them. In each profile they accurately created a Majoritarian Compromise-winner, a Borda-winner and a Condorcet-winner, making sure that the first is always distinct from the other two. They concluded that, although many of the participants consistently chose the Borda-winner, the Majoritarian-winner was also well supported.

Preference elicitation in Decision Aiding. The problem of information elicitation is extensively studied in the domain of decision aiding. In this context, there is a *Decision Maker (DM)* who seeks the help of an expert party, an *analyst*, to develop recommendations that the DM can later use to make her choices. In this process, there is a stage where the analyst has to formalize the problem in the form of a mathematical model. To do so, it is important to get the preferences of the DM, who, as a non-expert, often fails to identify the parameters relevant to the problem. It is the job of the analyst to ask the right questions to obtain the preference information. Different approaches to elicitation have been proposed based on the choice of those questions, but also the availability (and patience) of the DM, or whether the DM is a real person or a digital database. Please see Bouyssou et al. [2006] for a more detailed account on the matter and Mousseau [2005], in the context of Multiple Criteria Decision Aid, and Belahcene [2018] for an overview of these methods. In particular, the latter, much like Cailloux and Endriss [2014], proposes a framework for the explanation of the elicitation process itself. In fact, similarly to how Cailloux and Endriss [2014] tries to elicit the voting rule through concrete examples, Belahcene [2018] tries to do the same to elicit preference information instead of explaining the result of the recommendation. As a side note, it is also interesting to mention that it is often assumed that in an elicitation process the analyst tries to elicit information that the DM has but which is unknown by the analyst. This is not always the case. Bouyssou et al. [2000, Ch. 8]

presented a scenario of information uncertainty with a real case example. A company dealing with the production and distribution of electricity—which here represents the DM—asked the help of a research group—which embodies the analyst—for the realization of a decision-aiding tool that would help them with the company choices. Every three years the company had to reevaluate its plan of building or closing its various (coal, nuclear, and gas) power plants. When formalizing the model, the authors realized that the values to assign to the various criteria were mostly uncertain. For example, the fuel price fluctuates according to the market, the electricity demand may depend on the weather, and a specific pollution legislation may have an impact on the company policy. The authors proposed a model based on successive pairwise comparisons: at each step it compares the consequences of two different decision in the same scenario. The model was adopted by the company, however, the authors showed some drawbacks of the model, in particular how it can lead to inconsistencies, and discussed why they did not adopt a probabilistic approach. Although it is not in the scope of our research, because we assume that the unknown information exists somewhere but is unknown to us, this consideration opens the door to many other potentially interesting scenarios to investigate.

Preference Elicitation in Machine Learning. In general, when we think of machine learning problems we tend to divide them into classification problems—in which the goal is to predict the category into which a particular instance falls—and regression problems—in which the purpose is to predict the value of a particular variable associated with an instance. However, many real-life problems are based on preferences of individuals. Consider, as an example, the enormous amount of information on the Web. It is impossible for users to find something without efficient Information Retrieval (IR) mechanisms. Given a query submitted by the user, an IR mechanism provides a list of documents ranked in order of relevance. The ranking is the key element that defines a good system. We might approach the problem by predicting whether or not a document is relevant to the query, as in a classification problem, or by calculating the probability that a document is relevant, like a regression problem. However, we are leaving out an important piece of information that is the correlation between the documents that are not independent variables but are all potential answers to the query. This makes ranking a fundamentally different problem than classification and regression. Document retrieval is only one example of the use of rankings in IR, a detailed account of which can be found in Liu [2009], but other uses can be found like recommender systems.

The uses of preferences and rankings play a key role in many machine learning problems. The process of learning preference models from already existing and, available, preference information is denoted as *Preference Learning*. Fürnkranz and Hüllermeier [2010] provided a comprehensive overview of this sub-field of machine learning. Given a set of items and a set of preferences over these items, the goal of a *preference learning task* is to learn a function that predicts preferences for a different set of items producing, as a result, a total order of the whole set of items—this task is denoted as *Object Ranking*—, or that predicts preferences of new users over the same set of items—denoted as *Label Ranking*. The training preference information is usually given as a set of pairwise preferences. This looks quite similar to our elicitation problem, where we have incomplete information that we need to learn. However, there are many differences between the two approaches. First, in elicitation we try to learn about each individ-

ual preferences in order to find the most accurate approximation, in preference learning, on the other hand, the goal is to find a good prediction that is good on average on the entire population. Moreover, preference elicitation often assumes that, as seen in the psychological studies aforementioned, preferences are constructed concurrently with the elicitation process. This involves a third party who communicates interactively with each individual. In cases of inconsistencies this party could notice and correct them, perhaps by asking clarifying questions to the individuals who produced such inconsistency. In preference learning, on the contrary, a ground truth is assumed and the quality of the prediction is measured based on its differences from the truth. Moreover, these preferences exist in a database and are always accessible but there is no interactive process that can help correct inconsistencies. Note that, in Chapter 4, we assume the existence of a generic external party, that could represent the analyst in decision aiding, but we do not implement the idea of resolving inconsistencies.

There are, however, works that position themselves between these two approaches filling the gaps. For example, in Chapter 4 we assume the existence of a ground truth, and Bachrach et al. [2010], Lu and Boutilier [2011b] studied preference elicitation in a probabilistic setting. On the machine learning side, many studies try to put "humans-in-the-loop" to receive feedback on predictions and, in general, to get more accurate results [Wu et al., 2022]; and active learning relaxes the hypothesis under which data are always accessible and instead queries users to provide labels for instances [Settles, 2009].

PART II

CONTRIBUTIONS

SIMULTANEOUS ELICITATION OF SCORING RULE AND AGENT PREFERENCES FOR ROBUST WINNER DETERMINATION

Social choice deals with the problem of determining a consensus choice from the preferences of different agents. In the classical setting, the voting rule is fixed beforehand and full information concerning the preferences of the agents is provided. This assumption of full preference information has recently been questioned by a number of researchers and several methods for eliciting the preferences of the agents have been proposed. In this paper we argue that in many situations one should consider as well the voting rule to be partially specified. Focusing on positional scoring rules, we assume that the chair, while not able to give a precise definition of the rule, is capable of answering simple questions requiring to pick a winner from a concrete profile. In addition, we assume that the agent preferences also have to be elicited. We propose a method for robust approximate winner determination and interactive elicitation based on minimax regret; we develop several strategies for choosing the questions to ask to the chair and the agents in order to converge quickly to a near-optimal alternative. Finally, we analyze these strategies in experiments where the rule and the preferences are simultaneously elicited.

4.1 Introduction

Aggregation of preference information is a central task in many computer systems (recommender systems, search engines, etc). In many situations, such as in group recommender systems, the goal is to find a consensus choice; social choice theory can provide foundations for such applications. The traditional approach to social choice assumes that 1) the full preference orderings of the agents and 2) the social choice function are expressed beforehand. These represent two strong hypotheses. Requiring agents to express full preference orderings can be prohibitively costly (in terms of cognitive and communication cost). This observation has motivated several works assuming partial preference orders: one early work is by Conitzer and Sandholm [2005] who studied the complexity of communication when using different voting rules; Konczak and Lang [2005] studied the computation of possible and necessary winners for various voting rules; Xia and Conitzer [2008] then showed that, while the identification of a necessary co-winner in scoring rules is polynomial, the determination of possible co-winners is

NP-hard; additional complexity results were given by Walsh [2007] and Pini et al. [2007].

Since in many practical situations there would be too many possible winners but no necessary winners, several works addressed the problem of agent preferences elicitation using a variety of approaches (minimax regret, Bayesian methods, etc.) with the goal of converging to a necessary winner [Naamani-Dery et al., 2015, Kalech et al., 2011, Lu and Boutilier, 2011a, Pini et al., 2009, Benabbou et al., 2016, Dey and Misra, 2016]. Among those, Walsh [2008] and Conitzer [2009] analyzed when to stop the elicitation process.

A second concern is the ability of the chair (the person or organization supervising the voting process) to provide a precise definition of the voting rule, suggesting the relaxation of the second hypothesis. Indeed, it is often difficult for non-experts to formalize a voting rule on the basis of some generic preferences over a desired aggregation method. Here we provide two examples of such situations.

Consider, as a first example, a chair that is about to hire a new employee whose performances are evaluated by several experts. The members of the chair may not have a voting rule in mind at the start of the process, and might not wish to agree on a specific voting rule. However, they might be willing to answer a few questions requiring to select who should be the winner out of specific profiles.

Consider, as a second example, the reviewing process of a conference where the best paper must be elected. The agents express their preferences on the papers they reviewed, but they are not aware of the voting rule the Program Chair will apply when aggregating them. Nonetheless, reviewers are still willing to participate in the process. Also, the PC may not have a specific voting rule in mind, and she will find it hard to provide a precise scoring vector if asked. Maybe she strongly believes that being ranked once in the first position is “much more” valuable than being ranked two times second, but does not know exactly how much more (though she can judge example cases).

In this paper, we focus on positional scoring rules with convex weights, that are a particularly common method used to aggregate rankings. We develop methods, based on the notion of minimax regret, for determining a robust “winner” under uncertainty of both the voting rule and the agent preferences. We provide incremental elicitation methods that at each step of the elicitation question either one of the agents or the chair, and we discuss several heuristics to choose questions that quickly reduce the regret. Answers to questions are encoded as constraints; questions to the agents are comparisons between pairs of alternatives while questions to the chair ask to select a winner out of a synthetic profile.

While some previous works have considered partially specified aggregation methods [Stein et al., 1994, Llamazares and Peña, 2013, Viappiani, 2018], we do not know of any work considering both sources of uncertainty at the same time. Actually, very few works altogether have considered the problem of eliciting a voting rule by asking questions to the chair. We mention the work of Cailloux and Endriss [2014] that assumes a different representation for the rule. Additionally, some works address the manipulability of voting rules [Elkind and Erdélyi, 2012, Dey et al., 2018, Conitzer et al., 2011, Baumeister and Hogebe, 2019] and strategic behaviors [Endriss et al., 2016, Lev et al., 2019, Reijngoud and Endriss, 2012].

Our approach is evaluated on simulations with synthetic and real datasets where both the voting rule and the agent preferences are initially unknown to the system and incrementally re-

vealed through questioning. We assume the chair to be human, thus able to answer questions about a limited number of alternatives, so we focus on small scale social choice situations. We compare the effectiveness of several questioning strategies based on the current knowledge of the rule and preferences. To summarize our contributions: 1) we provide a novel mechanism for eliciting a voting rule by translating abstract questions about weights to a choice of an alternative given a concrete profile; 2) we show that with our elicitation method it is possible to reach low regret with a reasonable number of questions; 3) we present elicitation strategies that achieve good results within reasonable computation time; 4) we show that for the class of rules considered, asking a few questions to the chair suffice to reach low regret; 5) our experiments suggest that low degree of similarity among preferences (as in impartial culture) is a more challenging setting than less varied profiles.

4.2 Social choice with partial information

We now introduce some basic concepts. We consider a set A of m alternatives (products, restaurants, public projects, job candidates, etc.) and an infinite set $\mathbb{N}$ of potential agents.

A *profile* $(\succ_j, j \in N)$ considers a finite subset of agents $N \subset \mathbb{N}$ and associates to each agent a preference order $\succ_j \in \mathcal{L}(A)$, a linear order over the alternatives. A profile is equivalently represented by $\mathbf{v} = (v_j, j \in N)$ where $v_j(x) \in \{1, \dots, m\}$ denotes the rank of alternative x in the preference order $\succ_j$. A social choice function $f : \cup_{\emptyset \neq N \subset \mathbb{N}, N \text{ finite}} \mathcal{L}(A)^N \rightarrow \mathcal{P}^*(A)$ associates to each profile a set of (tied) winners, where $\mathcal{P}^*(A)$ is the powerset of A excluding the empty set. Among the many possible social choice functions, we consider convex *positional scoring rules* (PSRs). A PSR $f^{\mathbf{w}}$ is parameterized by a *scoring vector* $\mathbf{w}$ associating weights $w_r \in [0, 1]$ to positions, with $1 = w_1 \geq w_2 \geq \dots \geq w_m = 0$. Let α_r^x be the number of times that alternative x was ranked in the r -th position. Given $\mathbf{v}$ and $\mathbf{w}$, an alternative $x \in A$ obtains the score

$$s(x; \mathbf{v}, \mathbf{w}) = \sum_{j \in N} w_{v_j(x)} = \sum_{r=1}^m \alpha_r^x w_r. \quad (4.1)$$

The winners $f^{\mathbf{w}}(\mathbf{v})$ are the alternatives with highest score.

An important class of PSRs is the one using convex weights [Stein et al., 1994, Llamazares, 2016], meaning that the difference between the weight of the first position and the weight of the second position is at least as large as the difference between the weights of the second and third positions, etc.

$$\forall r \in \{1, \dots, m-2\} : w_r - w_{r+1} \geq w_{r+1} - w_{r+2}. \quad (4.2)$$

The constraint above is a natural and common assumption, often used when aggregating rankings in sport competitions (such as F1 racing, alpine skiing world cup): losing ranks at the top is more damaging than losing ranks at the bottom. Let $\mathcal{W}$ denote the set of such convex weight vectors.

We consider a specific finite set of agents $N^* \subset \mathbb{N}$ and let $\mathbf{v}^* = (\succ_j^*, j \in N^*)$ and $\mathbf{w}^*$ denote the profile and weight vector, unknown to us, that represent the preferences of the agents in N^* and of the chair.

At a given time, our knowledge of agent j 's preference is encoded by a partial order $\succ_j^p \subseteq \succ_j^*$ over the alternatives, a transitive and asymmetric relation (we assume that preference information is truthful). An incomplete profile $\mathbf{p} = (\succ_j^p, j \in N^*)$ maps each agent to a partial preference. Let $C(\succ_j^p) = \{\succ \in \mathcal{L}(A) \mid \succ_j^p \subseteq \succ\}$ denote the set of possible completions of $\succ_j^p$ and $C(\mathbf{p}) = \prod_{j \in N} C(\succ_j^p)$ the set of complete profiles extending $\mathbf{p}$. Note that $\mathbf{v}^* \in C(\mathbf{p})$.

The vector $\mathbf{w}^*$ is also unknown but we assume that the chair is able to specify additional preference information taking the form of linear constraints about $\mathbf{w}^*$. Let $W \subseteq \mathcal{W}$ denote the set of weight vectors compatible with the preferences expressed by the chair about the scoring vector. We will show in Section 4.4 that the additional preferences we use can be elicited by showing a complete profile of a synthetic election and asking who should be elected in this case.

4.3 Robust winner determination

It is desirable in an elicitation protocol such as ours to be able to stop before reaching full knowledge of the agent preferences or of the preferences of the chair about the voting rule. As, often, there are no necessary winners and too many possible winners, it is useful to declare a winner given partial information. As a decision criterion to determine a winner, we propose to use minimax regret [Savage, 1954]. This decision criterion has been used for robust optimization under data uncertainty [Kouvelis and Yu, 1997] as well as in decision-making with uncertain utility values [Salo and Hämäläinen, 2001, Boutilier et al., 2006]. In particular, Lu and Boutilier [2011a] have adopted minimax regret for winner determination in social choice where the preferences of agents are partially known, while the social choice function is known.

We consider the simultaneous presence of incomplete knowledge in agent preferences and in the weights of the PSR. We use *maximum regret* to quantify the worst-case error, and let the alternatives that minimize this quantity win, giving some robustness in face of ignorance. Intuitively, the quality of a proposed alternative a is how far a is from the optimal one in the worst case, given the current knowledge.

Given $\mathbf{p}$ and W (that represent the current knowledge about agent preferences and the PSR), the maximum regret is considered by assuming that an adversary can both 1) extend the partial profile $\mathbf{p}$ into a complete profile, and 2) instantiate the weights choosing among any weight vector in W . We formalize the notion of minimax regret in multiple steps. First of all, $\text{Regret}(x, \mathbf{v}, \mathbf{w})$ is the “regret” of selecting x as a winner instead of the optimal alternative under $\mathbf{v}$ and $\mathbf{w}$:

$$\text{Regret}(x, \mathbf{v}, \mathbf{w}) = \max_{y \in A} s(y; \mathbf{v}, \mathbf{w}) - s(x; \mathbf{v}, \mathbf{w}).$$

The *pairwise maximum regret* of x relative to y given the partial profile $\mathbf{p}$ and the set of weights W is the worst-case loss of choosing x instead of y under all possible realizations of the full profile *and* all possible instantiations of the weights:

$$\text{PMR}(x, y; \mathbf{p}, W) = \max_{\mathbf{w} \in W} \max_{\mathbf{v} \in C(\mathbf{p})} s(y; \mathbf{v}, \mathbf{w}) - s(x; \mathbf{v}, \mathbf{w}).$$

The maximum regret is the worst-case loss of x :

$$\text{MR}(x; \mathbf{p}, W) = \max_{y \in A} \text{PMR}(x, y; \mathbf{p}, W) = \max_{\mathbf{w} \in W} \max_{\mathbf{v} \in C(\mathbf{p})} \text{Regret}(x; \mathbf{v}, \mathbf{w}). \quad (4.3)$$

$\text{MR}(x; \mathbf{p}, W)$ is the result of an adversarial selection of the complete profile $\mathbf{v} \in C(\mathbf{p})$ and of the scoring vector $\mathbf{w} \in W$ that jointly maximize the loss between x and the true winner under $\mathbf{v}$ and $\mathbf{w}$. Finally, $\text{MMR}(\mathbf{p}, W) = \min_{x \in A} \text{MR}(x; \mathbf{p}, W)$ is the value of *minimax regret* under $\mathbf{p}$ and W , obtained when recommending a *minimax optimal* alternative $x_{\mathbf{p}, W}^* \in A_{\mathbf{p}, W}^* = \text{argmin}_{x \in A} \text{MR}(x; \mathbf{p}, W)$. Picking as consensus choice an alternative associated with minimax regret provides a recommendation that gives worst-case guarantees. In cases of ties, we can return all minimax alternatives $A_{\mathbf{p}, W}^*$ as winners or pick one of them using some tie-breaking strategy.

Observe that if $\text{MMR}(\mathbf{p}, W) = 0$, then any $x_{\mathbf{p}, W}^* \in A_{\mathbf{p}, W}^*$ is a necessary winner: any valid completion of the profile and choice of $\mathbf{w} \in W$ gives to $x_{\mathbf{p}, W}^*$ the highest score.

We note that our notion of regret gives some cardinal meaning to the scores: instead of just being used to select winners under the corresponding PSR, their differences are considered as representing the regret of the chair.

Computation of minimax regret Given a voting rule and a partially specified profile, Xia and Conitzer [2008] determine necessary winners by showing constructions that attempt to maximize the score difference between a proposed winner and a chosen alternative. This reasoning was later adopted by Lu and Boutilier [2011a] who used the considerations on the worst-case completions for computing the minimax regret.

In order to compute pairwise maximum regret, and therefore minimax regret, we decompose the PMR into the contributions associated to each agent by adapting this same reasoning to our setting. The context is however more challenging due to the presence of uncertainty in the weights.

Recall that, in the computation of $s(x; \mathbf{v}, \mathbf{w})$, $w_{v_j(x)}$ represents the score that x obtains in the ranking v_j (see Eq. (4.1)). Since scoring rules are additively decomposable, we can consider separately the contribution of each agent to the total score. Thus, we can write the actual regret of choosing x instead of y as $s(y; \mathbf{v}, \mathbf{w}) - s(x; \mathbf{v}, \mathbf{w}) = \sum_{j \in N} w_{v_j(y)} - w_{v_j(x)}$, and we obtain

$$\text{PMR}(x, y; \mathbf{p}, W) = \max_{\mathbf{w} \in W} \sum_{j \in N} \max_{v_j \in C(\succ_j^{\mathbf{p}})} [w_{v_j(y)} - w_{v_j(x)}].$$

The following propositions show that the procedure for completing a partial profile, proposed by Lu and Boutilier [2011a] when considering a fixed weight vector, also applies in our setting. We write $a \succeq_j^{\mathbf{p}} b$ iff $a \succ_j^{\mathbf{p}} b \vee a = b$ and adopt the canonical notation when considering a relation as a function, writing $\succeq_j^{\mathbf{p}}(x)$ for $\{y \mid x \succeq_j^{\mathbf{p}} y\}$.

Proposition 4.1. *There exists a completion $\hat{\mathbf{v}} \in C(\mathbf{p})$ of the partial profile $\mathbf{p}$ such that $\text{PMR}(x, y; \mathbf{p}, W) = \max_{\mathbf{w} \in W} [s(y; \hat{\mathbf{v}}, \mathbf{w}) - s(x; \hat{\mathbf{v}}, \mathbf{w})]$ and such that the linear order $\hat{v}_j$ of each agent j satisfies:*

$$a \succ_j x \iff \neg(x \succeq_j^{\mathbf{p}} a); \quad (4.4)$$

$$y \succ_j a \iff \neg(a \succeq_j^{\mathbf{p}} y) \wedge \neg((x \succeq_j^{\mathbf{p}} y) \wedge \neg(x \succeq_j^{\mathbf{p}} a)). \quad (4.5)$$

Proof Sketch. Consider our knowledge $\succeq_j^{\mathbf{p}}$ about the preference of the agent j . The adversary's goal is to make the score of y as high as possible and the score of x as low as possible. To do

this, he should complete $\succ_j^p$ to $\succ_j$ by placing above x as many alternatives as possible; that is, all the alternatives except those that are known to be worse than x (those a such that $x \succeq_j^p a$); and similarly, he should put below y all the alternatives he can. Two conditions must be excluded for a to go below y . The alternatives such that $a \succeq_j^p y$ can't be put below y . Furthermore, the first objective must take priority over the second one: when an alternative should go above x according to the first objective (because $\neg(x \succeq_j^p a)$), and x is known to be better than y (thus $x \succeq_j^p y$), then a should be put above x (irrespective of whether $a \succeq_j^p y$), which will move both x and y one rank lower than if a had been put below y . This maximizes the adversary's interests: because the weight vector is convex, the score difference will be lower when both alternatives are ranked lower (Equation 4.2), and that difference of scores is in favor of x when $x \succ_j^p y$, thus to be minimized from the the adversary's point of view. $\square$

Proposition 4.2. *The rank of x in the PMR-maximizing linear orders of agent j is $\hat{v}_j(x) = 1 + |A| - |\succeq_j^p(x)|$, and the rank of y is $\hat{v}_j(y) = 1 + |\prec_j^p(y)| + |\beta|$, where $|\beta| = |A \setminus (\succeq_j^p(x) \cup \prec_j^p(y))|$ if $(x \succeq_j^p y)$ and $|\beta| = 0$ otherwise.*

Proof. The rank of x is directly obtained from Eq.(4.4). The rank of y is obtained by complementing Eq.(4.5), obtaining $a \succeq_j y \Leftrightarrow (a \succeq_j^p y) \vee ((x \succeq_j^p y) \wedge \neg(x \succeq_j^p a))$, and, observing that $a \succ_j y \Leftrightarrow a \neq y \wedge a \succeq_j y$, obtaining that $a \succ_j y$ if and only if

$$(a \neq y) \wedge [(a \succeq_j^p y) \vee ((x \succeq_j^p y) \wedge \neg(x \succeq_j^p a))], \quad (4.6)$$

or equivalently, if and only if

$$(a \succ_j^p y) \vee ((x \succeq_j^p y) \wedge \neg(x \succeq_j^p a)). \quad (4.7)$$

Indeed, (4.6) $\Rightarrow$ (4.7), and (4.7) $\Rightarrow$ (4.6) because $(x \succeq_j^p y) \wedge \neg(x \succeq_j^p a) \Rightarrow a \neq y$ (as when $a = y$, $(x \succeq_j^p y)$ and $\neg(x \succeq_j^p a)$ are opposite claims). Suffices now to rewrite Eq. (4.7) to let the two disjuncts designate disjoint sets:

$$a \succ_j y \Leftrightarrow (a \succ_j^p y) \vee ((x \succeq_j^p y) \wedge \neg(x \succeq_j^p a) \wedge \neg(a \succ_j^p y)). \quad (4.8)$$

$\square$

Note that in Proposition 4.2, in the case $(x \succeq_j^p y)$, β is the number of alternatives incomparable with both x and y .

Proposition 4.3. *The PMR can be written as:*

$$\text{PMR}(x, y; \mathbf{p}, W) = \max_{w \in W} \sum_{j \in N} w_{\hat{v}_j(y)} - w_{\hat{v}_j(x)} = \max_{w \in W} \sum_{r=1}^m (\hat{\alpha}_r^y - \hat{\alpha}_r^x) w_i, \quad (4.9)$$

where $\hat{\alpha}_r^y$ (resp. $\hat{\alpha}_r^x$) is the number of times y (resp. x) has rank r in the complete profile $\hat{v}$ defined in Proposition 4.2.

Proposition 4.3 shows that PMR is linear in the weights. The pairwise max regret $\text{PMR}(x, y; \mathbf{p}, W)$ can thus be obtained by solving the following linear program defined on the variables $w_1, \dots, w_m$:

$$\max_{\mathbf{w}} \sum_{r=1}^m (\hat{\alpha}_r^y - \hat{\alpha}_r^x) w_r \quad \text{s.t. } w_1 = 1 \geq \dots \geq w_m = 0, \text{ Eq. (4.2) and } \mathbf{w} \in W. \quad (4.10)$$

The max regret $\text{MR}(x; \mathbf{p}, W)$ is determined by computing the pairwise regret of x with all other alternatives in A , and the recommended alternatives are the ones with least max regret. Observe that when the PMR of an alternative x (against some other alternative y) exceeds the best MR value found so far, we do not need to further evaluate x . This idea can be exploited using a minimax-search tree [Braziunas, 2012].

4.4 Interactive Elicitation

We propose an incremental elicitation method based on minimax regret. At each step, the system may ask a question either to one of the agents about her preferences or to the chair about the voting rule. The goal is to obtain relevant information to reduce minimax regret as quickly as possible. The elicitation can be terminated either after a given number of questions, or when the minimax regret is lower than a threshold (or when it drops to zero if we wish optimality).

Question types We distinguish between questions asked to the agents and questions asked to the chair. As *questions asked to the agents* we consider comparison queries relating two alternatives. The effect of a response to a question asked to an agent is the increase in our knowledge about the agent rankings, thus augmenting the partial profile $\mathbf{p}$. If agent j answers a comparison query stating that alternative a is preferred to b , then the partial order $\succ_j^{\mathbf{p}}$ is augmented with $a \succ_j^{\mathbf{p}} b$ and by transitive closure.

A bit more discussion is needed about *questions asked to the chair*. Such questions aim at refining our knowledge about the scoring rule; a response gives us a constraint on the weight vector $\mathbf{w}$. In particular, we want to obtain constraints of the type $w_r - w_{r+1} \geq \lambda(w_{r+1} - w_{r+2})$ for $r \in \{1, \dots, m-2\}$, relating the difference between the importance of ranks r and $r+1$ with the difference between ranks $r+1$ and $r+2$.

Building concrete questions for the chair Even if the chair might be considered able to answer directly such abstract questions, we want to ensure that these questions can also, in principle, be asked in a more concrete way: in terms of winners of example profiles. Such questions have clear semantics whose understanding can be assumed to be shared by the chair, contrary to abstract questions about weights. Moreover, this way of questioning the chair is independent of the voting rule that is being elicited; whereas questions about weights only make sense when considering PSRs. Asking who should win in specific profiles has been used in experimental settings investigating the feeling of justice of individuals [Girtligil et al., 2005], but, to the best of our knowledge, the use of such questions to systematically guide an elicitation process about voting rules is novel. This is similar to favor, in decision theory, direct choice

questions ("please choose either a or b") compared to, say, questioning the decision maker about the shape of her utility function. The former are considered "observable": acts of choice are translated to preference statements [Mas-Colell et al., 1995, Ch. 1].

Although questioning in terms of profiles and in terms of weights is logically equivalent in our setting, there is no a priori certainty that questioning the chair using different phrasing would yield logically equivalent answers: research in experimental psychology shows that participants' answers differ widely when changing the phrasing of preference-related questions [Lichtenstein, 2006]. To get out of such conundrums, we need a language considered "fundamental". Questions of the form "In this profile, who should win?" arguably provides such a natural language.

Thus, our task is to build a profile, given λ and $r \leq m - 2$, in such a way that the set of (tied) winners picked by the chair reveals whether $w_r - w_{r+1} \geq \lambda(w_{r+1} - w_{r+2})$.

Proposition 4.4. *Given a rational $\lambda = p/q > 1$ and a rank r between 1 and $m - 2$, there exists a profile P such that, for any weight vector $w \in \mathcal{W}$, $a \in f(P)$ iff $w_r - w_{r+1} \geq \lambda(w_{r+1} - w_{r+2})$ and $b \in f(P)$ iff $w_r - w_{r+1} \leq \lambda(w_{r+1} - w_{r+2})$, where f is the PSR parameterized with w .*

Proof. Define a linear order $>_1$ over A as placing a at rank r , b at rank $r + 1$, and the remaining alternatives arbitrarily. Define $>_2$ over A as placing a at rank $r + 2$, b at rank $r + 1$, and the remaining alternatives arbitrarily. Define an arbitrary linear ordering $>$ over $A \setminus \{a, b\}$. Define a linear order $>_3$ as placing a first, b second, and following the order of $>$ for the remaining positions. Finally, define a linear order $>_4$ as placing b first, a second, and following the *inverse* order of $>$ for the remaining positions.

Define P as the profile of $3(p + q)$ agents containing q times $>_1$, p times $>_2$, and $>_3$ and $>_4$ each $p + q$ times. As a result, a obtains the following ranks: q times r , p times $r + 2$, $p + q$ times first, and $p + q$ times second. The alternative b obtains the ranks $r + 1$, 2 and 1, each $p + q$ times. Consider any alternative $c \in A \setminus \{a, b\}$. Its score is maximal when it comes first in $>_1$, first in $>_2$ and first in $>$, by convexity of the weights. In that case, c is positioned at the ranks 1, 3 and m , each $p + q$ times.

Letting $s(x)$ denote the score of x at P , we obtain $s(a) = qw_r + pw_{r+2} + (p + q)w_1 + (p + q)w_2$, thus, $s(a) \geq (p + q)w_m + (p + q)w_1 + (p + q)w_2$; $s(b) = (p + q)w_{r+1} + (p + q)w_2 + (p + q)w_1$; and, $\forall c \in A \setminus \{a, b\}$, $s(c) \leq (p + q)w_1 + (p + q)w_3 + (p + q)w_m$. It follows that a or b maximize s (as $s(a) \geq s(c)$). We conclude by observing that $a \in f(P) \Leftrightarrow s(a) \geq s(b) \Leftrightarrow qw_r + pw_{r+2} \geq (p + q)w_{r+1} \Leftrightarrow w_r - w_{r+1} \geq (p/q)(w_{r+1} - w_{r+2})$, and similarly for $b \in f(P)$. $\square$

Example 4.1. *Suppose we want to ask the following question to the chair: $w_2 - w_3 \geq 2(w_3 - w_4)$. We show the profile in Figure 4.1a to the chair and ask who should win (each column is the preference of one agent). Both a and b have scores higher than c and d for all convex weights, thus either a or b will be picked under our hypothesis; and $s(a) \geq s(b) \Leftrightarrow w_2 + 2w_4 \geq 3w_3$. Figure 4.1b represents the same profile using a compressed view, the numbers in bold indicating the number of agents having the preference in the corresponding column. As the proof shows, constructed profiles require only four different linear orders.*

Elicitation strategies We develop several strategies for simultaneous elicitation of agent preferences and of the PSR. While it is of course possible to first fully elicit the agent preferences

Figure 4.1: Profile representing a question to the chair in extended (a) and compact (b) form.

										1	2	3	3
<i>c</i>	<i>d</i>	<i>d</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>b</i>	<i>b</i>	<i>b</i>		<i>c</i>	<i>d</i>	<i>a</i>	<i>b</i>
<i>a</i>	<i>c</i>	<i>c</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>a</i>	<i>a</i>	<i>a</i>		<i>a</i>	<i>c</i>	<i>b</i>	<i>a</i>
<i>b</i>	<i>b</i>	<i>b</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>d</i>	<i>d</i>	<i>d</i>		<i>b</i>	<i>b</i>	<i>c</i>	<i>d</i>
<i>d</i>	<i>a</i>	<i>a</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>c</i>	<i>c</i>	<i>c</i>		<i>d</i>	<i>a</i>	<i>d</i>	<i>c</i>
(a)										(b)			

and afterwards elicit weights, we want to investigate approaches that are able to recommend winning alternatives before obtaining complete knowledge of the profile or the rule. We define here various strategies; a strategy tells us, given the current partial knowledge $(\mathbf{p}, W)$, which question to ask next.

The *Random* strategy is used as a baseline. It first chooses equiprobably whether to question the chair or the agents. In the first case, it draws one rank in $1 \leq r \leq m - 2$ equiprobably, takes the middle of the interval of values for λ that are still possible considering our knowledge so far, and asks whether $w_r - w_{r+1} \geq \lambda(w_{r+1} - w_{r+2})$. In the second case, it draws equiprobably among the agents whose preference is not known entirely; it then draws an alternative a among those involved in some incomparabilities in $\succ_j^P$ and an alternative b among those incomparable with a in $\succ_j^P$.

Let $(x^*, \bar{y}, \bar{v}, \bar{w})$ be the current solution of the minimax regret, where x^* is the minimax optimal alternative and $\bar{y}, \bar{v}, \bar{w}$ the corresponding adversarial choices. The *Pessimistic* strategy considers a set of $n + (m - 2)$ candidate questions: one per agent, and one per rank (excluding the first and the last one which are known).

The candidate questions to the agents are chosen by extending the idea of Lu and Boutilier [2011a], that privilege learning about the relationship of x^* and $\bar{y}$ to the other alternatives if possible. Given $j \in N^*$, if x^* and $\bar{y}$ are incomparable in $\succ_j^P$, the candidate question concerns the pair $(x^*, \bar{y})$, otherwise, it concerns the pair (x^*, z) for some z incomparable to x^* (randomly chosen), or if none such z exist, the pair $(\bar{y}, z)$ for some z incomparable to $\bar{y}$, or, if both x^* and $\bar{y}$ are comparable to every alternatives in $\succ_j^P$, any incomparable pair is picked at random.

The candidate questions to the chair are determined as in the Random strategy.

Once having selected $n + m - 2$ candidate questions, the Pessimistic strategy selects the one that leads to minimal regret in the worst case. Assume that a question q_1 has type t_1 (being “chair” or “agent”), and leads to the new knowledge states $(\mathbf{p}_1, W_1)$ if answered positively and $(\mathbf{p}'_1, W'_1)$ if the answer is negative. Define

$$R_1^{\max} = \max\{\text{MMR}(\mathbf{p}_1, W_1), \text{MMR}(\mathbf{p}'_1, W'_1)\}$$

and

$$R_1^{\min} = \min\{\text{MMR}(\mathbf{p}_1, W_1), \text{MMR}(\mathbf{p}'_1, W'_1)\}\epsilon_t + \epsilon'_t.$$

The terms ϵ_t and ϵ'_t are real numbers associated to the type t of question; these parameters are used to fine tune the choice of the question type. Define similarly t_2 , $R_2^{\max}$ and $R_2^{\min}$

for q_2 . Pessimistic considers question q_1 to be better than q_2 iff $R_1^{\max} < R_2^{\max}$ or $[R_1^{\max} = R_2^{\max} \text{ and } R_1^{\min} < R_2^{\min}]$. In other words if the maximal *a posteriori* MMR of two questions are (approximately) equal, then it considers the (penalized) minimal MMR values.

The *Extended pessimistic* strategy uses the same criterion as the pessimistic strategy, but extending it to a bigger set of candidate questions, the same as those considered by the Random strategy. These candidate questions are then evaluated using the same operator as for the Pessimistic strategy. Extended pessimistic is applicable only to very small problem instances: its complexity is in $O(n^2m^5)$, because we consider $O(m^2)$ questions for each agent and need for each question to compute MMR twice, whose complexity is $O(nm^3)$.

The *Two phases* strategy is developed in order to investigate the effect of varying the proportion of questions of the two types, when asking all questions to the chair at the beginning or at the end. It is parameterized by q_c , the number of questions to be asked to the chair. The *Two phases-ca* variant first asks q_c questions to the chair, then $k - q_c$ questions to the agents, using in both cases Pessimistic to select the specific questions; whereas the *Two phases-ac* variant starts with $k - q_c$ questions to the agents, then questions the chair.


Finally, the *Elitist* strategy aims at uncovering as quickly as possible the top alternatives of all agents. For any agent j , it asks to compare an alternative currently undominated in $\succ_j^p$ with one that is currently incomparable. Thus, the top alternative for j will be known after having asked exactly $m - 1$ questions to j . After having asked $n(m - 1)$ questions to the agents, it questions the chair only, using the same approach as Pessimistic. This strategy can be expected to perform well when the chair assigns a large weight to the first rank, as compared to the other ranks. It is used to further challenge Pessimistic, which is not specifically tailored to such a situation.

4.5 Empirical Evaluation

We performed several numerical experiments using both real data and randomly generated profiles in order to validate our approach and test the performance of our elicitation strategies.

Given a problem size (m, n) , a number of questions k and a strategy to test, we first create an “oracle”, representing the true preferences of the agents (randomly generated or coming from real data) and the weights associated with the chair’s scoring rule (randomly generated). We start with empty knowledge ($\mathbf{p} = \emptyset, W = \mathcal{W}$) about the preference orderings of the agents and the weights of the chair. We obtain the first question to be asked using the strategy under test. We then use the oracle to answer the question and update the system’s knowledge, which is thus used to obtain the next question. This is repeated until k answers have been obtained, computing the resulting MMR values along the way for various values of k . We repeat this whole experiment a variable number of times, for a given (m, n, k) , and report the average resulting MMR and standard deviation sd . The sizes of the considered scenarios are comparable to the ones used by Cailloux and Endriss [2014].

The oracle is built as follows. For the real preferences, we used three datasets from PrefLib [Mattei and Walsh, 2013]: *T Shirt* (researchers voted on tee shirt designs; $m = 11, n = 30$), *Courses* (students voted on courses; $m = 9, n = 146$; referred to as AGH on PrefLib) and *Skate* (judges voted on skaters at the Euros Pairs Short Program; $m = 14, n = 9$). For the synthetic

Figure 4.2: Average MMR in problems of size $(5, 10)$ after k questions.Table 4.1: Average MMR in problems of size $(10, 20)$ after k questions assuming geometric weights.

k	Pes. $\pm$ sd	Eli. $\pm$ sd
0	20.0 $\pm$ 0.0	20.0 $\pm$ 0.0
50	16.0 $\pm$ 0.5	17.3 $\pm$ 0.4
100	12.5 $\pm$ 0.9	15.6 $\pm$ 0.4
150	9.6 $\pm$ 1.4	13.9 $\pm$ 0.8
200	7.4 $\pm$ 1.3	11.0 $\pm$ 1.1
250	5.3 $\pm$ 1.5	6.6 $\pm$ 0.8
300	3.5 $\pm$ 1.3	6.6 $\pm$ 0.8

datasets, we follow an Impartial Culture (IC) assumption: the linear order of each agent is drawn i.i.d. uniformly. We believe IC to be a challenging situation and expect the number of questions to ask, in order to reach a certain level of regret, to decrease with less varied profiles. To generate the scoring rule weights, we first draw $m-1$ numbers uniformly at random (in the interval $[0, 1]$ representing weight “differences”), normalize and sort them; a sequence of convex decreasing weights is then obtained by a decumulative sum. The penalty parameters for the Pessimistic and Extended pessimistic strategies are $\epsilon_{\text{chair}} = 1.1$, $\epsilon'_{\text{chair}} = 10^{-6}$, $\epsilon_{\text{agent}} = 1.0$ and $\epsilon'_{\text{agent}} = 0$.

Comparison of strategies Our first experiment concerns small size situations. Figure 4.2 compares some of our strategies in the case $m = 5, n = 10$ (variations around this size yield similar conclusions), where the results are averaged over 200 runs. We see that asking random questions does not allow to reach a low regret level even after having asked 100 questions, whereas a low regret level (MMR = 1) is reached by Pessimistic before having asked 60 questions. This also holds for other problem sizes. For instance, for $m = 10, n = 20$ and 500 questions, Random strategy reaches an average regret (over 20 runs) of 9.3 (± 0.7) and Pessimistic 0.5 (± 0.5). We notice that Pessimistic performs slightly better than Extended pessimistic, showing that Pessimistic chooses candidate questions wisely; this is good news since Pessimistic is

Table 4.2: Questions asked by Pessimistic strategy on several datasets to reach $\frac{n}{10}$ regret, columns 4 and 5, and zero regret, last two columns.

dataset	m	n	$q_c^{\text{MMR} \leq n/10}$	$q_a^{\text{MMR} \leq n/10}$	$q_c^{\text{MMR}=0}$	$q_a^{\text{MMR}=0}$
m5n20	5	20	0.0	[4.3 5.0 5.8]	5.3	[5.4 6.2 7.2]
m10n20	10	20	0.0	[13.9 16.1 18.4]	32.0	[19.7 21.8 24.7]
m11n30	11	30	0.0	[16.6 19.0 22.3]	45.2	[23.1 25.7 28.9]
tshirts	11	30	0.0	[13.1 16.6 19.6]	43.2	[28.2 32.0 35.6]
courses	9	146	0.0	[6.0 7.0 7.0]	0.0	[6.8 7.0 7.0]
m14n9	14	9	5.4	[30.3 33.5 36.7]	64.1	[37.6 40.5 44.3]
skate	14	9	0.0	[11.4 11.6 12.3]	0.0	[11.5 11.8 12.8]
m15n30	15	30	0.0	[25.0 29.5 33.7]		

much faster: it takes on average only 16s for a complete elicitation session (for $m = 5$, $n = 10$ and 100 questions), while Extended pessimistic takes 50s. Although their performance is close, Pessimistic performs systematically better in multiple runs of the experiment.


We also compared the Pessimistic strategy against Elitist in a situation specifically tailored to advantage Elitist. For that experiment specifically, instead of drawing the weights of the oracle randomly, we fix it to a “geometric” weight vector, such that $w_r - w_{r+1} = 2(w_{r+1} - w_{r+2})$, for all $r \leq m - 2$, so as to dramatically increase the importance of the weights associated to the top ranks. Even in that case, we see in Table 4.1 that Pessimistic performs better than Elitist.

Evaluation of Pessimistic Strategy Our next set of experiments evaluate the Pessimistic strategy in absolute terms. We first wonder how many questions should be asked in order to achieve low regret, fixed at $n/10$: this is equivalent to the difference of score of an alternative x that results from switching from a profile P to a profile P' where a tenth of the agents rank x last instead of first. Table 4.2, first five columns, contains the result: it displays, for each dataset, the number of questions asked to the chair ($q_c^{\text{MMR} \leq n/10}$), and the quartiles of the number of questions asked to the agents ($q_a^{\text{MMR} \leq n/10}$), averaged over 20 runs. It is interesting to note that about twenty or thirty questions per agent on average suffice to reach a low regret in those instances. We find also noteworthy that the Pessimistic strategy chooses to ask zero questions to the chair but still achieves low regret, in most of those instances.

Another interesting measure is the average number of questions asked to the chair ($q_c^{\text{MMR}=0}$) and to the agents ($q_a^{\text{MMR}=0}$) before reaching zero regret. The results for various sizes are displayed in the last two columns of Table 4.2. Here, we see that the Pessimistic strategy does choose to question the chair when reaching low enough regret values. The m15n30 dataset did not reach zero regret in 1000 questions.

Figure 4.3 shows the decrease in MMR according to the number of questions asked for various problem sizes. In particular, this shows important differences between some real datasets and the problems generated using IC. In the *Skate* problem, the value $\text{MMR} = 1$ is reached after less than 100 questions, while the IC case of the same size ($m = 14$, $n = 9$) requires more than 200 questions to reach that value. This reasoning also applies to the *Courses* dataset but not to the *T Shirt* dataset. This can be explained by the high degree of similarity in the preference

Figure 4.3: Average MMR (normalized by n) after k questions with Pessimistic strategy for different datasets.



rankings of the *Skate* and the *Courses* problems, which helps reducing the regret faster. For example, in *Skate* the top-2 alternatives are the same for all agents, and 8 out of 9 agents rank the same alternative at position 3. By contrast, in *T Shirt*, the alternatives are evenly distributed in the preference rankings.

Comparison with Two Phases The experiments so far let the strategy free to question either the chair or an agent at each step. One may wonder what is lost in terms of regret by asking different proportions of questions to the chair and the agents. Such restrictions may be useful because of (partial) unavailability of the chair, or because the estimated cognitive costs may differ sensibly.

Table 4.3 shows the MMR value reached in problems of size $m = 10, n = 20$ after 500 questions, using the Two phases strategy, in the “ca” (chair then agents) and in the “ac” (agents then chair) variants. These numbers are to be compared with the MMR value reached after 500 questions with the Pessimistic strategy (displayed in Fig. 4.3), which is 0.7; the Pessimistic strategy asks on average $13 (\pm 13)$ questions to the chair in this setting. The line $q_c = 0$, where no question is asked to the chair, suggest that it is possible to obtain a good-quality recommendation while knowing only that the voting rule is a scoring rule with convex weights, which is our basic hypothesis. However, we observe that asking no questions to the chair does not permit to reach $\text{MMR} = 0$. The strategy, indeed, obtains full knowledge of the profile after an average of 500 questions to the agents but never reaches zero.

Table 4.3: Average MMR in problems of size (10, 20) after 500 questions, among which q_c to the chair.

q_c	2 ph. ca $\pm$ sd	2 ph. ac $\pm$ sd
0	0.6 ± 0.6	0.6 ± 0.6
15	0.5 ± 0.5	0.6 ± 0.6
30	0.3 ± 0.5	0.4 ± 0.5
50	0.1 ± 0.2	0.1 ± 0.3
100	0.4 ± 0.6	0.2 ± 0.5
200	2.1 ± 1.6	2.4 ± 1.2
300	5.8 ± 1.7	6.7 ± 1.5
400	11.3 ± 1.1	11.9 ± 1.2
500	20.0 ± 0.0	20.0 ± 0.0

4.6 Conclusions

In this paper we have considered a social choice setting with partial information about the agent preferences and voting rule. We have proposed the use of minimax regret both as a means of robust winner determination and as a guide to the process of simultaneous elicitation of preferences and voting rule. Our experimental results suggest that regret-based elicitation is effective and allows to quickly reduce worst-case regret significantly. They also show that, in our setting, good quality (low regret) recommendations can be achieved short of having full knowledge of weights or profile.

As part of our contribution, we provide an open-source library that can be found at <https://github.com/oliviercailloux/minimax>, to reproduce our experiments and perform many more.


Some directions for future works include developing new elicitation strategies, considering alternative heuristics; extending the elicitation to voting rules beyond scoring rules; eliciting preferences while restraining to concrete and easy questions.

JAVA PACKAGE TO SUPPORT ELICITATION IN SOCIAL CHOICE

5.1 UML

The Unified Modeling Language (UML) is a standard language for modeling, building and documenting the design of a software system. It is often defined as the language in which the system's blueprint is written. In fact, various aspects of the system can be described by UML diagrams: use cases, interactions between different components, various logical and physical elements. In this chapter, we will describe the logical system of the preference elicitation software developed to accompany the article in Chapter 4. In particular, we will illustrate the classes that constitute the various components of the system and the relationships between them. A class describes the attributes and the properties of a group of objects. An object is an instance of a class. Classes may inherit from other classes their actions and attributes. They can implement abstract methods and add new functionality. Associations between classes represent these behaviors.

A class is represented as follows:



As we can see, the description of a class is divided into three sections, in the first there is its name, in the second its attributes and in the third its methods or behaviors. Attributes and methods can be either private (–) to the class itself, protected (#) to the class and its children, or public (+) to every other class.

In the example presented here there is a class named *Class* with three attributes: a public attribute named *d* of type *Double*, a private attribute named *i* of type *Integer* and a protected attribute named *b* of type *Boolean*. Moreover, *Class* has three methods:

- a method *equals* which takes an element of the class *Class* as input, compares it to the object on which this method is called and returns *True* if they are the same objects, *False*

- otherwise. The implementation of the method itself determine the meaning of similarity.
- a method *toString* which does not take any input and returns the String corresponding to the current object. This value is often used for visualization purposes.
 - a method *hashCode* which does not take any input and produces an Integer value associated to the current object. This value is often used by the method *equals* since same objects must return the same hash value.

We described these three methods because they are basic methods that are automatically inherited by all classes. In fact, every class is an implicit child of a more general *Object* class. If not specified these methods return the default mechanism that can lead to unexpected results. Although specified in our code, we decided to omit these methods from the uml diagrams as a matter of the clarity of the figure.

5.1.1 Preference Representation

Figure 5.1 describes the basic elements of our elicitation software. The two basic elements are represented by the *Alternative* and *Voter* objects. A *Preference* is a list of alternatives in which the order represents the preference order. Multiple alternatives can have the same rank but an alternative cannot be associated to multiple ranks, i.e. it represents a linear order. Methods allow, for example, to know the alternative in a given rank or, conversely, to know the rank of a given alternative. A *StrictPreference* is an extension of the *Preference* object. This means that it inherits all of its behaviors and attributes but can redefine a given mechanism. In particular, *StrictPreference* does not allow two alternatives to have the same rank, i.e. it represents a strict order. A *VoterStrictPreference* links a *Voter* to her *StrictPreference*. A *PrefGraph* is a graph of preferences that can be added or modified. Specifically, each node represents an alternative, and an arc from a node *a* to a node *b* means that the alternative *a* is preferred over the alternative *b*. A *VoterPartialPreference* links a *Voter* to her *PrefGraph*.

5.1.2 Knowledge Representation

Figure 5.2 describes our approach to the representation of knowledge. We defined *PreferenceKnowledge* as an interface, i.e. a completely abstract entity that define the structure that concrete objects must implement. An instance of *PreferenceKnowledge* must implements, among others, methods to return the alternatives, the voters, the profile and the constraints on the weights associated to rank positions. A possible implementation is given by *UpdateablePreferenceKnowledge* which keeps, as attributes, a partial profile and the constraints on weights and allows you to modify them. The *ConstraintsOnWeights* are equations relating the difference between the weights of two consecutive pairs of ranks. It provides method to add constraints and to return an optimal set of weights that satisfies the constraints. Those weights are represented by *PSR-Weights* which is a list of rational numbers associated to the Positional Scoring Rule. Since in our framework we assume the convexity of weights the satisfaction of this condition is automatically checked when creating a new instance of this class. Finally, in the diagram we can see an *Oracle* class. This represents "true" knowledge, which exists but is unknown to us and which we try to elicit by asking questions. It contains the set of alternatives, the preferences of the




Figure 5.1: UML class diagram of preferences representation.

chair—which are the weights of the scoring rule—and the preferences of the voters—which is the profile. During our elicitation procedure we ask questions to this entity.

5.1.3 Questions Representation

Figure 5.3 shows the diagram related to the questions implementation. A *Question* can be either a question to the chair, about their preferences on the voting rule, or a question to one of the voters, about their preferences over the alternatives. This is represented by the use of an enumeration *QuestionType* that admits only two values: *VOTER_QUESTION* and *COMMITTEE_QUESTION*. The method *getType()* of *Question* returns one of those two values. A *Question-Voter* is a comparison query relating two alternatives a and b , where *getPositiveInformation()* returns the *VoterPreferenceInformation* (see Figure 5.4) associated to that specific voter and the alternative a and b . Conversely, *getNegativeInformation()* returns the information regarding the alternative b and a . In other words, the first method returns the information when the answer to the question " $a \succ b$?" is yes, thus a is preferred to b , and the second method that when the answer is no, thus $b \succ a$. A *QuestionCommittee* tries to improve the knowledge about the scoring rule by refining the multiplier λ in the constraint $w_r - w_{r+1} \geq \lambda(w_{r+1} - w_{r+2})$ for some $r \in \{1, \dots, m-2\}$. Given a λ and a rank defined by an elicitation strategy, *getPositiveInformation()* returns the response to the following question " $w_r - w_{r+1} \geq \lambda(w_{r+1} - w_{r+2})$?". The complementary method, *getNegativeInformation()*, returns the information of the opposite question " $w_r - w_{r+1} \leq \lambda(w_{r+1} - w_{r+2})$ ". Positive and negative information is used by the elicitation strategies for the choice of the next question to ask. In particular, the Pessimistic strategy considers, for each potential question, the information in case of positive and negative answer and picks the question that leads to minimal regret in the worst case. This is described in Figure 5.6, in particular in the class diagram of *MmrLottery*.

5.1.4 Preference Information

Figure 5.4 illustrates the diagram related to the information about preferences acquired by asking questions. With a similar representation to the questions themselves, *PreferenceInformation* can be of two types: *CommitteePreferenceInformation* or *VoterPreferenceInformation*. The type of the information depends on the type of the question asked.






Figure 5.3: UML class diagram of questions representation.




Figure 5.4: UML class diagram of elicited knowledge obtained by asking questions.

5.1.5 Regret Computation

Figure 5.5 is the diagram representing the procedure described in Section 4.3. The *Pairwise-MaxRegret* is the building block of the regret computation. Given two alternatives, their ranks and the weights associated to ranks, it computes the regret of choosing the first alternative instead of the second one. This is used by the *Regrets* instance, which calculate the minimal max regret value and it can return all the alternatives associated to such value. A small constant epsilon is introduced to express the granularity of this similarity. In fact, we may want to consider two weights equals if, for example, they are equal up to the third decimal place. *RegretComputer* is the class that handles the proper calling of the methods specified by the classes aforementioned. For example by providing the pairwise maximum regret computation with the worst-case scenario: the worst completion of profile and scoring vector.

5.1.6 Strategies

Figure 5.6 describe our design of elicitation strategies. A *Strategy* is an interface that specify two methods: *setKnowledge* that sets and updates the knowledge collected so far, and *nextQuestion* that returns the next question to ask. Every strategy, i.e. every implementation of *Strategy*, must define those two methods. We created an enumeration for the *StrategyType* based on the different strategies we implemented and tested. In particular, the ones discussed in Section 4.4 are *StrategyByMmr* that, depending on the parameters with which it is invoked, corresponds to *Pessimistic*, *Extended pessimistic* or *Two Phases*, *StrategyElitist* that corresponds to *Elitist* and *StrategyRandom* that corresponds to *Random*. Please see Section 4.4 for a detailed explanation of their behaviors and performances. From a design point of view, each strategy defines the interface methods and it makes use of a *StrategyHelper* for all those functions common to each strategy. For example, to get all the pairs of alternatives whose order is already known and therefore it does not make sense to ask about; or to find the voters to whom we can still ask questions and, eventually, all those possible questions. Finally, *StrategyFactory* is an object that handles the creation of the desired strategy without disclosing to the user the strategy's creation logic.

The java package containing all those classes and also the UML diagrams can be found at <https://github.com/oliviercailloux/minimax>.




Figure 5.5: UML class diagram of regret computation.




Figure 5.6: UML class diagram of strategies.

MAJORITY JUDGMENT WINNER DETERMINATION UNDER INCOMPLETE INFORMATION

MJ is a voting system where voters assign grades to candidates using an ordinal scale. The winner is the candidate with the highest majority-grade —which is the median of the grades received. This method has attracted increasing attention of french associations and political parties which have started to use MJ for internal decisions or local elections. In particular LaPrimaire.org is a french association that uses MJ to choose its candidate for the french presidential election. The vote is conducted in two rounds: in the first one the voters judge five candidates randomly picked; the five candidates with the highest medians pass at the second round as finalists and the voters are asked to judge them. Is the random selection of candidates a good elicitation technique? In this paper we explore the consequences of profile incompleteness and prove that this method can fail to elect the winning candidate of the complete profile. Furthermore, we perform experiments on randomly generated profiles and profiles following the grade distribution of a real voting scenario. We find that the probability of not selecting the winner of the full profile greatly decreases as the number of voters increases and we investigate how much of the grade vector we should know in order for this probability to be low.

6.1 Introduction

MJ is a voting method proposed by Balinski and Laraki [2007, 2011] to elect one out of m candidates based on the judgments of n voters. The latter express their preferences by assigning to each candidate one of the following adjectives: Excellent, Very good, Good, Average, Mediocre, Inadequate, To be rejected. Those adjectives represent a common language whose semantic is assumed to be a shared knowledge among the voters carrying thus an absolute meaning. For each candidate the median of the grades she received is computed, this is called *majority-grade*. The candidate with the highest majority-grade is elected. Ties are broken by considering the majority-grade of first order: one vote associated with the majority-grade of each tied candidates is removed and their medians are recomputed. The candidate with the highest new median is elected. If there is still a tie the process is repeated until a unique winner is found.

In the last few years MJ has being adopted by a progressively larger number of french political parties including: Le Parti Pirate, Génération(s), LaPrimaire.org, France Insoumise and La République en Marche. "Mieux Voter" [MieuxVoter, 2022a] is a french association that promotes the use of MJ as voting method whenever a collective choice has to be selected: public

administration, associations, companies. On their website it is possible to find all the citizens lists –party lists that are not affiliated to any national political party—that used MJ to rank their candidates during the local elections of 2020. In two cases, Bordeaux et Annecy, the candidate selected using MJ was then elected as a mayor.

In particular, LaPrimaire.org [LaPrimaire.org, 2016] is a french political initiative whose goal is to select an independent candidate for the french presidential election using MJ as voting rule. The association Democratech implemented the platform for the first time in 2016 in view of the 2017 presidential elections. The number of voters who participated in the election was 10676 during the first round, and 32685 during the second round.

The procedure that they adopted consists of two rounds. In the first round each voter is asked to express her judgment, using MJ, on five random candidates. At the end of this phase the five candidates with the highest medians are considered the finalists who qualify for the second round. In the second round each voter is asked to express her judgment, using MJ, on all the five finalists. The candidate with the best median at the end of this phase is selected as representative for the presidential election.

In this paper we analyze this elicitation process of voters preferences. In particular, we investigate the consequences of randomness when asking the voters to judge candidates. We evaluate the cost of this procedure—which can be quantified as number of questions per voter—and its fairness for candidates—which reflects the idea that a potential winner should not loose for lack of information.

6.1.1 Related work

One of the earliest uses of the median as an aggregator in voting theory can be identified in the *middlemost* method proposed by Galton [1907a,b]. In particular, in situations where a group of people had to assess a damage, he suggested the median as the only method that does not suffer from over- or under-exaggerations.

We can consider the median grade as the highest level at which a candidate obtains the support of the majority of the voters. Starting for the highest grade δ we check if the majority of the voters assigned at least δ to some alternative. If this is not the case, we descend in the grading scale until such level δ^* is found. The grade δ^* is the median of grades associated to some alternative, and, since it is the first level we stopped at, it corresponds to the best possible median. This method was rediscovered several times and proposed under the name of Bucklin’s rule [Hoag and Hallett, 1926], Majoritarian Compromise [Sertel, 1986, Sertel and Yılmaz, 1999] and q-approval fallback bargaining [Brams and Kilgour, 2001]. Moreover, when the number of grades is equal to two (approve, disapprove) then it is reduced to Approval Voting.

More recently Bassett and Persky [1999] proposed the use of the median as voting rule for elections advocating for its robustness. Several studies on the use of the median have been conducted. In particular, Bassett and Persky [1994] and Gehrlein and Lepelley [2003] study its manipulation. Barthelmy and Monjardet [1981] survey mathematical problems and properties related to the notion of median in the context of cluster analysis and social choice theory. Nehring and Pivato [2022] analyze the median rule in judgment aggregation and they define a weighted median rule that is equivalent to MJ except for the treatment of ties.

However, there is a significant concern that has never been analyzed when considering MJ. Asking voters to provide a grade for each of the candidates can have a high cognitive and communicative cost. In situations where the set of alternatives is very large, voters may not be able to assign a significant and informed grades. Conitzer and Sandholm [2005] studied the complexity of communication of some of the most common voting rules. In these cases, elicitation strategies exist to retrieve the most relevant information. Konczak and Lang [2005] introduced the notion of possible and necessary winners. This paved the way for procedures that attempted to find necessary winners by asking voters the fewest questions possible [Kalech et al., 2011]. Considering scoring rules, Lu and Boutilier [2011a] suggested the use of minimax regret as a guide to the elicitation procedure. Bachrach et al. [2010] proposed a probabilistic approach. Several authors studied the complexity of determining when to stop the elicitation process for some of the most common voting rules [Conitzer and Sandholm, 2002, Walsh, 2008]. However, to the best of our knowledge, there are no works on preferences elicitation considering MJ.

Very recently, Varlout and Laraki [2022] considered a version of MJ under uncertainties in order to study its strategyproof properties. Their premise, however, is based on the fact that voters who are uncertain about the degree to assign to an alternative would instead assign to it a probability distribution. In other words, if a voter does not know whether to assign good or excellent to a certain alternative, then she may assign to it, for example, good with $\frac{1}{3}$ of the probability and excellent with $\frac{2}{3}$ of the probability. This approach requires the submission of even more information and is very far from our idea of incomplete knowledge.

6.2 Notation

Consider a finite set $N = \{i_1, \dots, i_n\}$ of voters (or judges) and a finite set $A = \{j_1, \dots, j_m\}$ of alternatives (or competitors). A *common language* $\Delta = \{\delta_1, \delta_2, \dots\}$ is a set of strictly ordered grades, and the notation $\delta_1 \geq \delta_2$ indicates that δ_1 is a better or equivalent grade than δ_2 . A profile $P : A \times N \rightarrow \Delta$ is a m by n matrix of grades. Given any $N' \subseteq N$, the operator $\rho : \Delta^{N'} \rightarrow \Delta^{|N'|}$ defines an ordering function that given a vector of grades P_j returns the vector ordered by decreasing grades.

Consider a set of alternatives $S \subseteq A$, we denote by $P^S \in \Delta^{S \times N}$ a restriction of the profile P to only the alternatives in S , $P^S \subseteq P$. Note that when $S = A$ then $P^S = P$.

We let $f_{maj} : \Delta^N \rightarrow \Delta$, denote the *majority-grade* function that associates to a vector of grades $\emptyset \neq v \in \Delta^{N'}$, $N' \subseteq N$ its median grade value: $f_{maj}(v) = \rho(v)_{\lfloor \frac{|v|}{2} \rfloor + 1}$. Note that in case $|v|$ is even, two medians could be used, but, as in Balinski and Laraki [2011] definition, the lower grade is picked.

Given any $S \subseteq A$, the winner function $F_{maj}^S : \Delta^{S \times N} \rightarrow A$ is a function that given $|S|$ alternatives and their median grade, selects the alternative with the highest median grade as winner. We can define it as $F_{maj}^S(P^S) = \operatorname{argmax}_{j \in S} f_{maj}(P_j^S)$ assuming it is a singleton. For brevity, we will write F_{maj} when considering $S = A$. To avoid adding further complexity to the notation, we describe only informally what happens in case multiple alternatives are tied for the highest median grade h , thus, when $\operatorname{argmax}_{j \in S} f_{maj}(P_j^S)$ is not a singleton. In this case, ties are broken by removing one h grade from the vectors of grades of each tied alternative, recomputing the new median grade and repeating the process until one unique winner is found or there are

no more grades to remove. When n is odd, this is equivalent to take the next element after the median, i.e. the one at index $(\lfloor \frac{n}{2} \rfloor + 1) + 1$. If there is still a tie we then look at the previous element before the median, i.e. the one at index $(\lfloor \frac{n}{2} \rfloor + 1) - 1$, and keep alternating until the tie is broken or there are no more elements in the vector. When n is even the element before the median, i.e. the one at index $\lfloor \frac{n}{2} \rfloor$, is taken first and, if there is still a tie, then the element after the median is considered, i.e. the one at index $\lfloor \frac{n}{2} \rfloor + 2$, and so on. If after applying the mechanism there are still ties we break them using an arbitrary ordering defined on all alternatives, e.g. lexicographical order.

Similarly to the winner function, given any $x \in \llbracket 1, m \rrbracket$, we can define a more general *selection function* $F_x^S : \Delta^{S \times N} \rightarrow \mathcal{P}^*(A)$, where $\mathcal{P}^*(A)$ is the powerset of A excluding the empty set, that selects exactly the x alternatives with the x highest median grades. Thus, $F_x^S(P^S) = F_{maj}^S(P^S) \cup F_{maj}^{S'}(P^{S'}) \cup \dots \cup F_{maj}^{S^{k-1}}(P^{S^{k-1}})$ where $S' = S \setminus F_{maj}^S(P^S)$, $S'' = S \setminus F_{maj}^S(P^S) \setminus F_{maj}^{S'}(P^{S'})$ etc.

6.2.1 Incomplete knowledge

In order to analyze the elicitation procedure used by LaPrimaire.org, we need to adapt the notation just described to incomplete profiles. Let $\bar{P}$ be our knowledge about the profile P . The voters have full knowledge of their own judgments but we ignore them; our goal is to elicit them by questioning the voters starting from zero knowledge. We introduce an additional grade $\bar{\delta}$, and, given a language Δ , $\bar{\delta} < \delta, \forall \delta \in \Delta$. The common language in the incomplete knowledge setting is then $\bar{\Delta} = \Delta \cup \bar{\delta}$. Voters cannot use this grade to express their judgment over an alternative and it does not count in the computation of the median grade as we will explain formally. We refer to the grades in Δ , the ones used by voters, as "defined" grades.

Let $\bar{P} \in \bar{\Delta}^{A \times N}$ denote an *incomplete profile*, i.e. a matrix $m \times n$ of grades some of which are "undefined". The starting knowledge is represented by a matrix $m \times n$ of $\bar{\delta}$ grades. Given $k \in \mathbb{N}$, we define with $K_i \subseteq A$ a set of k alternatives that we ask the voter $i \in N$ to evaluate. After having asked every voter in N to judge k candidates, we obtain an incomplete profile $\bar{P}^k$, that is a matrix $m \times n$ of grades of which kn are "defined". Note that when $k = m$, the resulting profile $\bar{P}^k$ corresponds to the complete profile P . Let $C(\bar{P})$ be the set of all completions of $\bar{P}$ obtained by substituting all $\bar{\delta}$ grades with "defined" ones, note that $P \in C(\bar{P})$.

Let $g : \bar{\Delta}^N \rightarrow \bigcup_{N' \subseteq N} \Delta^{N'}$ be the function that given an incomplete vector of grades $q \in \bar{\Delta}^N$, thus $q \subseteq N \times \bar{\Delta}$, returns the vector composed only of the "defined" grades.

The *majority-grade* for incomplete profile $\bar{f}_{maj} : \bar{\Delta}^N \rightarrow \Delta$ corresponds to f_{maj} that only considers the "defined" grades in the computation of the median. Consider an incomplete grade vector $\bar{P}_j$ and let $P'_j = g(\bar{P}_j)$ be the partial vector of $\bar{P}_j$ of only "defined" grades, then

$$\bar{f}_{maj}(\bar{P}_j) = \begin{cases} f_{maj}(P'_j) = \rho(P'_j)_{\lfloor \frac{|P'_j|}{2} \rfloor + 1} & \text{for } P'_j \neq \emptyset \\ \bar{\delta} & \text{otherwise} \end{cases}$$

In a similar fashion, we can define $\bar{F}_{maj}^S(\bar{P}) = \operatorname{argmax}_{j \in S} \bar{f}_{maj}(\bar{P}_j^S)$ and, given a value $x \in \llbracket 1, m \rrbracket$, the selection function of the x best alternatives is $\bar{F}_x^S(\bar{P}^S) = \bar{F}_{maj}^S(\bar{P}^S) \cup \bar{F}_{maj}^{S'}(\bar{P}^{S'}) \cup$

$\dots \cup \bar{F}_{maj}^{S^{x-1}}(\bar{P}^{S^{x-1}})$ where $S' = S \setminus \bar{F}_{maj}^S(\bar{P}^S)$, $S'' = S \setminus \bar{F}_{maj}^S(\bar{P}^S) \setminus \bar{F}_{maj}^{S'}(\bar{P}^{S'})$ etc.

To summarize the elicitation process we want to investigate, starting with zero knowledge, each voter is asked to evaluate k candidates. Given the partial information at our disposal, we are able to define the "known" median grade for each alternative by applying $\bar{f}_{maj}(\bar{P}_j^k)$, $\forall j \in A$. We can then use the *selection function* $\bar{F}_k$ to select the k alternatives with the highest "known" median grades. We denote this set of alternatives with $K = \bar{F}_k \subseteq A$. The alternatives in K will be presented to all voters, who must provide a grade for each of them, thus obtaining a restriction of the complete profile to this subset of k alternatives. From here, we can use $F_{maj}^K(P^K)$ to select the winning alternative of the restricted profile P^K .

Remark 6.1. *Because we are interested into investigating MJ, we are going to use an alphabet with the same size of the one proposed by Balinski and Laraki [2011] which is composed of the following adjectives: To be Rejected, Inadequate, Mediocre, Average, Good, Very Good, Excellent. For brevity we are gonna rename those adjectives respectively from δ_1 , corresponding to To be Rejected, to δ_7 , corresponding to Excellent. Therefore, $\Delta = \{\delta_1, \delta_2, \delta_3, \delta_4, \delta_5, \delta_6, \delta_7\}$*

6.3 Reasoning on incompleteness

Let us now consider the effects of incomplete knowledge on the selection of a winner. In particular, we want to prove that if we consider an incomplete profile and take the k alternatives with the highest median grade, with any k being between 1 and $m - 1$, it is possible to miss the winner of one of its completion. This is important because in the elicitation process implemented by LaPrimaire.org, only the full grade vectors of the k alternatives with the highest median grade are considered. Clearly, if the winning alternative in the complete profile is in this set $K \subset A$, then she will also be the winner of the incomplete profile. This is because the voters will be asked to provide a grade for all those k alternatives, providing thus a restriction of the complete profile to only the alternatives in K . If a candidate is a winner for a complete profile, then she is also a winner for any of its restrictions that include her. Note that this is not true in case of multiple winners where ties that cannot be broken by tie-breaking procedures and for which arbitrary ordering is, therefore, necessary. In what follows we assume the absence of such ties. We want to show that it is possible for the winning alternative of the complete profile not to be included in this subset K of alternatives.

Theorem 6.1. *Given a set A of alternatives $m \geq 2$, a set of voters N and a value $k \in \llbracket 1, m - 1 \rrbracket$, there exist a complete profile P and an incomplete profile $\bar{P}$ such that $P \in C(\bar{P})$ and $F_{maj}(P) \notin \bar{F}_k(\bar{P})$.*

Proof. Consider the following complete profile P

$$\begin{array}{ccccc}
 & i_1 & i_2 & \dots & i_n \\
 j_1 & \delta_7 & \delta_7 & \dots & \delta_7 \\
 j_2 & \delta_6 & \delta_6 & \dots & \delta_6 \\
 \vdots & \delta_6 & \delta_6 & \dots & \delta_6 \\
 j_m & \delta_6 & \delta_6 & \dots & \delta_6
 \end{array} ,$$

where all voters judge all alternatives δ_6 except for the alternative j_1 which is judged δ_7 by everyone. The median grade of the alternative j_1 is $f_{maj}(P_{j_1}) = \delta_7$ and the one of all the other alternatives is $f_{maj}(P_x) = \delta_6, \forall x \in A \setminus \{j_1\}$. Thus, the winner in the profile P is $F_{maj}(P) = \{j_1\}$. Consider now the following incomplete profile $\bar{P}$:

	i_1	i_2	$\dots$	i_n
j_1	$\bar{\delta}$	$\bar{\delta}$	$\dots$	$\bar{\delta}$
j_2	δ_6	δ_6	$\dots$	δ_6
$\vdots$	δ_6	δ_6	$\dots$	δ_6
j_m	δ_6	δ_6	$\dots$	δ_6

The median grade of j_1 under the profile $\bar{P}$ is $f_{maj}(\bar{P}_{j_1}) = \bar{\delta}$ and the one of all the other alternatives is $f_{maj}(\bar{P}_x) = \delta_6, \forall x \in A \setminus \{j_1\}$. Because $\bar{\delta} < \delta_6$, $\bar{F}_k(\bar{P})$ will be any subset of k elements of $\{j_2, j_3, \dots, j_m\}$, for any $k \in \llbracket 1, m-1 \rrbracket$. Since all those alternatives have the same grade vectors, some orderings can be used to define the members of $\bar{F}_k(\bar{P})$. Thus, $j_1 \notin \bar{F}_k(\bar{P})$, and $F_{maj}(P) \subsetneq \bar{F}_k(\bar{P})$, for any $k \in \llbracket 1, m-1 \rrbracket$. $\square$

From Theorem 6.1, considering $k = 1$ and recalling that $\bar{F}_{k=1}(\bar{P}) = \bar{F}_{maj}(\bar{P})$ we can conclude that there exist an incomplete profile and one of its completion that do not have the same winner.

Remark 6.2. *Given a set of alternatives A and a set of voters N , there exist a complete profile P and an incomplete profile $\bar{P}$ such that $P \in C(\bar{P})$ and $F_{maj}(P) \neq \bar{F}_{maj}(\bar{P})$.*

We have shown that it is possible to find a complete profile P whose winner is not in the set K of alternatives with the highest medians for an incomplete profile $\bar{P}$, with $P \in C(\bar{P})$. If we look at the procedure used by LaPrimaire.org, however, the situation is slightly more complex. We are not considering just any incomplete profile, but a k value is set from the beginning and the incomplete profile is formed by asking each voter to rate k random chosen candidates. They set $k = 5$. The k alternatives with the highest medians on this incomplete profile form K .

Note that, the example profile $\bar{P}$ considered in the proof of Theorem 6.1 could be the result of a random process of questioning the voters about $k = m - 1$ alternatives, and they never got the chance to grade the alternative j_1 .

We recall that $\bar{P}^k$ is an incomplete profile where every voter judges $k \in \llbracket 1, m-1 \rrbracket$ alternatives. If we consider the matrix $m \times n$ of grades, then the columns, i.e. the vectors of grades expressed by each voter, are composed of only k "defined" grades.

Theorem 6.2. *Given a set A of alternatives $m \geq 2$, a set of voters N and a value $k \in \llbracket 1, m-1 \rrbracket$, there exist a complete profile P and an incomplete profile $\bar{P}^k$ such that $P \in C(\bar{P}^k)$ and $F_{maj}(P) \not\subseteq \bar{F}_k(\bar{P}^k)$.*

Proof. Consider the complete profile P defined in the proof of Theorem 6.1, we will show how to construct an incomplete profile for any value of $k \in \llbracket 1, m-1 \rrbracket$ such that the statement is true. Let us call $K = \bar{F}_k(\bar{P}^k)$. Note that we can select any complete profile where there is one alternative j graded the highest by everyone, and all the other alternatives are considered worst by everyone.

When $k = m - 1$, the proof follows from the proof of Theorem 6.1. The set of the $m - 1$ alternatives with the highest medians is $K = \{j_2, j_3, \dots, j_m\}$. Thus, $F_{maj}(P) = j_1 \notin K$.

For any other $k < m - 1$ we can follow the same reasoning by building $\bar{P}^k$ making sure that we never ask anyone the grade of j_1 . The grade vector of j_1 is $\bar{P}_{j_1}^k = (\bar{\delta}, \dots, \bar{\delta})$ and $\bar{f}_{maj}(\bar{P}_{j_1}^k) = \bar{\delta}$. Two situations are now possible, either all voters grade the same k alternatives and their associated grade vectors are complete; or each voter evaluates a set of k different alternatives from $A \setminus j_1$. In the first case we have exactly k alternatives whose majority grade is defined and they would form the set K . In the second case there are at least k alternatives with at least one "defined" grade and, thus, with a defined majority grade. Therefore $\bar{f}_{maj}(\bar{P}_{j_1}^k) = \bar{\delta}$ will be lower than at least k others and it will not be selected in K . $\square$

From Theorem 6.2, considering $k = 1$ we have that:

Remark 6.3. *Given a set A of alternatives $m \geq 2$, a set of voters N and a value $k \in \llbracket 1, m - 1 \rrbracket$, there exist a complete profile P and an incomplete profile $\bar{P}^k$ such that $P \in C(\bar{P}^k)$ and $F_{maj}(P) \neq \bar{F}_{maj}(\bar{P}^k)$.*

In the proofs of Theorems 6.1 and 6.2 we showed the existence of an incomplete profile whose $k \in \llbracket 1, m - 1 \rrbracket$ best alternatives did not include the winner of one of its completions. To do so, we assumed that no questions were asked about an alternative j winner in the complete profile P . Although the theoretical result remains, how likely is this to happen in practice?

Proposition 6.1. *Given a set A of m alternatives, a set N of n voters, a value $k \in \llbracket 1, m - 1 \rrbracket$ and considering an elicitation strategy in which each voter is asked independently to evaluate k alternatives picked equiprobably, the probability that an alternative $j \in A$ is never asked about is $(1 - \frac{k}{m})^n$.*

Proof. To prove this let us first consider e_i the event of asking a voter $i \in N$ to grade the alternative j in k questions, and let us compute the probability $\mathcal{P}$ of this event to happen. There are $\binom{m-1}{k-1}$ ways to select k alternatives among m that include the alternative j . This is because once we select the alternative j we can still pick $k - 1$ alternatives to ask the voter among the remaining $m - 1$. Moreover, there are $\binom{m}{k}$ ways to select k alternatives among m with no constraints. Thus, the probability of asking a voter to grade the alternative j in k questions is:

$$\mathcal{P}(e_i) = \frac{\binom{m-1}{k-1}}{\binom{m}{k}} = \frac{\frac{(m-1)!}{(k-1)!(m-1-k+1)!}}{\frac{m!}{k!(m-k)!}} = \frac{(m-1)!}{(k-1)!(m-k)!} \cdot \frac{k(k-1)!(m-k)!}{m(m-1)!} = \frac{k}{m}.$$

The probability that one voter is never questioned about an alternative is then:

$$\mathcal{P}(\bar{e}_i) = 1 - \mathcal{P}(e_i) = 1 - \frac{k}{m}.$$

Because questionings different voters are independent events, we can express the probability that no voter is asked to grade an alternative j as the product of the individual probabilities. Denoting this event with e_j we have that:

$$\mathcal{P}(e_j) = \mathcal{P}(e_i)^n = \left(1 - \frac{k}{m}\right)^n.$$

□

Remark 6.4. After asking each of the n voters to grade $k \in \llbracket 1, m-1 \rrbracket$ of the m alternatives, the average size of the grade vectors is $q = n \cdot \frac{k}{m}$.

This observation comes from the proof of Proposition 6.1. Given an alternative $j \in A$ and a voter $i \in N$ the probability that i is asked about j , thus that $\bar{P}_j^k(i)$ is defined, is $\frac{k}{m}$. This gives us an approximation on the number of elements of the grade vector, because for each of the n elements, we have $\frac{k}{m}$ probability that the element is defined.

The value found in Proposition 6.1 is very low to occur in real examples, if, for example, we consider $k = 3$, $n = m = 10$ the probability that none is asked to grade the winner of the complete profile is 0.7^{10} (which approximated is 0.0282%). However, this is also very restrictive. For an alternative j , winner of the complete profile, not to be the winner of the incomplete profile it suffices for her median to be smaller than the one of other k alternatives. We do not need the vector of the alternative considered to be empty, just to be misrepresented enough to have a median lower than k others.

Example 6.1. Consider the following profile P :

	i_1	$\dots$	$i_{\lfloor \frac{n}{2} \rfloor + 1}$	$i_{\lfloor \frac{n}{2} \rfloor + 2}$	$\dots$	i_n
j_1	δ_7	$\dots$	δ_7	δ_5	$\dots$	δ_5
j_2	δ_6	$\dots$	δ_6	δ_6	$\dots$	δ_6
$\vdots$	δ_6	$\dots$	δ_6	δ_6	$\dots$	δ_6
j_m	δ_6	$\dots$	δ_6	δ_6	$\dots$	δ_6

Exactly $\lfloor \frac{n}{2} \rfloor + 1$ voters grade the alternative j_1 with the highest grade δ_7 and the rest of the voters grade her δ_5 . All voters grade any other alternative δ_6 . The median grade of j_1 is $f_{maj}(P_{j_1}) = \delta_7$ and the one of all the other alternatives is $f_{maj}(P_x) = \delta_6, \forall x \in A \setminus \{j_1\}$. Thus, j_1 is the winner of the profile P . For j_1 not to be the winner of any $\bar{P}^k$ such that $P \in C(\bar{P}^k)$, it must be that $\bar{f}_{maj}(j_1) = \delta_5$. This happens when the incomplete grade vector of j_1 has a number of δ_5 grades at least as great as the number of δ_7 grades. Let us denote with $q = |\bar{P}_{j_1}^k|$ the size of the incomplete grade vector of j_1 . Assume, for the sake of simplicity, that n and q are both even, although this reasoning can easily be adapted in case they are odd. An incomplete grade vector whose median is δ_5 , could be formed by $q/2$ δ_7 grades and $q/2$ δ_5 grades. To compute the probability of this vector to be formed, we must consider that there are $\binom{\frac{n}{2}+1}{\frac{q}{2}}$ ways of taking $q/2$ of the $\frac{n}{2} + 1$ δ_7 grades of the complete vector. Moreover, there are $\binom{\frac{n}{2}-1}{\frac{q}{2}+i}$ ways of taking $q/2$ of the $\frac{n}{2} - 1$ δ_5 grades. If we consider no restrictions there are $\binom{n}{q}$ ways of taking a vector of q grades out of a vector of n elements. However, this is not the only vector whose median is δ_5 . In fact, we could have $q/2 - 1$ δ_7 grades and $q/2 + 1$ δ_5 grades. If we iterate this reasoning on all the possible division of q we have that:

$$\mathcal{P} = \sum_{i=0}^{q/2} \frac{\binom{\frac{n}{2}+1}{\frac{q}{2}-i} \cdot \binom{\frac{n}{2}-1}{\frac{q}{2}+i}}{\binom{n}{q}}$$

Because we know from Remark 6.4 that the average size of any grade vector is $q = n \cdot \frac{k}{m}$, we experimentally computed this probability with different values of n , m and k . Except when k is very close to m , and then we have almost the whole vector, we have that this probability is $\mathcal{P} \approx \frac{1}{2}$.

Again, the profile considered here is rather peculiar and not of practical use. In the next section we investigate experimentally different profile distributions.

6.4 Experimental results

In our experiments we assume the existence of an Oracle who knows the complete profile, to whom we ask questions to elicit voters preferences. Specifically, we want to investigate what is the probability that starting from zero knowledge and asking k questions to each voter the winner of the complete profile is not within the k alternatives with the best medians of the incomplete profile. One of the first observations, also from the previous examples, is that given a number n of voters, this probability is highly dependent on how well each grade vector is represented. Because the average size of each vector is equal to $q = n \cdot \frac{k}{m}$, we are particularly interested in how the probability of "missing the winner" varies as a function of k/m . This, in fact, can be interpreted as the percentage of representation of the whole grade vector of n elements. By knowing this relation we therefore know the percentage of each vector, on average, that must be known in order to have a given probability of missing the winner. From this, we can infer important information, for example how many questions, k , to ask each voter so that the probability of missing the winner is less than a given threshold. Another measure we are interested in follows somewhat the opposite reasoning, given a number of questions k that I want to ask each voter, which can be fixed for reasons of cognitive and computational complexity, how large the electorate n must be for the vectors to be well represented?

In what follows we consider two different scenarios. In the first we assume that each grade vector is randomly selected from a uniform distribution, somewhat like in the *Impartial Culture* model used to evaluate ranking methods. In the second we consider that the grade vectors associated with the alternatives are distributed following the results of a real survey carried out during the French presidential election in 2022. The code to run the experiments is available at <https://github.com/xoxor/elicitationMJ>.

6.4.1 Profile uniformly distributed

For the first experiment we consider a profile where each voter has equal probability of assigning a grade $\delta \in \Delta$ to each alternative. Because we observed that with a low number of voters the probability of missing the winner $\mathcal{P}$ is very high, and with an high number of voters it is very low, we arbitrarily take $n = 100$ and $m = 50$. Moreover, in order to avoid the use of tie breaking mechanisms, that can be non resolute in case of ties that cannot be broken, we assume a profile P where there is an alternative j_1 whose median is always δ_7 and there are no other alternatives with the same median. Given such a profile P , we ask k questions to each voter for all $k \in \{1, m/2\}$. We consider this interval because we observed that already with $k = m/2$ the probability $\mathcal{P}$ is very close to 0. After the elicitation procedure is complete and the incomplete

Table 6.1: Average probability of missing the winner $\mathcal{P}$, given $n = 100, m = 50$ considering 1000 complete profiles uniformly distributed and 1000 incomplete profiles for each of them.

k	avg $\mathcal{P}$	$\pm$	sd
1	0.64	$\pm$	2.43×10^{-4}
2	0.52	$\pm$	7.75×10^{-4}
3	0.39	$\pm$	2.40×10^{-3}
4	0.28	$\pm$	3.03×10^{-3}
5	0.19	$\pm$	3.10×10^{-3}
6	0.14	$\pm$	2.77×10^{-3}
7	0.11	$\pm$	2.60×10^{-3}
8	0.09	$\pm$	2.21×10^{-3}
9	0.06	$\pm$	1.60×10^{-3}
10	0.05	$\pm$	1.01×10^{-3}
11	0.03	$\pm$	5.85×10^{-4}
12	0.02	$\pm$	2.33×10^{-4}
13	8.68×10^{-3}	$\pm$	1.01×10^{-4}


grade vectors are found, we compute the median for each of those vectors. The k alternatives with the highest medians form the set K and we check if j_1 is in this set. If yes, then we know it will be the winner in the next phase when the complete preferences are asked to each voters. Otherwise, j_1 will not be the winner and we count this incomplete profile $\bar{P}^k$ for our $\mathcal{P}$. In fact, given the profile P we repeat this operation 1000 times, and $\mathcal{P}$ will be the number of profiles $\bar{P}^k$ for which j_1 is not a winner, over the total number of incomplete profiles considered which is 1000. Moreover, we repeat this operations 1000 times. Thus, we generate 1000 complete profiles P and for each we consider 1000 incomplete profiles. At the end we obtain 1000 probabilities for which we consider average and standard deviation.

Table 6.1 shows the average probability for each value k with the standard deviation over the 1000 runs. We varied k from 1 to 25 but omitted the results after $k = 13$ because they are too small to be relevant. The complete results can be found at <https://github.com/xoxor/elicitationMJ> and they are displayed on Figure 6.1. In fact, if we take the first line of Table 6.1, we have that when we ask $k = 1$ question to each voter we have 64% of probability not to find j_1 in the set K . This happens when the incomplete vectors are, on average, $k/m = 1/50$ (2%) of the complete vector. If we take $k = 5$, we know on average 1/10 of each grade vector, and the probability drops at 0.19%. Figure 6.1 shows this correlation.

6.4.2 Profile real scenario

In the next experiment we want to investigate a different problem. Given a number of alternatives m , we assume that we want to ask a limited number of questions k because asking questions is costly. How large the electorate n must be for an incomplete grade vector of $n \cdot k/m$ elements to be sufficiently representative of a complete vector? Because randomly constructed profiles

Figure 6.1: Average probability of missing the winner under uniform distribution of preferences, given $n = 100$, $m = 50$ and $k \in \llbracket 1, 25 \rrbracket$ for 1000 batches of 1000 runs.




are often seen as unrepresentative of reality we decided to construct the grade vectors for each alternative following the grade distribution obtained in a real experiment conducted during the 2022 French presidential election. Given 12 political candidates, voters were asked to give a grade from *To be rejected*, corresponding to our δ_1 , to *Excellent*, corresponding to δ_7 , to each of the alternatives. The winner is determined using MJ.

The same experiment was carried out independently by the association *Mieux Voter* and the experimental study *Un Autre Vote* (UAV). The results, that can be found respectively at MieuxVoter [2022b] and Delemazure and Bouveret [2022], are extremely similar. We chose to use the distribution obtained from the candidates in the UAV experiment, see Figure 6.2, for no particular reason. The code can be easily adapted to the results obtained from MieuxVoter. Note that in the UAV study, some voters were asked to judge the alternatives on a 5-grade scale and others to provide judgments on a 7-grade scale. We consider the profile obtained in the second case. Also, if you have followed the French presidential election you may notice that the distribution shown in Figure 6.2 is not an accurate portrait of the country’s political situation. This is because the study did not capture a representative sample of the population. In their analysis of the results, Delemazure and Bouveret [2022] try to correct this bias while acknowledging, nonetheless, that a discrepancy may still remain. However, this is not relevant to us for the purpose of our work. We take a voting profile and consider what happens when we do not know all the grades. We are not interested in the identity of the candidates nor in the purpose of the vote, taking the profile out of context it could represent voters opinions on a set of restaurants or movies. Our study focuses on considering a distribution of preferences in a realistic profile, a profile in which grades are not uniformly distributed but some alternatives are generally preferred over others.

The grade vector associated to the candidate Jean-Luc Mélenchon is composed for: 29% of Excellent grades, 22% of Very good, 16% of Good, 11% of Average, 7% of Mediocre, 5%

Figure 6.2: Result obtained from UAV experiment by asking $n = 1147$ voters to judge $m = 12$ alternatives.



of Inadequate and for 9% of To be rejected. Because it is the winner of the complete profile this alternative will be our j_0 and its complete grade vector will be formed following the same proportions: 29% of δ_7 , 22% of δ_6 , 16% of δ_5 , 11% of δ_4 , 7% of δ_3 , 5% of δ_2 and for 9% of δ_1 . We follow the same reasoning for the other grade vectors in order to have the complete profile.

Given $m = 12$ alternatives we fix a value of n and see what is the probability that j_0 is not in K for k from 1 to 5. Using the same reasoning as in the previous experiment we repeat this a total of 1,000,000 times and we average over the probability. After that we increase n and repeat the same experiment. We note that the probability decreases dramatically as n increases so, although we varied n from 10 to 1500, in Figure 6.3 we show the results only for a limited number of n values. Table 6.2 shows the data for the same values of n , the complete dataset can be found at <https://github.com/xoxor/elicitationMJ>. We see that when we have $n = 50$ voters by asking $k = 3$ questions per voter already we have an almost zero probability of not electing j_0 . This result is obtained by asking 2 questions per voter when $n = 100$ and 1 question when $n = 500$.

These results suggest that in small profiles with few voters this method can pose a serious risk of not electing the "correct" alternative unless many questions are asked. This risk, however, is almost absent for elections in which a large number of voters are involved. Although we cannot use these findings to reach a definite conclusion, working in this direction might help in the formulation of a more comprehensive result.

6.5 Conclusions

In this paper we analyzed a procedure proposed by the association *Mieux Voter* and used in a real life application by *LaPrimaire.org* for eliciting judgments of a set of voters over a set of alternatives. The procedure consists of randomly asking each voter to judge k candidates on the m available. The k alternatives with the best medians are then presented to all voters in a

Figure 6.3: Average probability of missing the winner using a real case distribution of preferences, given $m = 12$ for $n \in \{10, 25, 50, 100, 250, 500\}$ and $k \in \llbracket 1, 5 \rrbracket$.




Table 6.2: Average probability of missing the winner using a real case distribution of preferences, given $m = 12$ for different values of k and n .

k	avg $\mathcal{P}_{n=10}$	avg $\mathcal{P}_{n=25}$	avg $\mathcal{P}_{n=50}$	avg $\mathcal{P}_{n=100}$	avg $\mathcal{P}_{n=250}$	avg $\mathcal{P}_{n=500}$
1	0.71	0.50	0.34	0.17	0.03	0.00
2	0.50	0.22	0.07	0.01	0.00	0.00
3	0.34	0.10	0.01	0.00	0.00	0.00
4	0.18	0.04	0.00	0.00	0.00	0.00
5	0.07	0.01	0.00	0.00	0.00	0.00

second round of voting, and the best alternative is selected by MJ as the winner. We have proven that given a complete profile it is possible to find an incomplete profile (such that a possible completion is equal to the considered profile) with a different winner from the complete profile. We also showed that the probability of this happening depends on the profile considered.

Moreover, we have performed experiments on profiles in which each grade vector has a randomly generated distribution of grade. Fixing the number of voters, we analyzed how much of the grade vector we must know in order for the probability missing the winner to be low. We have found that when knowing only the 20% of the vector (thus after having asked 10 questions) the probability of missing the winner is close to zero. We have also considered a profile created from the grade distribution of a real voting scenario. We have found that the probability of not selecting the winner of the complete profile greatly decreases as the number of voters increases, making this method suitable for elections with large numbers of voters but less appropriate for small profiles.

Future work on this topic could investigate the manipulability of this elicitation process. Especially on the side of the "analyst" who is in charge of questioning voters and how she can potentially manipulate the outcome by asking particular questions to specific groups of voters. For example, if she knows that voters in a given electoral district tend to vote predominantly for a given party she may decide never to ask those voters about that party's candidate (or vice-versa depending on the party she favors).

Another step might involve the explainability of this process to voters. In fact, voters may find difficult to believe that judging only one person gives a good approximation of the result that would be obtained by asking for the full profile. Also, they may not appreciate judging a random candidate and not their favorite.

The code to reproduce the experiments in this article and the files with the reported data are available at <https://github.com/xoxor/elicitationMJ>.

COMPROMISING AS AN EQUAL-LOSS PRINCIPLE

A social choice rule aggregates the preferences of a group of individuals over a set of alternatives into a collective choice. The literature admits several social choice rules whose recommendations are supposed to reflect a compromise among individuals. We observe that all these compromise rules can be better described as procedural compromises, i.e., they impose over individuals a willingness to compromise but they do not ensure an outcome where everyone has effectively compromised. We revisit the concept of a compromise in a collective choice environment with at least three individuals having strict preferences over a finite set of alternatives. Referring to a large class of spread measures, we view the concept of compromise from an equal loss perspective, favoring an outcome where every voter concedes as equally as possible. As such, being a compromise may fail Pareto efficiency, which we ensure by asking voters to concede as equally as possible among the Pareto efficient alternatives. We show that Condorcet consistent rules, scoring rules (except antiplurality) and Brams-Kilgour compromises (except fallback bargaining) all fail to ascertain an outcome which is a compromise. A slight restriction on acceptable spread measures suffices to extend the negative result to antiplurality and fallback bargaining. This failure also prevails for social choice problems with two individuals: all well-known two-person social choice rules of the literature, namely, fallback bargaining, Pareto and veto rules, short listing and veto rank, fail to pick ex-post compromises. We conclude that there is a need to propose and study rules that satisfy this equal loss, or outcome oriented, notion of a compromise.

7.1 Introduction

In a typical social choice problem, several individuals express their preferences over a set of alternatives and one shall be picked as the collective outcome. Although the literature admits several SCRs with different properties, there is a common understanding that collective choices must reflect “compromises”. One of the first to explicitly refer to a SCR as a compromise is Sertel [1986] introducing the *majoritarian compromise*. This SCR, further analyzed by Sertel and Yılmaz [1999], is a rediscovery of a method suggested by James W. Bucklin at the beginning of the 20th century [Erdélyi et al., 2015]. Starting from everyone’s ideal alternative, it falls back to the voters’ second, third and more generally k -th best, until one of the alternatives considered appears among the first k best for a majority. Brams and Kilgour [2001] generalize this concept and introduce a class of SCRs called *q-approval fall-back bargaining*, where q is the required quantity of support that can vary from a single voter up to unanimity. Different choices of q lead

to different SCRs. Considering n voters, the choice of $q = 1$ corresponds to the plurality rule [BN: not true], $q = \lceil \frac{n}{2} \rceil$ is closely related, but not identical, to the majoritarian compromise and $q = n$ represents a bargaining procedure called *fall-back bargaining*, which has been further analyzed by Kibris and Sertel [2007] and Congar and Merlin [2012]. We will refer to these rules as Brams and Kilgour (BK) compromises with threshold q .

As Özkal-Sanver and Sanver [2004] discuss, the concept of compromising is mostly understood as the trade off between the number of voters supporting an alternative (i.e., the quantity of support) and the distance of that alternative from the supporters' ideal alternative (i.e., the quality of support). This trade off, which is explicit for q -approval fall-back bargaining, is also the basis for several other SCRs such as the *median voting rule* proposed by Bassett and Persky [1999] and further analyzed by Gehrlein and Lepelley [2003] or the *Condorcet practical method* described by Nurmi [1999]. Merlin et al. [2019] identify and analyze a large class of *compromise rules* that explore this trade off.

One can argue that a collective choice *per se* implies a compromise. After all, except extreme cases such as dictatorships, a SCRs operates on the principle that all voters could fall back from their ideal position. Whether all voters effectively do fall back and whether this fall is “equal among voters” is the subject of our analysis. In what follows, we will present examples where they do not, which could be considered counter to the spirit of compromising.¹

Consider the following example.

Example 7.1. Let N be a set of $n \geq 3$ voters and A a set of alternatives. $\mathcal{L}(A)$ represents the set of linear orders over A . Consider the following preference profile $P \in \mathcal{L}(A)^N$,

$$\begin{array}{cccc} 1 & c & b & a \\ n-1 & a & b & c \end{array} ,$$

which represents one individual who prefers c to b , b to a , hence c to a ; and $n-1$ individuals who prefer a to b , b to c , hence a to c . At P , all BK compromises, except when $q = n$ (i.e. fall-back bargaining that would select b) and $q = 1$ (which would select both a and c), will ignore the single voter and will pick a as the collective outcome.

As a matter of fact, almost every SCR will ignore this “marginal minority” and choose a in this situation. While this choice is defensible on the grounds of qualified majoritarianism, it is questionable whether a can be qualified as a compromise. Observe that b receives unanimous support when each voter falls back one step from his ideal point. The question becomes more compelling when a remains the collective choice even if the ignored group is much larger.

Example 7.2. Consider the following preference profile with $n = 100$:

$$\begin{array}{cccc} 49 & c & b & a \\ 51 & a & b & c \end{array} .$$

When $q \in \llbracket 1, \frac{n}{2} - 1 \rrbracket$, all BK compromises pick $\{a, c\}$, and, when $q = \frac{n}{2}$ and $q = \frac{n}{2} + 1$, all BK compromises pick $\{a\}$. Again, it does not appear as a compromise as almost half of voters reach their best alternative while the remaining half have to be contented with their worst one.

¹This objection to the compromise nomenclature was raised by Jean-François Laslier during a CNRS workshop on compromising hosted by Istanbul Bilgi University in 2018.

Observe that all these SCRs impose to voters a willingness to compromise, but do not effectively ensure a compromise as the collective choice. In fact, the term “compromise” in this literature refers to procedural compromises that differ from outcome oriented compromises, a conceptual distinction that seems to be overlooked in the literature. This can also be viewed as a distinction between ex-ante and ex-post compromises, where the profile is the source of uncertainty.

To define an ex-post compromise, we adapt a concept of equal losses that considers allocation of continuous utilities. This principle is used for bargaining [Chun, 1988, Chun and Peters, 1991] and bankruptcy problems [Herrero and Villar, 2001]. We introduce two definitions of compromise. In both of them, we pick a spread measure that determines how equally a given vector of numbers is distributed and we propose to make a collective choice where voters give up from their ideal points “as equally as possible”. The difference between the two is that one, called *egalitarian compromise*, insists on equality at the expense of Pareto efficiency while the other, called *Paretian compromise*, is constrained to pick among the Pareto efficient alternatives. We show that Pareto efficient SCRs cannot ensure egalitarian compromises under any spread measure. We prove that several well-known SCRs such as Condorcet extensions, scoring rules, q -approval fall-back bargaining, all fail to be Paretian compromises under any spread measure. We provide examples for which being a Paretian compromise would necessitate to pick an alternative that is, although Pareto optimal, ranked very low by all voters. Such alternative would never be picked by any of the above-mentioned SCRs.

We conclude that the equal-loss principle appears adequate for collective choice problems with at least three individuals, when egalitarianism, in the sense of conceding equally, is a major concern. Imagine a situation where the head of a laboratory needs to decide which project to fund and she asks for the preferences of the laboratory members. The workplace harmony is extremely important and, in order to avoid conflicts, the winning project must be equally supported by all members.

Consider now a situation with only two voters. As the vast literature on the ultimatum game [Güth and Kocher, 2014] suggests, mutual consent is hard to obtain when one individual sees injustice at the levels of mutual losses. The equal-loss principle seems to be crucial in this new scenario.

Collective choice models with two individuals can be interpreted as bargaining or arbitration problems. While the bargaining interpretation necessitates an explicitly defined disagreement outcome [Kibris and Sertel, 2007], the arbitration interpretation [Sprumont, 1993] remains within the classical collective choice environment with no explicit disagreement outcome. In this paper, we consider the latter interpretation.

Arbitration rules are thoroughly discussed by Barberà and Coelho [2020]. As prominent examples, we have fallback bargaining proposed by Brams and Kilgour [2001], the veto-rank and short listing procedures analysed by de Clippel et al. [2014] and the Pareto-and-veto rules analysed by Laslier et al. [2020]. These models consider discrete alternatives which are not contained by the classical Nash [1950] bargaining environment with convex utilities. We make the same assumption. However, as Mariotti [1998] and Nagahisa and Tanaka [2002] illustrate, the two worlds can be interconnected, as we do for the equal-loss principle of Chun [1988] and Chun and Peters [1991]. The arbitration environment presents an instance where the equal-loss

principle could matter and it is rather surprising to discover that most interesting SCRs used as arbitration solutions fail to be Paretian compromises.

The rest of the paper is organized as follows. Section 7.2 presents the basic notions and notation. Section 7.3 introduces egalitarian compromises and Paretian compromises, two concepts that turn out to be logically incompatible. Section 7.4 shows that with at least three individuals, many SCRs fail to pick a compromise. Section 7.5 considers the two-individual case, showing that most SCRs of the literature fail to pick compromises. Section 7.6 makes some concluding remarks.

7.2 Basic notions and notation

Consider a finite set N of individuals with $\#N = n \geq 2$ and a finite set A of alternatives with $\#A = m \geq 3$. We write $\mathcal{L}(A)$ for the set of linear orders over A . A generic element $\succ_i$ of $\mathcal{L}(A)$ stands for a preference of $i \in N$. This implies that, given any $x \neq y \in A$, precisely one of $x \succ_i y$ and $y \succ_i x$ holds while $x \succ_i x$ holds for no $x \in A$. Moreover, $x \succ_i y$ and $y \succ_i z$ implies $x \succ_i z \forall x, y, z \in A$.

A *profile* $P : N \rightarrow \mathcal{L}(A)$ associates with each individual $i \in N$ a preference $P(i) = \succ_i$. A *Social Choice Rule* (SCR) is a mapping $f : \mathcal{L}(A)^N \rightarrow 2^A \setminus \{\emptyset\}$.

We write $r_{\succ_i}(x) = \#\{y \in A \mid y \succ_i x\} + 1$ for the *rank* of $x \in A$ at $\succ_i \in \mathcal{L}(A)$. We denote by $\lambda_{\succ_i}(x) = r_{\succ_i}(x) - 1$ the *loss* in terms of ranks for $i \in N$ with preference $\succ_i$, when x is elected instead of the best alternative for i . The mapping $\lambda_P : A \rightarrow \llbracket 0, m-1 \rrbracket^N$ assigns to each $x \in A$ the *loss vector* $\lambda_P(x) = (\lambda_{\succ_i}(x))_{i \in N}$ induced by the election of x . The double brackets denote intervals in the integers.

We are interested in measuring the spread of loss vectors. To this end, we define a *spread measure* $\sigma : \llbracket 0, m-1 \rrbracket^N \rightarrow \mathbb{R}_+$ as a function that associates a spread value to every possible loss vector. We write Σ for the set of spread measures σ that satisfy, for every $l \in \llbracket 0, m-1 \rrbracket^N$, $\sigma(l) = 0 \Leftrightarrow l_i = l_j \forall i, j \in N$. Thus, the spread of l gets its lowest value 0 in case of perfect equality and only in this case. Note that this condition incorporates the minimal requirement to identify a spread measure and leads to the largest set of spread measures one could define. As discussed in Section 7.4.4, this flexibility has the advantage of making our results more general.

Given any distinct $x, y \in A$, we say that x *Pareto dominates* y at $P \in \mathcal{L}(A)^N$ (or equivalently y is *Pareto dominated* by x at P) iff $x \succ_i y, \forall i \in N$. We denote by $\text{PO}(P) = \{x \in A \mid \forall y \in A \setminus \{x\}, \exists i \in N \mid x \succ_i y\}$ the set of *Pareto optimal* alternatives at P . A SCR f is *Paretian* iff $f(P) \subseteq \text{PO}(P) \forall P \in \mathcal{L}(A)^N$.

7.3 Egalitarian versus Paretian compromises

7.3.1 Egalitarian compromises

We let $\text{argmin}_X(\sigma \circ \lambda_P) = \{x \in X \mid \forall y \in X : \sigma(\lambda_P(x)) \leq \sigma(\lambda_P(y))\}$ denote the minimal elements of $X \subseteq A$ according to $(\sigma \circ \lambda_P)$. In other words, $\text{argmin}_X(\sigma \circ \lambda_P)$ denotes the alternatives in X whose loss vectors are the most equally distributed according to the spread measure σ .

In what follows, we define some classes of SCRs that we are interested in analyzing.

Definition 7.1. A SCR f is an Egalitarian Compromise (EC) iff

$$\exists \sigma \in \Sigma \mid \forall P \in \mathcal{L}(A)^N \text{ we have } f(P) \subseteq \min_A(\sigma \circ \lambda_P).$$

Definition 7.2. A SCR f is Egalitarian Compromise Compatible (ECC) iff

$$\exists \sigma \in \Sigma \mid \forall P \in \mathcal{L}(A)^N \text{ we have } f(P) \cap \min_A(\sigma \circ \lambda_P) \neq \emptyset.$$

Under a SCR that is EC (resp., ECC), *all* (resp., *some*) winners are among the alternatives with most equally distributed losses. Clearly, EC is a subclass of ECC. Perhaps less obviously, being ECC (or EC) is incompatible with being Paretian. This will be deduced from the following proposition, which will also be useful to prove other theorems.

Proposition 7.1. For $n \geq 2, m \geq 3$, there exists a profile $P \in \mathcal{L}(A)^N$ and an alternative a_m such that $\forall i \in N: r_{\succ_i}(a_m) = m$, and such that $\forall \sigma \in \Sigma: \min_A(\sigma \circ \lambda_P) = \{a_m\}$; hence, $\min_A(\sigma \circ \lambda_P) \cap \text{PO}(P) = \emptyset$.

Proof. Consider the following profile P :

$$\begin{array}{cccccc} \mathbf{1} & a_1 & a_2 & \dots & a_{m-1} & a_m \\ \mathbf{n-1} & a_{\pi(1)} & a_{\pi(2)} & \dots & a_{\pi(m-1)} & a_m \end{array},$$

where π is the following permutation over $\llbracket 1, m-1 \rrbracket$:

$$\pi(i) = \begin{cases} i+1 & \text{if } i \in \llbracket 1, m-2 \rrbracket \\ 1 & \text{if } i = m-1 \end{cases}.$$

In P , a_m is the only alternative such that $r_{\succ_i}(a_m)$ is independent of i ; hence, $\sigma(\lambda_P(a)) > 0$, $\forall a \in A \setminus \{a_m\}$, $\forall \sigma \in \Sigma$. Thus, the set $\min_A(\sigma \circ \lambda_P)$ consists of the sole element a_m , and, because a_m is Pareto dominated, $\min_A(\sigma \circ \lambda_P) \cap \text{PO}(P) = \emptyset$. $\square$

Our main result for Section 7.3.1 follows easily.

Theorem 7.1. For $n \geq 2, m \geq 3$, no Paretian SCR is ECC.

Proof. Proving this amounts to show that $\forall \sigma \in \Sigma, \exists P \in \mathcal{L}(A)^N \mid \text{PO}(P) \cap \min_A(\sigma \circ \lambda_P) = \emptyset$. Suffices to use Proposition 7.1, which asserts that there exists a profile P such that $\forall \sigma \in \Sigma: \min_A(\sigma \circ \lambda_P) \cap \text{PO}(P) = \emptyset$. $\square$

7.3.2 Paretian compromises

Having seen the tension for a SCR being Paretian and ECC, we investigate the consequences of inverting the order of priorities by insisting that at least some of the winning alternatives are Pareto optimal, and considering the most equally distributed loss vectors among those.

We consider two classes of SCRs. Observe that $\min_{\text{PO}(P)}(\sigma \circ \lambda_P)$ denotes the set of Pareto optimal alternatives whose loss vectors are the most equally distributed according to the spread measure σ .

Definition 7.3. A SCR f is a Paretian Compromise PC iff

$$\exists \sigma \in \Sigma \mid \forall P \in \mathcal{L}(A)^N \text{ we have } f(P) \subseteq \min_{PO(P)}(\sigma \circ \lambda_P).$$

Definition 7.4. A SCR f is Paretian Compromise Compatible PCC iff

$$\exists \sigma \in \Sigma \mid \forall P \in \mathcal{L}(A)^N \text{ we have } f(P) \cap \min_{PO(P)}(\sigma \circ \lambda_P) \neq \emptyset.$$

Again, it is clear that PC is a subclass of PCC. It will also probably come with no surprise that for a SCR, being PC is incompatible with being ECC, as being PC requires to be Paretian, which permits to use Theorem 7.1. On the other hand, it is less immediate that being EC is incompatible with being PCC, because being PCC does not require to be Paretian. This is however true.

Theorem 7.2. For $n \geq 2, m \geq 3$, no SCR is both EC and PCC.

Proof. Considering the profile P of Proposition 7.1, with a_m the alternative mentioned there, any EC f and any $\sigma \in \Sigma$, suffices to prove that $f(P) \cap \min_{PO(P)}(\sigma \circ \lambda_P) = \emptyset$.

First, from Proposition 7.1 we have that $\{a_m\} \cap PO(P) = \emptyset$, hence $\{a_m\} \cap \min_{PO(P)}(\sigma \circ \lambda_P) = \emptyset$.

Second, because f is EC, for some $\bar{\sigma}$, $f(P) \subseteq \min_A(\bar{\sigma} \circ \lambda_P)$. Using Proposition 7.1 again, we see that $\min_A(\bar{\sigma} \circ \lambda_P) = \{a_m\}$, hence $f(P) = \{a_m\}$.

That $f(P) \cap \min_{PO(P)}(\sigma \circ \lambda_P) = \emptyset$ follows from these two facts. $\square$

It is interesting to note that the incompatibility is not complete, however.

Remark 7.1. For $n \geq 2, m \geq 3$, there exist SCRs that are both ECC and PCC, such as the SCR that selects the whole set of alternatives at every profile. However, this SCR fails to be Paretian, as is any SCR that is ECC.

7.4 Which SCRs are compromises?

In this section we assume $n \geq 3$ and leave the analysis of $n = 2$ to the next section.

7.4.1 Condorcet consistent rules

An alternative $x \in A$ is a *Condorcet winner* at $P \in \mathcal{L}(A)^N$ iff for all $y \in A \setminus \{x\}$, $\#\{i \in N \mid x \succ_i y\} > \#\{i \in N \mid y \succ_i x\}$. So each profile admits either no or a unique Condorcet winner. An SCR f is *Condorcet consistent* iff $f(P) = \{x\}$ at each $P \in \mathcal{L}(A)^N$ that admits x as the (unique) Condorcet winner.

Theorem 7.3. Let $n \geq 3$ and $m \geq 3$. A Condorcet consistent SCR f is neither ECC nor PCC.

Proof. Consider the following profile P , where the dots represent the sequence a_4 to a_m :

$$\begin{array}{ccccccc} \mathbf{n-1} & a_1 & a_2 & a_3 & \dots & & \\ \mathbf{1} & a_3 & a_2 & \dots & a_1 & & \end{array}.$$

Consider any Condorcet consistent SCR f . Thus, $f(P) = \{a_1\}$. However, $\min_A(\sigma \circ \lambda_P) = \min_{\text{PO}(P)}(\sigma \circ \lambda_P) = \{a_2\} \forall \sigma \in \Sigma$, so there exists a profile P such that both $f(P) \cap \min_A(\sigma \circ \lambda_P)$ and $f(P) \cap \min_{\text{PO}(P)}(\sigma \circ \lambda_P)$ are empty. $\square$

Note that Condorcet consistent rules need not be Paretian so the fact that they all fail ECC does not follow from Theorem 7.1.

7.4.2 Scoring rules

A *score vector* is an m -tuple $w = (w_1, \dots, w_m) \in \llbracket 0, 1 \rrbracket^m$ with $w_1 = 1$, $w_m = 0$ and $w_i \geq w_{i+1} \forall i \in \llbracket 1, m-1 \rrbracket$. Given a score vector w , we write $s^w(x, P) = \sum_{i \in N} w_{r_{\succ_i}(x)}$ for the score of $x \in A$ at $P \in \mathcal{L}(A)^N$. Each vector w identifies a *scoring rule* f_n^w defined as $f_n^w(P) = \{x \in A : s^w(x, P) \geq s^w(y, P) \forall y \in A\}$ for every $P \in \mathcal{L}(A)^N$.

We first show that no scoring rule is ECC, for any value of n and m at least 3.

Theorem 7.4. *Let $n \geq 3$ and $m \geq 3$. No score vector w induces a scoring rule f_n^w that is ECC.*

Proof. Take any score vector w . Consider the profile P of Proposition 7.1. Observe that $\min_A(\sigma \circ \lambda_P) = \{a_m\} \forall \sigma \in \Sigma$. However, as $w_1 > w_m$, we have $s^w(a_1, P) > s^w(a_m, P)$ which implies $a_m \notin f_n^w(P)$. $\square$

We call *antiplurality score vector* the vector w such that $w_i = 1, \forall i \in \llbracket 1, m-1 \rrbracket$, and $w_m = 0$.

Theorem 7.5. *Let $m \geq 3$ and let w be the antiplurality score vector. The SCR f_n^w satisfies PCC for all $n \geq 3$.*

Proof. Define $\bar{\sigma} \in \Sigma$ as, $\forall l \in \llbracket 0, m-1 \rrbracket^N$: $\bar{\sigma}(l) = 1$ iff $\exists i, j \in N \mid l_i \neq l_j$; $\bar{\sigma}(l) = 0$ otherwise. We show the non-emptiness of $f_n^w(P) \cap \min_{\text{PO}(P)}(\bar{\sigma} \circ \lambda_P)$ for any profile P .

Let $k = \min_{\text{PO}(P)}\{(\bar{\sigma} \circ \lambda_P)(x)\}$ be the minimal value attained by $\bar{\sigma} \circ \lambda_P$ over $\text{PO}(P)$. By construction of $\bar{\sigma}$, k equals either 0 or 1.

For $k = 1$, take any $x \in f_n^w(P) \cap \text{PO}(P)$. This intersection is non-empty because whenever the antiplurality rule picks a Pareto dominated alternative z , it also picks all alternatives which Pareto dominate z . By definition of $\bar{\sigma}$, $\bar{\sigma}(x) \leq 1$, hence, $x \in \min_{\text{PO}(P)}(\bar{\sigma} \circ \lambda_P)$.

For $k = 0$, take any $x \in \min_{\text{PO}(P)}(\bar{\sigma} \circ \lambda_P)$. As $\bar{\sigma}(\lambda_P(x)) = 0$, we have, $\forall i, j \in N$: $\lambda_i^P(x) = \lambda_j^P(x)$, hence, $\forall i, j \in N$: $r_{\succ_i}(x) = r_{\succ_j}(x)$. The case $r_{\succ_i}(x) = m, \forall i \in N$ is ruled out by $x \in \text{PO}(P)$. Hence, $r_{\succ_i}(x) \leq m-1, \forall i \in N$, hence, $x \in f_n^w(P)$. $\square$

It is worth noting that the antiplurality rule f_n^w is not Paretian, hence fails PC for all $n \geq 3$. This can be seen by picking a unanimous profile $P \in \mathcal{L}(A)^N$ with $a_1 \succ_i a_2 \succ_i \dots \succ_i a_m \forall i \in N$, where $\min_{\text{PO}(P)}(\sigma \circ \lambda_P) = \{a_1\} \forall \sigma \in \Sigma$ while $f_n^w(P) = A \setminus \{a_m\}$.

Theorem 7.6. *Let $m \geq 3$. Take any score vector w which is not the antiplurality score vector. For some $n \geq 3$, the SCR f_n^w fails PCC.*

Proof. Take any $m \geq 3$ and any score vector w that is not the antiplurality score vector; therefore, $w_{m-1} < 1$. Pick any n such that $n \geq m-1$ and $n > \frac{1}{1-w_{m-1}}$. Consider a profile $P \in \mathcal{L}(A)^N$ conforming to

$$\begin{array}{ccccccc} i = 1 & a_2 & \dots & a_m & a_1 & & \\ 2 \leq i \leq m-2 & a_1 & \dots & a_m & a_i & & \\ m-1 \leq i \leq n & a_1 & \dots & a_m & a_{m-1} & & \end{array},$$

where all alternatives except a_m appear at least once in the last rank. Thus, for every $\sigma \in \Sigma$, we have $\sigma(\lambda_P(x)) > 0 \forall x \in A \setminus \{a_m\}$ while $\sigma(\lambda_P(a_m)) = 0$. Moreover, $a_m \in \text{PO}(P)$. Thus, $\min_{\text{PO}(P)}(\sigma \circ \lambda_P) = \{a_m\} \forall \sigma \in \Sigma$. On the other hand, $s^w(a_1; P) = n-1$, $s^w(a_m; P) = n \cdot w_{m-1}$ and as $n > \frac{1}{1-w_{m-1}}$ (or, equivalently, $n-1 > nw_{m-1}$), we have $s^w(a_1; P) > s^w(a_m; P)$, establishing $a_m \notin f^w(P)$, thus $f^w(P) \cap \min_{\text{PO}(P)}(\sigma \circ \lambda_P) = \emptyset \forall \sigma \in \Sigma$. $\square$

7.4.3 BK compromises

Given any $k \in \llbracket 1, m \rrbracket$, we write $n_k(x, P) = \#\{i \in N \mid r_{\succ_i}(x) \leq k\}$ for the k -support that x gets at P , that is, the number of individuals for whom the rank of alternative $x \in A$ is lower than or equal to k in the profile $P \in \mathcal{L}(A)^N$. Note that $n_k(x, P) \in \llbracket 1, n \rrbracket$ is non-decreasing on k and $n_m(x, P) = n$. For each $q \in \llbracket 1, n \rrbracket$, we define $\rho_q(x, P) = \min\{k \in \llbracket 1, m \rrbracket \mid n_k(x, P) \geq q\}$ as the minimal rank k at which the k -support that x gets at P is at least q . We write $\rho_q(P) = \min_{x \in A} \{\rho_q(x, P)\}$ for the minimal rank k at which the k -support that some alternative gets at P is at least q . A *BK compromise with threshold q* is the SCR f_q defined for each $P \in \mathcal{L}(A)^N$ as $f_q(P) = \{x \in A \mid n_{\rho_q(P)}(x, P) \geq q\}$. We can also define a tie breaking version of the BK compromise where among the winners only the alternatives with the greatest support are selected: $f'_q(P) = \{x \in A \mid n_{\rho_q(P)}(x, P) \geq n_{\rho_q(P)}(y, P), \forall y \in A\}$.

We first consider the BK compromise with threshold $q = n$, f_n , which corresponds to the rule also known as *fallback bargaining* [Brams and Kilgour, 2001].

Theorem 7.7. *Let $n \geq 3$ and $m \geq 3$. The BK compromise f_n satisfies PC.*

Proof. Define $\bar{\sigma} \in \Sigma$ as, $\forall l \in \llbracket 0, m-1 \rrbracket^N$: $\bar{\sigma}(l) = 1$ iff $\exists i, j \in N \mid l_i \neq l_j$; $\bar{\sigma}(l) = 0$ otherwise. Considering any $x \in f_n(P)$, let us show that $x \in \min_{\text{PO}(P)}(\bar{\sigma} \circ \lambda_P)$. Because $x \in f_n(P)$, $x \in \text{PO}(P)$, and therefore, suffices to show that $\forall y \in \text{PO}(P)$, $\bar{\sigma}(\lambda_P(y)) \geq \bar{\sigma}(\lambda_P(x))$. Given the choice of $\bar{\sigma}$, picking any $y \in \text{PO}(P)$ with $y \neq x$, suffices to show that $\bar{\sigma}(\lambda_P(y)) = 1$, equivalently, that $\exists i, j \in N \mid r_{\succ_i}(y) \neq r_{\succ_j}(y)$. Because $x \in f_n(P)$, $\rho_n(P) = \rho_n(x, P) = \max_N r_{\succ_i}(x)$. It follows from $\rho_n(P) = \min_{z \in A} \{\rho_n(z, P)\}$ that $\rho_n(y, P) \geq \rho_n(x, P)$, thus, $\exists i \in N \mid r_{\succ_i}(y) \geq \rho_n(P)$. Also, $y \in \text{PO}(P)$ implies that $\exists j \in N \mid r_{\succ_j}(y) < r_{\succ_j}(x)$, thus $\exists j \in N \mid r_{\succ_j}(y) < \rho_n(P)$. Therefore, $r_{\succ_i}(y) \neq r_{\succ_j}(y)$. $\square$

Theorem 7.8. *Let $n \geq 3$ and $m \geq 3$. The BK compromise f_n fails ECC.*

Proof. As f_n is Paretian, the proof comes straightforward from Theorem 7.1. $\square$

Theorem 7.9. *Let $n \geq 3$ and $m \geq 3$. A BK compromise f_q with threshold $q \in \llbracket 1, n-1 \rrbracket$ is neither ECC nor PCC.*

Proof. Consider the following profile P (also used in the proof of Theorem 7.3), where the dots represent the sequence a_4 to a_m :

$$\begin{array}{ccccccc} \mathbf{n-1} & a_1 & a_2 & a_3 & \dots & & \\ & 1 & a_3 & a_2 & \dots & a_1 & \end{array}.$$

When $q = 1$ we have that $f_1(P) = \{a_1, a_3\}$ and when $q \in \llbracket 2, n-1 \rrbracket$ we have that $f_q(P) = \{a_1\}$. Because $\sigma(\lambda_P(a_2)) = 0$ and $\sigma(\lambda_P(a_1)) > 0$ and $\sigma(\lambda_P(a_3)) > 0$, neither $\min_A(\sigma \circ \lambda_P)$ nor $\min_{PO(P)}(\sigma \circ \lambda_P)$ contain a_1 nor a_3 for any $\sigma \in \Sigma$.

Note that for $q \in \llbracket 1, n-1 \rrbracket$ we have that $f'_q(P) = \{a_1\}$ since the tie breaking version of the BK compromise selects the alternatives with the greatest support. This version of the rule was used in the proof for Theorem 7.9 published by Cailloux et al. [2022]. $\square$

7.4.4 Restrictions on sigma

The perfect equality recognition condition we adopt for spread measures, i.e., that the spread gets its lowest value 0 in case of perfect equality and only in this case, is very basic. Unless this condition is violated, Σ is the largest set of spread measures we could conceive. On the other hand, it is possible to let Σ shrink by imposing additional conditions over spread measures. Nevertheless, as the satisfaction of PC, PCC, EC, or ECC requires the existence of a spread measure, all of our negative results, namely, those expressed by Theorems 7.1, 7.2, 7.3, 7.4, 7.6, 7.8 and 7.9 prevail when Σ is restricted. In a similar vein, the positive results in Theorems 7.5 and 7.7 risk to be lost with additional conditions over spread measures. Indeed, this section shows that a mild restriction removes the positive results concerning the only two rules that we found to be compatible with any of our compromise concepts.

Definition 7.5. *Given any $m \geq 4$ and $n \geq \max\{4, m-1\}$, we say that a spread measure σ satisfies condition $C_{m,n}$ iff we have $\sigma(m-3, m-1, m-2, \dots, m-2) < \sigma(m-2, m-3, \dots, 1, 0, \dots, 0)$.*

As both vectors are n dimensional, the term $m-2$ repeats $n-2$ times in the first vector and the term 0 repeats $n-m+2$ times in the second vector.

This condition imposes a very reasonable requirement on spread measures for large values of m and n . Asking for $\sigma(1, 3, 2, 2)$ to be smaller than $\sigma(2, 1, 0, 0)$ is demanding while asking for $\sigma(5, 7, 6, 6, 6, 6, 6)$ to be smaller than $\sigma(6, 5, 4, 3, 2, 1, 0)$ reflects a mild assumption. In any case, as we state below, several well-known spread measures of the literature (see Allison [1978] for a comprehensive account) satisfy Definition 7.5. Letting $\bar{l} = \sum_{i=1}^n l_i / n$ denote the arithmetic mean of the values of $l = (l_1, \dots, l_n)$, we consider the following measures:

- the mean absolute difference $\sigma_{mad}(l) = \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n |l_i - l_j|$;
- the average absolute deviation $\sigma_{ad}(l) = \frac{\sum_{i=1}^n |l_i - \bar{l}|}{n}$;
- the standard deviation $\sigma_{sd}(l) = \sqrt{\frac{\sum_{i=1}^n (l_i - \bar{l})^2}{n}}$;
- the Gini coefficient $\sigma_G(l) = \frac{\sum_{i=1}^n \sum_{j=1}^n |l_i - l_j|}{2 \cdot n \cdot \sum_{i=1}^n l_i}$.

Remark 7.2. We checked experimentally that σ_{mad} , σ_{ad} , σ_{sd} and σ_G all satisfy condition $C_{m,n}$, for $m \in \llbracket 4, 1000 \rrbracket$ and $n \in \llbracket b, 1000 \rrbracket$ where $b = \max\{4, m - 1\}$.

We write $\Sigma^{C_{m,n}} \subseteq \Sigma$ for the set of spread measures that satisfy condition $C_{m,n}$.

Theorem 7.10. For all $m \geq 4$, $n \geq \max\{4, m - 1\}$, under $\Sigma^{C_{m,n}}$,

- 1) f_n^w fails PCC when w is the antiplurality score vector;
- 2) the BK compromise f_n fails PCC.

Proof. Given any $\sigma \in \Sigma^{C_{m,n}}$, let us show that there exists a profile P such that $f_n^w \cap \min_{PO(P)}(\sigma \circ \lambda_P) = \emptyset$ and $f_n \cap \min_{PO(P)}(\sigma \circ \lambda_P) = \emptyset$. To that aim, consider some $x, y \in A$ and some $P \in \mathcal{L}(A)^N$ with $r_{\succ_1}(x) = m - 2$, $r_{\succ_2}(x) = m$, $r_{\succ_i}(x) = m - 1 \ \forall i \in N \setminus \{1, 2\}$, and $r_{\succ_i}(y) = m - i \ \forall i \in \llbracket 1, m - 1 \rrbracket$, $r_{\succ_j}(y) = 1 \ \forall j \in \llbracket m, n \rrbracket$. Moreover, for each $z \in A \setminus \{x, y\}$, let $r_{\succ_i}(z) = m$ for some $i \in N$.

Note that $f_n^w(P) = f_n(P) = \{y\}$. On the other hand, $\lambda_P(x) = (m - 3, m - 1, m - 2, \dots, m - 2)$ and $\lambda_P(y) = (m - 2, m - 3, \dots, 1, 0, \dots, 0)$. As $\sigma(\lambda_P(x)) < \sigma(\lambda_P(y))$ (because $\sigma \in \Sigma^{C_{m,n}}$), we see that $y \notin \min_{PO(P)}(\sigma \circ \lambda_P)$. $\square$

7.5 Two-voters case

In addition to *fallback bargaining (FB)* [Brams and Kilgour, 2001] (defined in Section 7.4.3), we consider three prominent solutions of the literature.

Pareto-and-Veto rules (PV) [Moulin, 1983, Abreu and Sen, 1991, Laslier et al., 2020] distribute a veto power of v_1 and v_2 alternatives to voters 1 and 2, respectively, with $v_1 + v_2 = m - 1$. So, every voter $i = 1, 2$ (simultaneously) vetoes his worst v_i alternatives. The SCR picks all non-vetoes and Pareto optimal alternatives.

The *Veto-Rank mechanism (VR)* is commonly used in the selection of arbitrators [de Clippel et al., 2014]. Given a list of m (odd) alternatives (that are candidates to be arbitrators), each of the two voters (that are the two parties that must agree on an arbitrator) simultaneously vetoes his worst $\frac{m-1}{2}$ alternatives. The selected alternatives are the ones with the highest Borda score among the non-vetoes alternatives.

Again within the context of selecting arbitrators, de Clippel et al. [2014] propose and analyze *Shortlisting (SL)* where one of the two parties starts by vetoing her worst $\frac{m-1}{2}$ alternatives (m being odd), and then the second party chooses her best alternative out of the remaining ones. As the outcome of the procedure depends on the party that starts, symmetry among players is ensured by defining the solution as the union of the two outcomes where one and the other party starts.

Definition 7.6. Given any $m \geq 7$, a spread measure $\sigma \in \Sigma$ satisfies condition D_m iff $\sigma(\lceil \frac{m}{2} \rceil, \lceil \frac{m}{2} \rceil - 2) < \sigma(0, \lceil \frac{m}{2} \rceil - 1)$ and $\sigma(\lceil \frac{m}{2} \rceil - 2, \lceil \frac{m}{2} \rceil) < \sigma(\lceil \frac{m}{2} \rceil - 1, 0)$.

For $m = 7$ the condition requires $\sigma(4, 2) < \sigma(0, 3)$ and $\sigma(2, 4) < \sigma(3, 0)$ which is reasonable in our context. When the value of m is larger, the condition appears even more convincing. As m grows, the distance between 0 and $\lceil \frac{m}{2} \rceil - 1$ grows, while the distance between $\lceil \frac{m}{2} \rceil$ and

$\lceil \frac{m}{2} \rceil - 2$ remains constant. Requiring, for example, the spread of (15, 13) to be smaller than the spread of (0, 14) is very reasonable.

We write $\Sigma^{D_m} \subseteq \Sigma$ for the set of spread measures that satisfy condition D_m .

Theorem 7.11. *Let $m \geq 7$. Under Σ^{D_m} , FB and PV fail PCC. Furthermore, when m is odd, VR and SL also fail PCC.*

Proof. Take any $m \geq 7$ and any $\sigma \in \Sigma^{D_m}$. Define $\alpha = \lceil \frac{m}{2} \rceil - 1$ and $\beta = \lceil \frac{m}{2} \rceil - 2$. It follows from $\sigma \in \Sigma^{D_m}$, that $\sigma(\alpha + 1, \beta) < \sigma(0, \beta + 1)$ and $\sigma(\beta, \alpha + 1) < \sigma(\beta + 1, 0)$. For m odd, note that $\alpha + \beta + 2 = m$ and consider the profile P where voter i_1 has the preference $x \succ a_1 \succ \dots \succ a_\alpha \succ y \succ b_1 \succ \dots \succ b_\beta$ and voter i_2 has the preference $b_1 \succ \dots \succ b_\beta \succ y \succ x \succ a_1 \succ \dots \succ a_\alpha$. For m even, note that $\alpha + \beta + 3 = m$, and define the profile P in the same way, except that a supplementary alternative z is added at the bottom of both rankings.

Note that $\sigma(\lambda_P(y)) = \sigma(\alpha + 1, \beta)$ and that $\sigma(\lambda_P(x)) = \sigma(0, \beta + 1)$. Therefore, $\sigma(\lambda_P(y)) < \sigma(\lambda_P(x))$. As y is not Pareto-dominated, an SCR that uniquely picks x at P cannot be PCC. In a similar vein, at the profile P' which is obtained by the inversion of the preferences of i_1 and i_2 at P , an SCR that is PCC cannot pick x uniquely.

The proof will be concluded by showing that FB, PV, and (when m is odd) VR and SL all pick only x at P or at P' .

We readily see that FB picks only x at P (and at P') since x is the first alternative which reaches the unanimous consent. For PV, let $v_{i_1} \geq v_{i_2}$ (thus $v_{i_1} \geq \lceil \frac{m-1}{2} \rceil \geq \lceil \frac{m-2}{2} \rceil = \beta + 1$ and $v_{i_2} \leq \lfloor \frac{m-1}{2} \rfloor = \lceil \frac{m-2}{2} \rceil = \alpha$), and consider the profile P . Observe that the first voter vetoes at least y and every b_j ($1 \leq j \leq \beta$) while no voter vetoes x . As x Pareto-dominates every a_j ($1 \leq j \leq \alpha$), PV picks only x at P . When $v_{i_2} \geq v_{i_1}$, a similar reasoning yields that PV picks only x at P' .

Now let m be odd.

For VR, a reasoning similar to the one applied to PV yields x as the unique choice at P : each voter vetoes her worst $\frac{m-1}{2}$ alternatives, thus i_1 vetoes y and every b_j ($1 \leq j \leq \beta$) and i_2 vetoes every a_j ($1 \leq j \leq \alpha$). The alternative x is the only non-vetoed alternative, so it is selected as the sole winner.

Finally, SL also picks x , as it is the unique winner no matter which voter starts the veto phase. If i_1 starts, y and every b_j ($1 \leq j \leq \beta$) get vetoed, then i_2 chooses her best alternative out of the remaining ones which is x . If i_2 starts, every a_j ($1 \leq j \leq \alpha$) get vetoed, then i_1 chooses her best alternative which is x . $\square$

7.6 Concluding remarks

We define an ex-post compromise as an outcome where individuals give up as equally as possible from their ideal points. With three or more individuals, several well known SCRs fail to pick ex-post compromises, under any reasonable meaning attributed to “giving up equally”. Our findings cover Condorcet extensions and scoring rules but also BK compromises, which impose a willingness to compromise without ensuring a compromised outcome. In particular we find that: no Condorcet procedure is ECC or PCC (Theorem 7.3), no scoring rule is ECC (Theorem 7.4)

or PCC except for the antiplurality rule (Theorems 7.5 and 7.6), BK compromises are neither ECC or PCC (Theorems 7.8 and 7.9) except for fallback bargaining that is PC (Theorem 7.7).

Our impossibility results are stated for the set of spread measures Σ , but they prevail for any subset of Σ . As Σ is the largest set of sensible spread measures, they are valid for any specific concept of equity one might pick. On the other hand, our possibility results on fallback bargaining and antiplurality are not propagated to subsets of Σ . In fact, as soon as a reasonably mild restriction on Σ is imposed, both rules are no longer PCC (Theorem 7.10). With two individuals and a similar restriction on Σ , all well-known two-person SCRs of the literature, namely, fallback bargaining, Pareto and veto rules, short listing and veto rank, fail to pick ex-post compromises (Theorem 7.11).

The exclusion of the equal-loss principle by almost all SCRs of the literature leads to ask whether the principle is uninteresting in a discrete social choice context. This seems to be the case for voting situations where the number of voters exceeds the number of candidates and usually every candidate is ranked last by at least one voter. In these cases, the main concern is about the support of alternatives rather than equality. On the other hand, two-person collective choice problems are typically interpreted as arbitration or bargaining situations where mutual consent is a critical element in reaching a solution. Thus, the equal-loss principle appears to be valid for two-person collective choice problems and our analysis raises the question of designing new discrete arbitration rules compatible with the equal-loss principle.

We also want to mention that different notions of compromise can be conceived. Börgers [1991] defines as compromises all Pareto-optimal alternatives that are not the top choice of any individual. In this context a compromise does not always exist. This is also a possible approach to adopt and we thank an anonymous reviewer for this remark.

We close by noting, as one anonymous reviewer to whom we are grateful remarked, that the tension between equity and efficiency is not new in economics. In our paper we try to cast this tension in a context where it does not seem to have been considered yet. We certainly hope that this is only the beginning of a discussion that may lead to further progress in the future. In particular, viewing a compromise through the equal loss principle can be especially interesting in richer informational settings with a status-quo point, cardinal individual preferences or a continuum of alternatives.

PART III

CONCLUSION

CONCLUSION

Throughout this manuscript, we have addressed several problems related to the field of social choice. In particular, we focused on two scenarios. Considering a classical model in which the preferences of a set of voters over a set of alternatives are known, we defined two classes of voting rules able to reflect a notion of compromise in which egalitarianism, in the sense of conceding equally, is a major concern. Moreover, we stepped back from this standard perspective in which preferences are assumed to be known from the beginning and investigated the problem of preference elicitation in different settings. In what follows we describe the results of our works in more detail.

8.1 Summary of the contributions

Simultaneous Elicitation of PSR and Agent Preferences In Chapter 4 we studied the second of the aforementioned scenarios. Considering positional scoring rules with convex weights, we developed methods for preference elicitation under uncertainty of both the voting rule and the agent preferences. Assuming that those preferences are initially unknown to the system, the goal of the procedures is to incrementally reveal them through questioning and quickly acquire the most relevant information. We proposed the use of minimax regret both as a means of robust winner determination and as a guide to the process of simultaneous elicitation of preferences and voting rule. This serves mainly for two reasons, first to give an indication of the relevance of potential questions, but also to give us a measure of how many questions it takes to get to a low regret or to stop the elicitation process once the regret reaches a certain threshold. We presented incremental elicitation methods that at each step of the elicitation question either one of the agents or the chair, and we discussed several heuristics to choose questions quickly reduce the regret. Answers to questions are encoded as constraints: questions to the agents are comparisons between pairs of alternatives, while questions to the chair ask to select a winner out of a synthetic profile. We compared the effectiveness of several questioning strategies based on the current knowledge of the rule and preferences. Our experimental results suggest that regret-based elicitation is effective and allows to quickly reduce worst-case regret significantly. In particular:

- we defined different elicitation strategies that achieve good results within reasonable computation time and we compared their performances;
- we showed that with our elicitation methods, in particular with the Pessimistic strategy, it is possible to reach low regret with a reasonable number of questions;

- we analyzed the strategies on both real data and randomly generated profiles and found that low degree of similarity among preferences (as in impartial culture) is a more challenging setting than less varied profiles (as real preferences profiles);
- we showed that for the class of rules considered, asking a few questions to the chair suffice to reach low regret.

Moreover, as a part of our contribution, we proposed a novel mechanism for eliciting a voting rule by translating abstract questions about weights to a choice of the winning alternative on a concrete profile.

Preference Elicitation Under Majority Judgment In Chapter 6 we analyzed the preference elicitation mechanism used in a real voting scenario by the french political initiative LaPrimaire.org. In fact, they proposed a voting procedure to select an independent candidate for the french presidential election using MJ as voting rule. The procedure that they adopted consists of two rounds. In the first round each voter is asked to express her judgment, using MJ, on five random candidates. At the end of this phase the five candidates with the highest medians are considered the finalists who qualify for the second round. In the second round each voter is asked to express her judgment, using MJ, on all the five finalists. The candidate with the best median at the end of this phase is selected as representative for the presidential election. We investigated the consequences of randomness in the described preference elicitation procedure. We analyzed its cost, quantified as number of questions per voter, and its fairness, stating that a winner for the complete profile should not loose for lack of information. In particular, we have proven that this latter case it is possible. Given a complete profile it is possible to find an incomplete profile (such that a possible completion is equal to the considered profile) with a different winner. We also showed that the probability of this happening depends on the profile considered and we showed, as an extreme example, a case for which this probability is almost $1/2$. By denoting this as the probability of "missing the winner", we computed the average size of a grade vector after a given number of questions to each voter and studied what percentage of this vector we need to know so that this probability is low. Considering randomly generated vector of grades following a uniform distribution, we have found that when knowing only the 20% of the vector (thus after having asked 10 questions) the probability of missing the winner is close to zero. We have also considered a profile created following the grade distribution of a real voting scenario and analyzed the probability of missing the winner as a function of the number of voters. We have found that this greatly decreases as the number of voters increases, suggesting that this method is more suitable for elections with large numbers of voters but less appropriate for small profiles.

Compromising as an Equal-Loss Principle In Chapter 7 we analyzed the concept on compromise in the literature. We observed that almost all the voting rules that are known as *compromise rules* impose over individuals a willingness to compromise but they do not ensure an outcome where everyone has effectively compromised. We revisited the notion of compromise from an *equal-loss* perspective, favoring an outcome where every voter concedes as equally as possible, and we proposed some class of rules reflecting this concept. We denoted a voting

rule as *Egalitarian Compromise* (EC) (resp., *Egalitarian Compromise Compatible* (ECC)), if *all* (resp., *some*) winners are among the alternatives with the most equally distributed losses. Furthermore, we denoted a voting rule as *Paretian Compromise* (PC) (resp., *Paretian Compromise Compatible* (PCC)), if *all* (resp., *some*) winners are among the Pareto optimal alternatives with the most equally distributed losses. We proved that no Condorcet procedure is ECC or PCC, no scoring rule is ECC or PCC except for the antiplurality rule, BK compromises are neither ECC or PCC except for fallback bargaining that is PC. We showed that, although our results are stated for a broad set of spread measures Σ , they prevail for any subset of Σ . Moreover, as soon as we consider such subset where a reasonably mild restriction on Σ is imposed, fallback bargaining and antiplurality are no longer PCC. We also considered a specific voting scenario with two individuals and a similar restriction on Σ , and we proved that all well-known two-person voting rules of the literature, namely, fallback bargaining, Pareto and veto rules, short listing and veto rank, are not PCC.

8.2 Future work

There are many interesting directions for future works for each of the problems considered.

Regarding the simultaneous elicitation of voting rule and agent preferences, one could develop more strategies with different heuristics to compare to the ones we proposed. For example, less computationally costly strategies would allow testing of larger datasets, with more voters and more alternatives. It would also be interesting to expand the elicitation of the voting rules other than scoring rules or to relax the convexity constraint. Furthermore, we think that transforming questions into example profiles is a very interesting concept that is not explored much in the literature and could be applied in different settings.

The chapter on the analysis of the elicitation procedure using MJ is not yet fully mature so there are many ideas for possible extensions. One could investigate, for example, the manipulability of the elicitation process itself. In particular, how the party in charge of questioning the voters can manipulate the outcome by directing questions to specific voters. Another direction may involve the explainability of this process. Voters may find difficult to believe that judging only one person gives a good approximation of the result that would be obtained by asking for the full profile.

Finally, when considering the concept of compromise, different notions of compromise can be conceived. One approach might consider choosing alternatives that are neither the best nor the worst of any voter. It all depends on the idea of justice one wants to represent and the priorities in the scenario under consideration. Furthermore, the trade-off between equity and efficiency can be explored in other, and richer, settings.

To conclude, we like to think that each of our contributions is merely the beginning of a long series of future expansions.

BIBLIOGRAPHY

- D. Abreu and A. Sen. Virtual implementation in nash equilibrium. *Econometrica*, 59(4):997–1021, 1991. doi:10.2307/2938171.
- P. D. Allison. Measures of inequality. *American Sociological Review*, 43:865, 1978. ISSN 0003-1224. doi:10.2307/2094626.
- G. A. Aristotle trd. Kennedy. *On Rhetoric: A Theory of Civic Discourse*. New York: Oxford University Press, 2006. URL <https://books.google.fr/books?id=0mxVngEACAAJ>.
- K. Arrow, A. K. Sen, and K. Suzumura, editors. *Handbook of Social Choice and Welfare*, volume 1. Elsevier, 1 edition, 2002.
- K. Arrow, A. K. Sen, and K. Suzumura, editors. *Handbook of Social Choice and Welfare*, volume 2. Elsevier, 1 edition, 2011.
- K. J. Arrow. A difficulty in the concept of social welfare. *Journal of Political Economy*, 58(4):328–346, 1950. doi:10.2307/1828886.
- K. J. Arrow. *Social Choice and Individual Values*. New Haven: Cowles Foundation, 1st edition, 1951.
- R. J. Aumann and M. Maschler. Game theoretic analysis of a bankruptcy problem from the talmud. *Journal of Economic Theory*, 36(2):195–213, 1985. ISSN 0022-0531. doi:10.1016/0022-0531(85)90102-4.
- Y. Bachrach, N. Betzler, and P. Faliszewski. Probabilistic possible winner determination. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence*, AAAI’10, page 697–702. AAAI Press, 2010. doi:10.5555/2898607.2898720.
- M. Balinski. Réponse à des critiques du jugement majoritaire. *Revue économique*, 70(4):589–610, 2019. doi:<https://doi.org/10.3917/reco.704.0589>.
- M. Balinski and R. Laraki. A theory of measuring, electing, and ranking. *Proceedings of the National Academy of Sciences*, 104(21):8720–8725, 2007. doi:10.1073/pnas.0702634104.
- M. Balinski and R. Laraki. *Majority Judgment: Measuring, Ranking, and Electing*. The MIT Press, 01 2011. ISBN 9780262295604. doi:10.7551/mitpress/9780262015134.001.0001.
- K. Barber. *A Right to Representation: Proportional Election Systems for the Twenty-first Century*. Urban life and urban landscape series. Ohio State University Press, 2000. ISBN 9780814208540. URL <https://books.google.fr/books?id=bDCwpD1RiocC>.
- S. Barberà and D. Coelho. Compromising on compromise rules. working paper, 08 2020.

BIBLIOGRAPHY

- J. P. Barthelemy and B. Monjardet. The median procedure in cluster analysis and social choice theory. *Mathematical Social Sciences*, 1(3):235–267, 1981. ISSN 0165-4896. doi:[https://doi.org/10.1016/0165-4896\(81\)90041-X](https://doi.org/10.1016/0165-4896(81)90041-X).
- J. Bartholdi, C. A. Tovey, and M. A. Trick. Voting schemes for which it can be difficult to tell who won the election. *Social Choice and Welfare*, 6(2):157–165, 1989. ISSN 01761714, 1432217X. doi:10.2307/41105913.
- G. W. Bassett and J. Persky. Rating skating. *Journal of the American Statistical Association*, 89(427):1075–1079, 1994. doi:10.1080/01621459.1994.10476845.
- G. W. Bassett and J. Persky. Robust voting. *Public Choice*, 99(3-4):299–310, 1999. doi:10.1023/A:1018324807861.
- D. Baumeister and T. Hoguebe. Manipulative design of scoring systems. In *Proceedings of AAMAS’19*, 2019. URL <https://dl.acm.org/doi/10.5555/3306127.3331928>.
- D. Baumeister and J. 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. ISSN 0020-0190. doi:<https://doi.org/10.1016/j.ipl.2011.11.016>.
- D. Baumeister, M. Roos, and J. Rothe. Computational complexity of two variants of the possible winner problem. In *The 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 2, AAMAS ’11*, page 853–860, Richland, SC, 2011. International Foundation for Autonomous Agents and Multiagent Systems. doi:10.5555/2031678.2031739.
- K. Belahcene. *Towards accountable decision aiding : explanations for the aggregation of preferences*. PhD thesis, Université Paris-Saclay, 12 2018. URL <https://www.theses.fr/2018SACL101>.
- N. Benabbou, S. Di Sabatino Di Diodoro, P. Perny, and P. Viappiani. Incremental preference elicitation in multi-attribute domains for choice and ranking with the borda count. In *Proceedings of SUM 2016*, 2016. doi:10.1007/978-3-319-45856-4_6.
- D. Black. *The Theory of Committees and Elections*. Cambridge University Press, 1958. doi:10.1007/978-94-009-4225-7.
- J.-C. de Borda. Mémoire sur les élections au scrutin. *Histoire de l’Académie Royale des Sciences, Année 1781*, 102:657–665, 1784.
- C. Boutilier, R. Patrascu, P. Poupart, and D. Schuurmans. Constraint-based optimization and utility elicitation using the minimax decision criterion. *Artificial Intelligence*, 2006. doi:10.1016/j.artint.2006.02.003.
- D. Bouyssou, T. Marchant, M. Pirlot, P. Perny, A. Tsoukias, and P. Vincke. *Evaluation and Decision Models: A Critical Perspective*. International Series in Operations Research & Management Science. Springer US, 2000. ISBN 9780792372509. doi:10.1007/0-387-31099-1.

- D. Bouyssou, T. Marchant, M. Pirlot, A. Tsoukiàs, and P. Vincke. *Evaluation and Decision Models with Multiple Criteria: Stepping Stones for the Analyst*, volume 86. Springer New York, 01 2006. ISBN 0-387-31098-3. doi:10.1007/0-387-31099-1.
- S. Brams and P. Fishburn. *Approval Voting*, volume 72. Springer New York, NY, 01 2007. ISBN 978-0-387-49895-9. doi:10.1007/978-0-387-49896-6.
- S. J. Brams and D. M. Kilgour. Fallback bargaining. *Group Decision and Negotiation*, 10(4): 287–316, Jul 2001. doi:10.1023/A:1011252808608.
- F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia. *Handbook of Computational Social Choice*. Cambridge University Press, USA, 1st edition, 2016. ISBN 1107060435.
- D. Braybrooke. The possibilities of compromise. *Ethics*, 93(1):139–150, 1982. doi:10.2307/2380768.
- D. Braziunas. *Decision-Theoretic Elicitation of Generalized Additive Utilities*. PhD thesis, University of Toronto, 2012. URL <https://dl.acm.org/doi/book/10.5555/2519022>.
- T. Börgers. Undominated strategies and coordination in normalform games. *Social Choice and Welfare*, 8(1):65–78, 1991.
- O. Cailloux and U. Endriss. Eliciting a suitable voting rule via examples. In *Proceedings of ECAI 2014*, 2014. doi:10.3233/978-1-61499-419-0-183.
- O. Cailloux, B. Napolitano, and M. R. Sanver. Compromising as an equal loss principle. *Review of Economic Design*, May 2022. ISSN 1434-4750. doi:10.1007/s10058-022-00302-w.
- R. Camps, X. Mora, and L. Saumell. Choosing by means of approval-preferential voting. the revised approval choice. *CoRR*, abs/1411.1367, 2014. doi:10.48550/arXiv.1411.1367.
- C. Chang and Y.-A. Hwang. A characterization of the leximin solution of the bargaining problem. *Mathematical Methods of Operations Research*, 49(3):395–400, Jul 1999. ISSN 1432-5217. doi:10.1007/s001860050056.
- Y. Chun. The equal-loss principle for bargaining problems. *Economics Letters*, 26(2):103–106, 1988. doi:10.1016/0165-1765(88)90022-5.
- Y. Chun and H. Peters. The lexicographic equal-loss solution. *Mathematical Social Sciences*, 22:151–161, 1991. ISSN 0165-4896. doi:10.1016/0165-4896(91)90004-b.
- G. A. Churchill and J. P. Peter. Research design effects on the reliability of rating scales: A meta-analysis. *Journal of Marketing Research*, 21(4):360–375, 1984. doi:10.2307/3151463.
- J. S. Coggins and C. F. Perali. 64% majority rule in ducal venice: Voting for the doge. *Public Choice*, 97(4):709–723, Dec 1998. doi:10.1023/A:1004947715017.
- N. de Condorcet. *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. L'imprimerie royale, 1785. URL <https://gallica.bnf.fr/ark:/12148/bpt6k417181>.

BIBLIOGRAPHY

- N. de Condorcet. Sur la forme des élections. *Oeuvres*, 9, 1789. URL <https://gallica.bnf.fr/ark:/12148/bpt6k417270/f4.item>.
- N. de Condorcet. Sur les élections. *Journal d'Instruction Sociale*, 1793. URL <https://gallica.bnf.fr/ark:/12148/bpt6k41759m/f5.item>.
- R. Congar and V. Merlin. A characterization of the maximin rule in the context of voting. *Theory and Decision*, 72, 01 2012. doi:10.1007/s11238-010-9229-0.
- V. Conitzer. Eliciting single-peaked preferences using comparison queries. *Journal of Artificial Intelligence Research*, 2009. doi:10.1613/jair.2606.
- V. Conitzer and T. Sandholm. Vote elicitation: Complexity and strategy-proofness. *Proceedings of the National Conference on Artificial Intelligence*, 06 2002. doi:10.48550/arXiv.cs/0205073.
- V. Conitzer and T. Sandholm. Communication complexity of common voting rules. In *Proceedings of EC'05*, 2005. doi:10.1145/1064009.1064018.
- V. Conitzer, T. Walsh, and L. Xia. Dominating manipulations in voting with partial information. In *AAAI'11*, 2011. URL <https://dl.acm.org/doi/10.5555/2900423.2900525>.
- A. n. c. Constitution de 1791, 1791. URL www.conseil-constitutionnel.fr/les-constitutions-dans-l-histoire/constitution-de-1791.
- W. D. Cook and M. Kress. A data envelopment model for aggregating preference rankings. *Management Science*, 36(11):1302–1310, 1990. doi:10.2307/2632607.
- E. P. Cox. The optimal number of response alternatives for a scale: A review. *Journal of Marketing Research*, 17(4):407–422, 1980. doi:10.2307/3150495.
- A. Davenport and J. Kalagnanam. A computational study of the kemeny rule for preference aggregation. In *Proceedings of the 19th National Conference on Artificial Intelligence, AAAI'04*, page 697–702. AAAI Press, 2004. ISBN 0262511835. doi:10.5555/1597148.1597260.
- J. P. Day. Compromise. *Philosophy*, 64(250):471–485, 1989. doi:10.1017/S0031819100044247.
- G. de Clippel, K. Eliaz, and B. Knight. On the selection of arbitrators. *American Economic Review*, 104:3434–3458, 2014. ISSN 0002-8282. doi:10.1257/aer.104.11.3434.
- T. Delemazure and S. Bouveret. Un autre vote résultats 2022, 2022. URL <https://mieuxvoter2022.fr/rapport-premier-tour.pdf>.
- C. Demeulenaere-Douyère. *L'académie des sciences des origines au milieu du XXe siècle*, pages 203–224. Hors collection. Perrin, Paris, 1995. doi:10.3917/perri.insti.1995.01.0203.
- P. Dey and N. Misra. Preference elicitation for single crossing domain. In *Proceedings of IJCAI'16*, 2016. URL <https://dl.acm.org/doi/10.5555/3060621.3060653>.

- P. Dey, N. Misra, and Y. Narahari. Complexity of manipulation with partial information in voting. *Theoretical Computer Science*, 2018. doi:10.1016/j.tcs.2018.03.012.
- G. Doglioni and Z. Ziotti. *Le Cose Notabili Et Maravigliose Della Citta' Di Venetia ... ampliate da Zuanne Zittio*. Gio. Batt. Cestari, 1666. URL <https://books.google.fr/books?id=VZhSAAAAcAAJ>.
- F. Y. Edgeworth. The pure theory of taxation. *The Economic Journal*, 7(25):46–70, 1897. doi:10.2307/2956958.
- E. Elkind and G. Erdélyi. Manipulation under voting rule uncertainty. In *Proceedings AA-MAS'12*, 2012. URL <https://dl.acm.org/doi/10.5555/2343776.2343786>.
- U. Endriss, S. Obraztsova, M. Polukarov, and J. S. Rosenschein. Strategic voting with incomplete information. In *IJCAI'16*, 2016. URL <https://dl.acm.org/doi/10.5555/3060621.3060655>.
- G. Erdélyi, M. R. Fellows, J. Rothe, and L. Schend. Control complexity in bucklin and fallback voting: A theoretical analysis. *Journal of Computer and System Sciences*, 81:632–660, 2015. doi:10.1016/j.jcss.2014.11.002.
- D. Felsenthal and M. Machover. The majority judgment voting procedure: A critical evaluation1. *Homo oeconomicus*, 25, 01 2008. URL <http://eprints.lse.ac.uk/24213/>.
- P. C. Fishburn. Condorcet social choice functions. *SIAM Journal on Applied Mathematics*, 33(3):469–489, 1977. doi:10.2307/2100704.
- P. C. Fishburn. *The Theory of Social Choice*. Princeton University Press, 2015. ISBN 9781400868339. doi:10.1515/9781400868339.
- P. C. Fishburn and S. J. Brams. Paradoxes of preferential voting. *Mathematics Magazine*, 56(4): 207–214, 1983. doi:10.2307/2689808.
- J. Fürnkranz and E. Hüllermeier. *Preference Learning*. Springer Berlin Heidelberg, 2010. ISBN 9783642141256. doi:10.1007/978-3-642-14125-6.
- W. Gaertner. *Primer in Social Choice Theory*. Oxford University Press UK, 2006.
- W. Gaertner and Y. Xu. A general scoring rule. *Mathematical Social Sciences*, 63(3):193–196, 2012. ISSN 0165-4896. doi:<https://doi.org/10.1016/j.mathsocsci.2012.01.006>.
- F. Galton. One vote, one value. *Nature*, 75(1948):414–414, Feb 1907a. doi:10.1038/075414a0.
- F. Galton. Vox populi. *Nature*, 75(1949):450–451, Mar 1907b. doi:10.1038/075450a0.
- W. V. Gehrlein and D. Lepelley. On some limitations of the median voting rule. *Public Choice*, 117(1/2):177–190, 2003. ISSN 00485829, 15737101. doi:10.1023/A:1026147319249.

BIBLIOGRAPHY

- A. Giritligil, M. Sertel, and A. Kara. Does majoritarian approval matter in selecting a social choice rule? an exploratory panel study. *Social Choice and Welfare*, 2005. doi:10.1007/s00355-005-0024-8.
- W. Güth and M. G. Kocher. More than thirty years of ultimatum bargaining experiments: Motives, variations, and a survey of the recent literature. *Journal of Economic Behavior and Organization*, 108:396 – 409, 2014. ISSN 0167-2681. doi:10.1016/j.jebo.2014.06.006.
- G. Hägele and F. Pukelsheim. Llull’s writings on electoral systems. In *Studia Lulliana*, 2001.
- A. Hamlin. Podcast 2012-10-06: Interview with nobel laureate dr. kenneth arrow, 2012. URL <https://electionscience.org/commentary-analysis/voting-theory-podcast-2012-10-06-interview-with-nobel-laureate-dr-kenneth-arrow/>.
- E. Hemaspaandra, H. Spakowski, and J. Vogel. The complexity of kemeny elections. *Theoretical Computer Science*, 349(3):382–391, 2005. ISSN 0304-3975. doi:10.1016/j.tcs.2005.08.031.
- C. Herrero and A. Villar. The three musketeers: four classical solutions to bankruptcy problems. *Mathematical Social Sciences*, 42(3):307–328, 2001. doi:10.1016/S0165-4896(01)00075-0.
- C. G. Hoag and G. H. Hallett. *Proportional representation*. The Macmillan Company, New York, 1926.
- W. H. Holliday and E. Pacuit. Strategic voting under uncertainty about the voting method. *Electronic Proceedings in Theoretical Computer Science*, 297:252–272, jul 2019. doi:10.4204/eptcs.297.17.
- J. R. Irwin, P. Slovic, S. Lichtenstein, and G. H. McClelland. Preference reversals and the measurement of environmental values. *Journal of Risk and Uncertainty*, 6(1):5–18, Jan 1993. ISSN 1573-0476. doi:10.1007/BF01065347.
- E. Kalai. Proportional solutions to bargaining situations: interpersonal utility comparisons. *Econometrica: Journal of the Econometric Society*, 45(7):1623–1630, 1977. doi:10.2307/1913954.
- E. Kalai and M. Smorodinsky. Other solutions to nash’s bargaining problem. *Econometrica: Journal of the Econometric Society*, 43(3), 1975. doi:10.2307/1914280.
- M. Kalech, S. Kraus, G. A. Kaminka, and C. V. Goldman. Practical voting rules with partial information. *Autonomous Agents and Multi-Agent Systems*, 2011. doi:10.1007/s10458-010-9133-6.
- H. Keller. Electoral systems and conceptions of community in italian communes. (twelfth–fourteenth centuries). *Revue française de science politique*, 64(6):1083–1107, 2014. doi:https://doi.org/10.3917/rfsp.646.1083.
- Ö. Kibris and M. R. Sertel. Bargaining over a finite set of alternatives. *Social Choice and Welfare*, 28(3):421–437, 2007. doi:10.1007/s00355-006-0178-z.

- K. Konczak and J. Lang. Voting procedures with incomplete preferences. In *IJCAI'05*, 2005.
- P. Kouvelis and G. Yu. *Robust Discrete Optimization and Its Applications*. Kluwer, 1997.
- J. Lang. Logical preference representation and combinatorial vote. *Annals of Mathematics and Artificial Intelligence*, 42(1):37–71, Sep 2004. ISSN 1573-7470. doi:10.1023/B:AMAI.0000034522.25580.09.
- J. Lang. Collective decision making under incomplete knowledge: Possible and necessary solutions. In C. Bessiere, editor, *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 4885–4891. International Joint Conferences on Artificial Intelligence Organization, 7 2020. doi:10.24963/ijcai.2020/680. Survey track.
- LaPrimaire.org, 2016. URL <https://laprimaire.org/>.
- J.-F. Laslier. The strange "majority judgment". unpublished, 2018. URL <https://hal.archives-ouvertes.fr/hal-01965227>.
- J.-F. Laslier, M. Nunez, and M. Remzi Sanver. A solution to the two-person implementation problem. working paper or preprint, May 2020. URL <https://hal.archives-ouvertes.fr/halshs-02173504>.
- O. Lev, R. Meir, S. Obraztsova, and M. Polukarov. Heuristic voting as ordinal dominance strategies. *Proceedings of the AAAI'19*, 2019. doi:10.1609/aaai.v33i01.33012077.
- S. Lhuillier. Examen du mode d'élection proposé à la convention nationale de france en février 1793, et adopté à genève. *Mathématiques et Sciences Humaines*, 54:7–24, 1793. URL <http://eudml.org/doc/94179>.
- P. Lichtenstein, Sarahand Slovic. *The Construction of Preference*. Cambridge University Press, 2006. doi:10.1017/CBO9780511618031.
- M. Lines. Approval voting and strategy analysis: A venetian example. *Theory and Decision*, 20:155–172, 1986. doi:10.1007/BF00135090.
- T.-Y. Liu. Learning to rank for information retrieval. *Foundations and Trends in Information Retrieval*, 3(3):225–331, 2009. ISSN 1554-0669. doi:10.1561/15000000016.
- B. Llamazares. Ranking candidates through convex sequences of variable weights. *Group Decision and Negotiation*, 2016. doi:10.1007/s10726-015-9452-8.
- B. Llamazares and T. Peña. Aggregating preferences rankings with variable weights. *European Journal of Operational Research*, 2013. doi:10.1016/j.ejor.2013.04.013.
- T. Lu and C. Boutilier. Robust approximation and incremental elicitation in voting protocols. In *Proceedings of IJCAI'11*, 2011a. doi:10.5591/978-1-57735-516-8/IJCAI11-058.

BIBLIOGRAPHY

- T. Lu and C. Boutilier. Vote elicitation with probabilistic preference models: Empirical estimation and cost tradeoffs. In R. I. Brafman, F. S. Roberts, and A. Tsoukiàs, editors, *Algorithmic Decision Theory*, pages 135–149. Springer Berlin Heidelberg, 2011b. ISBN 978-3-642-24873-3. doi:10.1007/978-3-642-24873-3_11.
- D. Luban. Bargaining and compromise: Recent work on negotiation and informal justice. *Philosophy and Public Affairs*, 14(4):397–416, 1985. doi:10.2307/2265340.
- G. R. Maio, N. J. Roese, C. Seligman, and A. Katz. Rankings, ratings, and the measurement of values: Evidence for the superior validity of ratings. *Basic and Applied Social Psychology*, 18(2):171–181, 1996. doi:10.1207/s15324834basp1802_4.
- M. Mariotti. Nash bargaining theory when the number of alternatives can be finite. *Social Choice and Welfare*, 15:413–421, 1998. ISSN 0176-1714. doi:10.1007/s003550050114.
- A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford University Press, 1995. ISBN 978-0-195-10268-0.
- N. Mattei and T. Walsh. Preflib: A library of preference data PREFLIB.ORG. In *Proceedings of ADT 2013*, 2013. URL <https://www.preflib.org/>.
- I. McLean. The borda and condorcet principles: Three medieval applications. *Social Choice and Welfare*, 7(2):99–108, 1990. ISSN 1432-217X. doi:10.1007/BF01560577.
- I. McLean and F. Hewitt. *Condorcet: Foundations of Social Choice and Political Theory*. Edward Elgar Publishing, Incorporated, 1994. ISBN 9781781008119. URL <https://books.google.fr/books?id=e7IXFcvyEMUC>.
- I. McLean and A. Urken, editors. *Classics of Social Choice*. University of Michigan Press, 1995. doi:10.3998/mpub.12736.
- V. Merlin, I. Özkal Sanver, and M. R. Sanver. Compromise rules revisited. *Group Decision and Negotiation*, 28:63–78, 2019. ISSN 0926-2644. doi:10.1007/s10726-018-9598-2.
- Merriam-Webster. Arbitration. In *Merriam-Webster.com Dictionary*. Encyclopædia Britannica, Inc., 2022a. URL <https://www.merriam-webster.com/dictionary/arbitration>.
- Merriam-Webster. Compromise. In *Merriam-Webster.com Dictionary*. Encyclopædia Britannica, Inc., 2022b. URL <https://www.merriam-webster.com/dictionary/compromise>.
- MieuxVoter, 2022a. URL www.mieuxvoter.fr.
- MieuxVoter. Mieux voter 2022 que retenir de notre consultation sur les candidats à l’élection présidentielle ?, 2022b. URL <https://mieuxvoter2022.fr/rapport-premier-tour.pdf>.
- J. S. Mill. *Principles of political economy*, volume 2. London : John W. Parker, West Strand, 1849. URL <https://books.google.fr/books?id=TdUPAAAYAAJ>.

- H. Moulin. *The strategy of social choice*. Advanced textbooks in economics ; 18. North-Holland Pub. Co., 1983. ISBN 0444863710. doi:10.1016/C2013-0-11866-1.
- V. Mousseau. A general framework for constructive learning preference elicitation in multiple criteria decision aid. working paper or preprint, 2005. URL https://hal.archives-ouvertes.fr/hal-00958059/file/cahier_223.pdf.
- M. Mowbray and D. Gollmann. Electing the doge of venice: Analysis of a 13th century protocol. In *IEEE Computer Security Foundations Symposium*, pages 295 – 310, 08 2007. ISBN 0-7695-2819-8. doi:10.1109/CSF.2007.21.
- L. Naamani-Dery, I. Golan, M. Kalech, and L. Rokach. Preference elicitation for group decisions using the borda voting rule. *Group Decision and Negotiation*, 2015. doi:10.1007/s10726-015-9427-9.
- R. Nagahisa and M. Tanaka. An axiomatization of the kalai-smorodinsky solution when the feasible sets can be finite. *Social Choice and Welfare*, 19:751–761, 2002. ISSN 0176-1714. doi:10.1007/s003550100150.
- B. Napolitano, O. Cailloux, and P. Viappiani. Simultaneous elicitation of scoring rule and agent preferences for robust winner determination. In *Proceedings of Algorithmic Decision Theory - 7th International Conference, ADT 2021*, 2021. doi:10.1007/978-3-030-87756-9_4.
- J. F. Nash. The bargaining problem. *Econometrica: Journal of the Econometric Society*, 18: 155, 1950. ISSN 0012-9682. doi:10.2307/1907266.
- K. Nehring and M. Pivato. The median rule in judgement aggregation. *Economic Theory*, 73 (4):1051–1100, Jun 2022. ISSN 1432-0479. doi:10.1007/s00199-021-01348-7.
- S. Nitzan. *Collective Preference and Choice*. Cambridge University Press, 2009. doi:10.1017/CBO9780511803871.
- H. Nurmi. *Voting paradoxes and how to deal with them*. Springer, 1999. ISBN 978-3-540-66236-5. doi:10.1007/978-3-662-03782-9.
- I. Özkal-Sanver and M. R. Sanver. Efficiency in the degree of compromise: A new axiom for social choice. *Group Decision and Negotiation*, 13:375–380, 2004. ISSN 0926-2644. doi:10.1023/b:grup.0000042925.01972.ad.
- M. S. Pini, F. Rossi, K. B. Venable, and T. Walsh. Incompleteness and incomparability in preference aggregation. In *Proceedings of IJCAI’07*, 2007. doi:10.1016/j.artint.2010.11.009.
- M. S. Pini, F. Rossi, K. B. Venable, and T. Walsh. Aggregating partially ordered preferences. *Journal of Logic and Computation*, 2009. doi:10.1093/logcom/exn012.
- M. Pivato. Formal utilitarianism and range voting. *Mathematical Social Sciences*, 67:50–56, 2014. ISSN 0165-4896. doi:<https://doi.org/10.1016/j.mathsocsci.2013.08.002>.
- K. Popper. *The Open Society and its Enemies*. Routledge, 1945.

BIBLIOGRAPHY

- C. C. Preston and A. M. Colman. Optimal number of response categories in rating scales: reliability, validity, discriminating power, and respondent preferences. *Acta Psychologica*, 104(1):1–15, 2000. doi:[https://doi.org/10.1016/S0001-6918\(99\)00050-5](https://doi.org/10.1016/S0001-6918(99)00050-5).
- A. D. Procaccia. A note on the query complexity of the condorcet winner problem. *Information Processing Letters*, 108(6):390–393, 2008. ISSN 0020-0190. doi:[10.1016/j.ipl.2008.07.012](https://doi.org/10.1016/j.ipl.2008.07.012).
- J. Rawls. Justice as fairness. *The Philosophical Review*, 67(2):164–194, 1958. doi:[10.2307/2182612](https://doi.org/10.2307/2182612).
- J. Rawls. *Philosophy, Politics, and Society: Third Series; a Collection*, chapter 3, pages 58–82. A Blackwell paperback. Blackwell, 1967. ISBN 9780631048909. URL <https://books.google.fr/books?id=Ii-GAAAAMAAJ>.
- A. Reijngoud and U. Endriss. Voter response to iterated poll information. In *Proceedings of AAMAS’12*, 2012. URL <https://dl.acm.org/doi/10.5555/2343776.2343787>.
- J. Rothe, H. Spakowski, and J. Vogel. Exact complexity of the winner problem for young elections. *Theory of Computing Systems*, 36(4):375–386, Aug 2003. ISSN 1433-0490. doi:[10.1007/s00224-002-1093-z](https://doi.org/10.1007/s00224-002-1093-z).
- T. L. Saaty. Axiomatic foundation of the analytic hierarchy process. *Management Science*, 32(7):841–855, 1986. doi:[10.1287/mnsc.32.7.841](https://doi.org/10.1287/mnsc.32.7.841).
- A. A. Salo and R. P. Hämmäläinen. Preference ratios in multiattribute evaluation (PRIME)—elicitation and decision procedures under incomplete information. *IEEE Trans. on Systems, Man and Cybernetics*, 2001. doi:[10.1109/3468.983411](https://doi.org/10.1109/3468.983411).
- L. J. Savage. *The Foundations of Statistics*. Wiley, 1954.
- T. Schwartz. Rationality and the myth of the maximum. *Noûs*, 6(2):97–117, 1972. doi:[10.2307/2216143](https://doi.org/10.2307/2216143).
- A. Sen. The impossibility of a paretian liberal. *Journal of Political Economy*, 78(1):152–157, 1970. doi:[10.2307/1829633](https://doi.org/10.2307/1829633).
- A. Sen. *Rationality and Freedom*. Belknap Press of Harvard University Press, 2004. ISBN 9780674013513. doi:[10.2307/j.ctv1dv0td8](https://doi.org/10.2307/j.ctv1dv0td8).
- A. Sen. *Collective Choice and Social Welfare: Expanded Edition*. Penguin Books Limited, 2017. ISBN 9780241244609. URL <https://books.google.fr/books?id=9I4sDAAQBAJ>.
- M. R. Sertel. Lecture notes on microeconomics (unpublished). *Boğaziçi University, İstanbul*, 1986.
- M. R. Sertel and B. Yılmaz. The majoritarian compromise is majoritarian-optimal and subgame-perfect implementable. *Social Choice and Welfare*, 16:615–627, 1999. ISSN 0176-1714. doi:[10.1007/s003550050164](https://doi.org/10.1007/s003550050164).

- B. Settles. Active learning literature survey. *CS Technical Reports*, 2009. URL <https://minds.wisconsin.edu/handle/1793/60660>.
- W. D. Smith. Range voting, 2000. URL <https://www.rangevoting.org/WarrenSmithPages/homepage/rangevote.pdf>.
- E. I. Sparling and S. Sen. Rating: How difficult is it? In *Proceedings of the Fifth ACM Conference on Recommender Systems*, RecSys '11, page 149–156, New York, NY, USA, 2011. Association for Computing Machinery. ISBN 9781450306836. doi:10.1145/2043932.2043961.
- Y. Sprumont. Intermediate preferences and rawlsian arbitration rules. *Social Choice and Welfare*, 10(1):1–15, 1993. ISSN 01761714, 1432217X. doi:10.1007/BF00187429.
- W. E. Stein, P. J. Mizzi, and R. C. Pfaffenberger. A stochastic dominance analysis of ranked voting systems with scoring. *EJOR*, 1994. doi:10.1016/0377-2217(94)90205-4.
- A. D. Taylor. *Social Choice and the Mathematics of Manipulation*. Cambridge University Press, 2005. doi:10.1017/CBO9780511614316.
- L. L. Thurstone. A law of comparative judgment. *Psychological Review*, 34(4):273–286, 1927. doi:10.1037/h0070288.
- A. Tversky and D. Kahneman. Rational choice and the framing of decisions. *The Journal of Business*, 59(4):S251–S278, 1986. doi:10.2307/2352759.
- S. Uckelman and J. Uckelman. Strategy and manipulation in medieval elections. In *Proceedings of the 3rd International Workshop on Computational Social Choice (COMSOC)*, 2010.
- A. B. Urken. *The Origins of Social Choice Theory*, pages 725–727. Springer US, 2004. doi:10.1007/978-0-306-47828-4_148.
- E. M. Varloot and R. Laraki. Level-strategyproof belief aggregation mechanisms. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, EC '22, page 335–369, New York, NY, USA, 2022. Association for Computing Machinery. doi:10.1145/3490486.3538309.
- P. Viappiani. Positional scoring rules with uncertain weights. In *Scalable Uncertainty Management*, 2018. doi:10.1007/978-3-030-00461-3_21.
- T. Walsh. Uncertainty in preference elicitation and aggregation. In *Proceedings of AAAI'07*, 2007. URL <https://dl.acm.org/doi/10.5555/1619645.1619648>.
- T. Walsh. Complexity of terminating preference elicitation. In *Proceedings of AAMAS'08*, 2008. URL <https://dl.acm.org/doi/abs/10.5555/1402298.1402357>.
- A. M. Wolfson. The ballot and other forms of voting in the italian communes. *The American Historical Review*, 5(1):1–21, 1899. doi:10.2307/1832957.

BIBLIOGRAPHY

- X. Wu, L. Xiao, Y. Sun, J. Zhang, T. Ma, and L. He. A survey of human-in-the-loop for machine learning. *Future Generation Computer Systems*, 135:364–381, 2022. ISSN 0167-739X. doi:<https://doi.org/10.1016/j.future.2022.05.014>.
- L. Xia and V. Conitzer. Determining possible and necessary winners under common voting rules given partial orders. In *Proceedings of AAAI’08*, 2008. doi:10.1613/jair.3186.
- L. Xia and V. Conitzer. Determining possible and necessary winners given partial orders. *J. Artif. Intell. Res. (JAIR)*, 41, 2011. doi:10.1613/jair.3186.
- H. Young. An axiomatization of borda’s rule. *Journal of Economic Theory*, 9(1):43–52, 1974. ISSN 0022-0531. doi:[https://doi.org/10.1016/0022-0531\(74\)90073-8](https://doi.org/10.1016/0022-0531(74)90073-8).
- H. P. Young. Social choice scoring functions. *Siam Journal on Applied Mathematics*, 28:824–838, 1975. doi:10.1137/0128067.
- H. P. Young. Progressive taxation and the equal sacrifice principle. *Journal of Public Economics*, 32(2):203–214, 1987. ISSN 0047-2727. doi:10.1016/0047-2727(87)90012-0.
- P. L. Yu. A class of solutions for group decision problems. *Management Science*, 19(8):936–946, 1973. doi:10.2307/2629125.
- A. Zahid. Majority judgment theory and paradoxical results. In ., 2009.
- M. A. Zahid and H. de Swart. The borda majority count. *Information Sciences*, 295:429–440, 2015. ISSN 0020-0255. doi:<https://doi.org/10.1016/j.ins.2014.10.044>.
- F. Zappalá. Memoria storica dell’arbitrato. *Meritum, Revista de Direito da Universidade FUMEC*, 2018.
- R. Zimmermann. *The Law of Obligations: Roman Foundations of the Civilian Tradition*. Clarendon paperbacks. Oxford University Press, 1996. ISBN 9780198764267. URL https://books.google.fr/books?id=iFIT_NsmE7MC.

RÉSUMÉ LONG EN FRANÇAIS

8.1 Introduction

La prise de décision collective est un processus dans lequel les préférences individuelles sont agrégées pour former un choix collectif unique. Cette définition large permet à de nombreux cas, qui semblent très différents les uns des autres, d'entrer dans cette catégorie. Parmi les exemples, citons les *problèmes d'allocation équitable*—qui traite du problème de la répartition équitable de certaines ressources entre des individus qui ont des intérêts différents, par exemple l'allocation de maisons, la création d'un horaire de travail, etc.—*problèmes d'appariement*—qui traite du problème de l'appariement d'individus de deux groupes distincts en tenant compte de leurs préférences, par exemple, des étudiants aux écoles, des locataires aux maisons etc.—et *problèmes d'agrégation de jugements*—qui tentent de regrouper les croyances de différents individus en un jugement qui reflète la société en tant qu'entité unique. Dans cette thèse, nous nous concentrerons sur un autre exemple de prise de décision collective : le *vote*. Étant donné un ensemble d'individus qui expriment leurs préférences sur un ensemble donné de candidats, comment choisir le *meilleur* candidat pour le groupe ? Il s'agit d'un dilemme très ancien qui a été affronté de multiples fois au cours des siècles et que nous allons aborder tout au long de ce manuscrit. Comme nous allons le voir, la réponse dépend de nombreux facteurs, tels que les informations dont nous disposons—connaissons-nous les préférences de tous les électeurs par rapport à tous les candidats ?—et les propriétés que nous voulons que le résultat satisfasse—qu'entendons-nous par le meilleur candidat ? Qui décide de ce qui est le meilleur ?

Une méthodologie importante dans notre travail consiste à définir les propriétés souhaitables que nous voulons qu'une règle satisfasse. Cela nous permet, d'une part, de diviser les règles d'agrégation en classes de méthodes satisfaisant les mêmes propriétés, et, d'autre part, de faire le processus inverse : partir d'une méthode déjà définie pour aider le décideur à comprendre ses propriétés. Le manque d'information et l'aide aux décideurs sont deux des thèmes explorés au cours de ce travail.

Nous avons mentionné l'analyse des processus d'agrégation par la définition de desiderata, mais cette approche axiomatique est très récente et n'a commencé qu'après la publication de la thèse de doctorat de Kenneth Arrow [1951]. La théorie des processus de prise de décision collective est cependant beaucoup plus ancienne et, comme l'a écrit McLean [1990], “la théorie du vote est connue pour avoir été découverte trois fois et perdue deux fois”. En partant des premières traces des processus de décision, nous décrivons le début et l'évolution de la théorie du choix social.

Questions de recherche et contributions. En étudiant l’histoire des méthodes d’agrégation, nous avons réalisé qu’à partir d’un large éventail de procédures, nous pouvons identifier des classes dont les éléments satisfont des propriétés communes. Étant donné une méthode, nous pouvons analyser les axiomes qu’elle satisfait. L’intérêt de cette démarche est que nous pouvons aussi faire l’inverse : étant donné certaines préférences sur les propriétés que la règle doit respecter, nous pouvons aider le décideur à définir la méthode d’agrégation qu’il souhaite. C’est précisément l’un des points sur lesquels nous allons nous concentrer :

- comment pouvons-nous aider le décideur à définir formellement une règle électorale sur la base de préférences génériques sur ses propriétés ?

Dans Chapter 4, nous étudions une situation dans laquelle un comité de non-experts doit décider comment agréger les préférences des électeurs. Supposons que le comité attribue un score à chaque position dans l’ordre des préférences des électeurs. Ainsi, par exemple, le premier classé obtient 10 points, le second 5, etc. Imaginez maintenant que le comité souhaite que le score du meilleur choix soit “beaucoup plus élevé” que celui du troisième meilleur choix. Comment traduire ce “beaucoup plus” en une règle de vote ? Qu’est-ce que cela signifie exactement ? Dans cette thèse, nous nous concentrons sur une classe particulière de règles, les *Règles de notation positionnelle* —que nous décrivons en détail dans Chapter 2 —et développons une méthode pour poser des questions au comité en termes de choix des gagnants dans un exemple de profil. Ainsi, nous transformons une question complexe impliquant des différences de poids entre des positions consécutives en des questions auxquelles un comité humain peut plus facilement répondre, telles que “qui devrait gagner dans ce profil” ? De la réponse à cette question, nous déduisons la réponse à la question initiale.

Un autre problème important lors du choix d’une règle de vote est son coût. C’est un point que nous retrouvons également dans notre première question de recherche. Dans ce cas, il est représenté par le coût cognitif pour les non-experts de formaliser une règle de vote sur la base de certaines préférences génériques. Mais, comme nous le savons, le coût est également lié aux électeurs lorsqu’ils doivent ordonner un grand ensemble de préférences, et au calcul de la transmission de toutes ces informations. Cette prise de conscience nous amène à étudier les stratégies d’éllicitation des préférences des électeurs :

- comment pouvons-nous acquérir les informations les plus pertinentes au moindre coût ?

Nous étudions cette question pour différentes règles de vote (règles de notation positionnelle et jugement majoritaire) dans Chapters 4 and 6. En particulier, dans Chapter 4, nous combinons cette question de recherche avec la première et nous essayons de comprendre si, dans un contexte d’information nulle, l’imbrication de l’éllicitation des préférences des électeurs et de l’éllicitation de la règle de vote donne de meilleurs résultats qu’une approche plus linéaire. Dans Chapter 6, nous étudions des scénarios concrets tirés de situations de vote réelles, dans lesquels les préférences des électeurs sont obtenues en posant un nombre donné de questions aléatoires aux électeurs. Entre autres, nous cherchons comment identifier le nombre minimum de questions à poser pour que la probabilité d’obtenir le gagnant, en utilisant la règle du jugement majoritaire, soit élevée.

Nous avons commencé à nous demander comment différentes situations pouvaient justifier l’utilisation de différentes règles de vote qui cherchent à atteindre le même objectif. Dans la

littérature, diverses notions de compromis ont été proposées, et avec elles diverses règles de vote qui tentent plus ou moins de mettre en œuvre ces définitions. Lorsque nous parlons de compromis, ce que nous entendons dépend beaucoup du contexte. Lorsque nous voulons décider où dîner, nous sommes probablement d'accord pour que les électeurs essaient de trouver un terrain d'entente, même si le résultat ne s'avère pas être un compromis. C'est le cas avec certaines règles de vote dont nous parlerons plus tard, comme le Fallback Bargaining, où il peut arriver que le résultat soit le meilleur choix pour certains électeurs, mais qu'il déplaie à beaucoup d'autres. Cependant, lorsque nous avons affaire à des contextes où l'égalitarisme —que nous interpréterons comme le fait que tout le monde concède également—est une préoccupation majeure, cette notion de compromis n'est plus acceptable. Sur ces bases, nous développons les fondements de notre concept de compromis :

- comment pouvons-nous définir la notion de compromis où chaque électeur concède de manière aussi égale que possible ?

Dans Chapter 7, nous abordons cette question et analysons quelles règles existantes, le cas échéant, répondent à cette nouvelle définition. Nous verrons que différentes définitions ont un sens dans différents contextes et que l'une n'est pas nécessairement meilleure que l'autre.

Organisation de la thèse Chapter 2 rassemble les notions importantes qui sont à la base des thèmes abordés dans cette thèse et qui reviendront de manière récurrente dans les chapitres suivants. Nous y décrivons la différence entre le vote avec différents types de scrutins : les scrutins classés, que nous utiliserons dans Chapters 4 and 7, et les scrutins notés, que nous considérerons dans Chapter 6. Nous allons décrire les règles de vote que nous utiliserons dans nos contributions et certains des axiomes qui les caractérisent.

Dans Chapter 3, en utilisant la même approche que dans cette introduction, nous fouillerons dans le passé pour positionner notre travail. En particulier, dans Section 3.1, en partant de la signification du compromis, nous retraçons son utilisation à travers l'histoire. Nous analysons les définitions qui ont été données des règles de compromis dans la théorie du choix social et soulignons comment certaines d'entre elles ne représentent pas certaines idées du compromis. Nous montrons comment notre approche s'inscrit dans cette littérature. Constatant que la règle du jugement majoritaire est considérée comme une forme de compromis dans la littérature, nous approfondirons cette règle dans Section 3.2. Nous étudierons son introduction, ses utilisations et aussi ses critiques. Comme dernier sujet de ce chapitre, nous couvrons dans Section 3.3 un autre aspect important de cette thèse : l'éllicitation des préférences. En effet, si jusqu'à présent on a supposé connaître les préférences des électeurs et la procédure de vote, cela ne peut être considéré comme acquis. L'éllicitation des préférences est un problème bien étudié et nous décrivons les différentes approches par lesquelles il a été abordé dans la littérature. En outre, nous décrivons comment il a été abordé dans différents domaines, notamment dans la littérature sur l'aide à la décision et l'apprentissage automatique, et les similitudes avec notre approche.

Part II inclut nos contributions. Plus précisément, Chapter 4 est un article publié dans une conférence internationale [Napolitano et al., 2021] et Chapter 7 est un article publié dans une revue [Cailloux et al., 2022]. Chapter 5 accompagne Chapter 4 en décrivant le code fourni à l'appui de la contribution. Enfin, Chapter 6 est une contribution originale, non encore publiée,

dont le but est d'étudier les conséquences du processus d'élicitation dans une situation de vote utilisant le jugement majoritaire.

Pour conclure, Chapter 8 fournit un résumé des contributions de cette thèse et quelques perspectives pour les directions futures.

Contributions

8.2 Simultaneous Elicitation of PSR and Agent Preferences

Le cadre classique du choix social suppose une information complète sur les ordres de préférence de tous les électeurs, et sur le mécanisme de vote lui-même. Mais dans quelle mesure cette hypothèse est-elle raisonnable ? Si l'ensemble des alternatives est très vaste, il n'est pas raisonnable d'attendre des électeurs qu'ils fournissent un ordre complet de leurs préférences. D'abord d'un point de vue cognitif : des études de psychologie expérimentale décrivent comment, dans de nombreuses situations, nos préférences ne sont pas prédéfinies, mais que nous les construisons seulement au moment où nous devons les exprimer [Lichtenstein, 2006]. Demander aux électeurs d'exprimer leurs préférences lorsque l'éventail de choix est très large peut entraîner divers problèmes, tels que la confusion et une faible participation. Deuxièmement, mais non moins important, cela représente un énorme coût de communication. Conitzer and Sandholm [2005] ont étudié la complexité de communication de certaines des règles de vote les plus courantes et, pour chacune d'entre elles, ont fourni une limite supérieure et inférieure sur le nombre de bits d'information que les électeurs doivent communiquer avant que la règle puisse sélectionner un gagnant. D'autres mesures existent, comme le nombre de questions auxquelles les électeurs doivent répondre avant que la règle puisse déterminer un gagnant. Cette approche a été utilisée par Procaccia [2008] pour déterminer le nombre minimum de questions permettant de sélectionner le gagnant de Condorcet. Dans les élections présidentielles, par exemple, nous avons tendance à préférer la précision au coût, mais ce n'est pas vrai dans toutes les applications. Lorsqu'il s'agit de choisir un restaurant pour dîner avec des amis, nous nous contenterons probablement d'un gagnant approximatif plutôt que de devoir passer toute la soirée à classer tous les restaurants de la ville.

Cette observation a motivé plusieurs travaux supposant des ordres de préférence partiels : un des premiers travaux est celui de Conitzer and Sandholm [2005] qui a étudié la complexité de la communication lors de l'utilisation de différentes règles de vote ; Konczak and Lang [2005] a étudié le calcul des gagnants possibles et nécessaires pour différentes règles de vote ; Xia and Conitzer [2008] ont ensuite montré que, si l'identification d'un co-gagnant nécessaire dans les règles de notation est polynomiale, la détermination des co-gagnants possibles est NP-dure ; D'autres résultats de complexité ont été donnés par Walsh [2007] et Pini et al. [2007].

Puisque dans de nombreuses situations pratiques, il y aurait trop de gagnants possibles mais pas de gagnants nécessaires, plusieurs travaux ont abordé le problème de l'élicitation des préférences des agents en utilisant une variété d'approches (regret minimax, méthodes bayésiennes, etc.) dans le but de converger vers un gagnant nécessaire : [Naamani-Dery et al., 2015, Kalech et al., 2011, Lu and Boutilier, 2011a, Pini et al., 2009, Benabbou et al., 2016, Dey and

Misra, 2016]. Parmi ceux-ci, Walsh [2008] et Conitzer [2009] ont analysé quand arrêter le processus d'élicitation.

Une deuxième préoccupation concerne la capacité de la personne ou l'organisation qui supervise le processus de vote à fournir une définition précise de la règle de vote, ce qui suggère l'assouplissement de la deuxième hypothèse. En effet, il est souvent difficile pour des non-experts de formaliser une règle de vote sur la base de quelques préférences génériques sur une méthode d'agrégation souhaitée. Nous donnons ici deux exemples de telles situations.

Considérons, comme premier exemple, une commission qui est sur le point d'embaucher un nouvel employé dont les performances sont évaluées par plusieurs experts. Les membres de la commission peuvent ne pas avoir de règle de vote en tête au début du processus et ne pas vouloir se mettre d'accord sur une règle de vote spécifique. Cependant, ils pourraient être disposés à répondre à quelques questions exigeant de sélectionner le gagnant parmi des profils spécifiques.

Considérons, comme deuxième exemple, le processus d'évaluation d'une conférence où le meilleur article doit être élu. Les agents expriment leurs préférences sur les articles qu'ils ont évalués, mais ils ne sont pas conscients de la règle de vote que le président du programme appliquera lorsqu'il les regroupera. Néanmoins, les évaluateurs sont toujours prêts à participer au processus. De plus, le CP peut ne pas avoir de règle de vote spécifique en tête, et il lui sera difficile de fournir un vecteur de notation précis si on le lui demande. Peut-être croit-elle fermement que le fait d'être classé une fois en première position a "beaucoup plus" de valeur que le fait d'être classé deux fois en deuxième position, mais ne sait pas exactement de combien (bien qu'elle puisse juger des cas d'exemple).

Nous nous concentrons sur les règles de notation positionnelle avec des poids convexes, qui sont une méthode particulièrement courante utilisée pour agréger les classements. Nous développons des méthodes, basées sur la notion de regret minimax, pour déterminer un "gagnant" robuste en cas d'incertitude à la fois sur la règle de vote et sur les préférences des agents. Nous fournissons des méthodes d'élicitation incrémentales qui, à chaque étape de l'élicitation, interrogent l'un des agents ou le président, et nous discutons de plusieurs heuristiques pour choisir des questions qui réduisent rapidement le regret. Les réponses aux questions sont encodées sous forme de contraintes ; les questions aux agents sont des comparaisons entre des paires d'alternatives tandis que les questions au président demandent de sélectionner un gagnant parmi un profil synthétique.

Conclusions. Notre approche est évaluée sur des simulations avec des ensembles de données synthétiques et réelles où la règle de vote et les préférences de l'agent sont initialement inconnues du système et progressivement révélées par des questions. Nous supposons que le président est humain, donc capable de répondre à des questions sur un nombre limité d'alternatives, et nous nous concentrons donc sur des situations de choix social à petite échelle. Nous comparons l'efficacité de plusieurs stratégies de questionnement basées sur la connaissance actuelle de la règle et des préférences. Pour résumer nos contributions :


1) Nous fournissons un nouveau mécanisme pour éliciter une règle de vote en traduisant des questions abstraites sur les poids en un choix d'une alternative étant donné un profil concret. Supposons, par exemple, que nous voulions poser la question suivante à la commission : $w_2 - w_3 \geq 2(w_3 - w_4)$. Cela signifie que nous voulons savoir si la différence de poids entre la

Figure 8.1: Profil représentant une question posée à la commission sous forme étendue (a) et compacte (b).

										1	2	3	3
<i>c</i>	<i>d</i>	<i>d</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>b</i>	<i>b</i>	<i>b</i>		<i>c</i>	<i>d</i>	<i>a</i>	<i>b</i>
<i>a</i>	<i>c</i>	<i>c</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>a</i>	<i>a</i>	<i>a</i>		<i>a</i>	<i>c</i>	<i>b</i>	<i>a</i>
<i>b</i>	<i>b</i>	<i>b</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>d</i>	<i>d</i>	<i>d</i>		<i>b</i>	<i>b</i>	<i>c</i>	<i>d</i>
<i>d</i>	<i>a</i>	<i>a</i>	<i>d</i>	<i>d</i>	<i>d</i>	<i>c</i>	<i>c</i>	<i>c</i>		<i>d</i>	<i>a</i>	<i>d</i>	<i>c</i>

(a) (b)

Figure 8.2: MMR moyen (normalisé par n) après k questions avec la stratégie pessimiste pour différents ensembles de données.




deuxième et la troisième position est supérieure ou égale au double de la différence entre la troisième et la quatrième position. Nous montrons le profil de la figure 8.1a à la commission et demandons qui devrait gagner (chaque colonne est la préférence d'un agent). Les deux a et b ont des scores supérieurs à c et d pour tous les poids convexes, donc soit a soit b seront choisis selon notre hypothèse ; et $s(a) \geq s(b) \Leftrightarrow w_2 + 2w_4 \geq 3w_3$. La figure 8.1b représente le même profil en utilisant une vue compressée, les chiffres en gras indiquant le nombre d'agents ayant la préférence dans la colonne correspondante.

2) Nous montrons qu'avec notre méthode d'éllicitation, il est possible d'atteindre un faible regret avec un nombre raisonnable de questions, voir Figure 8.2

3) Nous présentons des stratégies d'éllicitation qui permettent d'obtenir de bons résultats en un temps de calcul raisonnable, voir Figure 8.3.

4) Nous montrons que pour la classe de règles considérée, il suffit de poser quelques questions à la commission pour obtenir un faible regret, voir Table 8.1.

Figure 8.3: MMR moyen dans des problèmes de taille (5, 10) après k questions.Table 8.1: MMR moyen dans des problèmes de taille (10, 20) après 500 questions, parmi lesquelles q_c à la commission.

q_c	2 ph. ca $\pm$ sd	2 ph. ac $\pm$ sd
0	0.6 ± 0.6	0.6 ± 0.6
15	0.5 ± 0.5	0.6 ± 0.6
30	0.3 ± 0.5	0.4 ± 0.5
50	0.1 ± 0.2	0.1 ± 0.3
100	0.4 ± 0.6	0.2 ± 0.5
200	2.1 ± 1.6	2.4 ± 1.2
300	5.8 ± 1.7	6.7 ± 1.5
400	11.3 ± 1.1	11.9 ± 1.2
500	20.0 ± 0.0	20.0 ± 0.0

8.3 Majority Judgment winner determination under incomplete information

Un exemple de compromis peut être représenté par Majority Judgment, un système de vote où les électeurs attribuent des notes aux candidats en utilisant une échelle ordinale. L'évaluation des candidats au lieu de leur classement permet d'obtenir plus d'informations grâce à une plus grande expressivité. Dans le système MJ, chaque électeur évalue, ou juge, chaque candidat et le gagnant est le candidat ayant la médiane la plus élevée des notes reçues. Cette méthode a été introduite par Balinski and Laraki [2007] dans la période la plus récente de l'histoire du choix social. Pourtant, elle a attiré l'attention croissante des associations et des partis politiques français qui ont commencé à utiliser MJ pour des décisions internes ou des élections locales.

De nombreux observateurs ont décrit la note médiane comme le niveau le plus élevé auquel un candidat obtient le soutien de la majorité des électeurs. En d'autres termes, en commençant par le niveau le plus élevé d , nous vérifions si la majorité des électeurs a attribué au moins d à une certaine alternative a . Si ce n'est pas le cas, nous descendons dans l'échelle de notation jusqu'à trouver un niveau d^* où un candidat a^* satisfait la moitié de la population. La note d^* est alors la médiane des notes a^* , et, comme c'est le premier niveau auquel on s'est arrêté, elle correspond à la meilleure médiane possible. Cette méthode a été redécouverte plusieurs fois et proposée sous le nom de règle de Bucklin [Hoag and Hallett, 1926], Compromis Majoritaire [Sertel, 1986, Sertel and Yilmaz, 1999] et "q-approval fallback bargaining" [Brams and Kilgour, 2001]. De plus, notez que lorsque le nombre de grades est égal à deux (approuver, désapprouver), cette méthode se réduit au vote par approbation.

En France, le MJ a été adopté par un nombre de plus en plus important d'associations et de partis politiques dont : Le Parti Pirate, Génération(s), LaPrimaire.org, la France Insoumise et La République en Marche. "Mieux Voter" [MieuxVoter, 2022a] est une association française qui promeut l'utilisation du MJ comme mode de scrutin dès lors qu'un choix collectif doit être sélectionné : administration publique, associations, entreprises. Sur leur site internet, il est possible de trouver toutes les listes de citoyens –partis non affiliés à un parti politique national—qui ont utilisé le MJ pour classer leurs candidats lors des élections locales de 2020. Dans deux cas, Bordeaux et Annecy, le candidat sélectionné à l'aide de MJ a ensuite été élu maire.

LaPrimaire.org [LaPrimaire.org, 2016] est une initiative politique française dont l'objectif est de sélectionner un candidat indépendant pour l'élection présidentielle française en utilisant MJ comme règle de vote. Tous les citoyens français de plus de 18 ans ayant le droit de vote peuvent participer en tant que candidats ou électeurs. L'association Democratech a mis en place la plateforme pour la première fois en 2016 en vue de l'élection présidentielle de 2017. Le nombre d'électeurs ayant participé à l'élection était de 10676 au premier tour et de 32685 au second tour.

Tout citoyen éligible peut soumettre sa candidature à la plateforme et ceux qui obtiennent au moins 500 soutiens sont les candidats de l'élection. Le vote se déroule en deux tours. Au premier tour, chaque électeur est invité à exprimer son jugement, en utilisant MJ, sur cinq candidats aléatoires. À la fin de cette phase, les cinq candidats ayant les médianes les plus élevées sont considérés comme les finalistes qui se qualifient pour le second tour. Au second tour, chaque électeur est invité à exprimer son jugement, en utilisant MJ, sur les cinq finalistes. Le candidat

ayant la meilleure médiane à la fin de cette phase est sélectionné comme représentant pour l'élection présidentielle. Cependant, la participation de ce candidat à l'élection présidentielle réelle en France n'est pas acquise. En effet, selon la loi française, un candidat doit recueillir au moins 500 signatures d'élus pour pouvoir participer à l'élection présidentielle. Le candidat sélectionné par les électeurs de LaPrimaire.org en 2016 n'a recueilli que 135 signatures et n'a pas participé à l'élection présidentielle de 2017.

La sélection aléatoire de candidats est-elle une bonne technique d'élicitation ? Nous explorons les conséquences de l'incomplétude du profil et prouvons que cette méthode peut échouer à élire le candidat gagnant du profil complet. De plus, nous réalisons des expériences sur des profils générés aléatoirement et des profils suivant la distribution des notes d'un scénario de vote réel. Nous constatons que la probabilité de ne pas sélectionner le gagnant du profil complet diminue fortement lorsque le nombre de votants augmente et nous étudions quelle partie du vecteur de notes nous devons connaître pour que cette probabilité soit faible.

Conclusions. Nous avons prouvé que, étant donné un profil complet, il est possible de trouver un profil incomplet (tel qu'un achèvement possible est égal au profil considéré) avec un gagnant différent du profil complet. Nous avons également montré que la probabilité que cela se produise dépend du profil considéré.

De plus, nous avons réalisé des expériences sur des profils dans lesquels chaque vecteur de note a une distribution de note générée aléatoirement. En fixant le nombre de votants, nous avons analysé quelle part du vecteur de note nous devons connaître pour que la probabilité de manquer le gagnant soit faible, voir Figure 8.4. Nous avons également considéré un profil créé à partir de la distribution des notes d'un scénario de vote réel illustré par Figure 8.5. Nous avons constaté que la probabilité de ne pas sélectionner le gagnant du profil complet diminue fortement lorsque le nombre d'électeurs augmente, ce qui rend cette méthode adaptée aux élections avec un grand nombre d'électeurs mais moins appropriée pour les petits profils, voir Figure 8.6.

Les travaux futurs sur ce sujet pourraient étudier la possibilité de manipuler ce processus d'élicitation. En particulier du côté de l'"analyste" chargé d'interroger les électeurs et de la manière dont il peut potentiellement manipuler le résultat en posant des questions particulières à des groupes spécifiques d'électeurs. Par exemple, si elle sait que les électeurs d'une circonscription électorale donnée ont tendance à voter majoritairement pour un parti donné, elle peut décider de ne jamais interroger ces électeurs sur le candidat de ce parti (ou vice-versa, selon le parti qu'elle favorise).

Une autre étape pourrait concerner l'explicabilité de ce processus aux électeurs. En effet, les électeurs peuvent ne pas comprendre comment le fait de juger une seule personne donne une bonne approximation du résultat qui serait obtenu en demandant le profil complet. De plus, ils peuvent ne pas apprécier de juger un candidat au hasard et non leur favori.

Le code permettant de reproduire les expériences de cet article et les fichiers contenant les données rapportées sont disponibles sur <https://github.com/xoxor/elicitacionMJ>.

Figure 8.4: Probabilité moyenne de ne pas trouver le gagnant dans le cadre d'une distribution uniforme des préférences, étant donné $n = 100, m = 50$ pour 1000 lots de 1000 exécutions et $k \in \llbracket 1, 25 \rrbracket$.




Figure 8.5: Résultat obtenu à partir d'une expérience de vote en demandant à $n = 1147$ électeurs de juger $m = 12$ alternatives.





Figure 8.6: Probabilité moyenne de ne pas trouver le gagnant en utilisant une distribution des préférences de cas réel, étant donné $m = 12$ pour $n \in \{10, 25, 50, 100, 250, 500\}$ et $k \in \llbracket 1, 5 \rrbracket$.



8.4 Compromising as an equal loss principle

Le mot compromis vient du latin *compromissus*, participe passé de *compromittere* : *com* ensemble et *promittere* promettre. L'idée de compromis se retrouve en effet dans des textes datant de l'époque romaine, où deux parties en conflit qui souhaitaient soumettre leur différend à l'arbitrage désignaient un tiers (un arbitre) et promettaient mutuellement de se conformer au jugement incontestable de ce dernier. Si l'une des parties avait rompu cette promesse, elle aurait dû payer une pénalité. Dans un fragment d'une lettre écrite par Proculus, qui est rapporté par Zimmermann [1996, p.529], il est clair que les deux parties au conflit ont accepté de donner à l'arbitre des pouvoirs illimités dans sa décision. Aucun recours n'était possible contre la sentence finale, qui était contraignante, aussi injuste et inégale soit-elle. Cette idée du compromis comme simple processus d'arbitrage semble très différente de la notion que nous en avons aujourd'hui. Considérons les définitions des deux verbes dans un dictionnaire moderne : arbitrer "régler un différend entre deux personnes ou groupes après avoir entendu les arguments et les opinions de chacun" ; compromettre "parvenir à un accord par concession mutuelle". La concession mutuelle est l'élément qui nous fait immédiatement penser à un compromis et c'est précisément le facteur manquant dans la description du compromis romain et de l'arbitrage moderne. En citant Braybrooke [1982] "C'est tout simplement une mauvaise plaisanterie que de dire qu'une personne est partie à un compromis alors qu'elle n'en a rien retiré."

Dans Chapter 7, nous présentons deux versions différentes de compromis, en particulier une qui privilégie l'égalité aux dépens de l'efficacité de Pareto. De plus, nous considérerons le concept de "equal-loss" (perte égale) [Chun, 1988, Chun and Peters, 1991] comme la base du compromis dans toutes les situations où l'égalitarisme, au sens de concéder également, est une prérogative importante. Il s'agit d'une nouveauté dans la littérature des règles de choix social, qui jusqu'à présent n'ont fait qu'imposer la volonté de compromis sans réellement garantir que toutes les parties concèdent quelque chose. Le philosophe Day [1989] tente d'expliquer ce phénomène en pensant qu'il découle de la négation de l'adjectif "intransigeant" : "*Une personne intransigente est une personne qui n'est pas disposée à faire des concessions, donc (on en déduit à tort) une personne compromettante est une personne souple et disposée à faire des concessions—indépendamment du fait qu'elle reçoive une concession en retour. Quoi qu'il en soit, il faut insister, car il est généralement admis que le compromis implique nécessairement des concessions mutuelles.*" Cela pourrait expliquer pourquoi toutes les règles de vote qui tentent de faire un compromis se contentent en fait de la volonté des agents de faire un compromis.

Merlin et al. [2019] discutent des plus célèbres de ces procédures en proposant de les rassembler dans la même classe de *règles de compromis*. Dans Section 2.1.2, nous avons déjà défini le Majoritarian Compromise, introduit par Sertel [1986] et analysé plus en détail par Sertel and Yilmaz [1999]. Basée sur une révision de la règle de Condorcet-Bucklin, cette procédure part du choix idéal de chacun pour trouver une alternative soutenue par la majorité des votants. Si une telle alternative n'existe pas, elle se rabat sur le deuxième, le troisième et plus généralement le k -e meilleur choix des votants, jusqu'à ce qu'au moins une des alternatives considérées figure parmi les k premiers meilleurs pour une majorité. Si au lieu de considérer un accord pour la majorité des votants, nous souhaitons sélectionner comme gagnants les alternatives soutenues par l'unanimité des votants, alors la procédure correspond à la règle Fallback Bargaining. Plus

généralement, si nous considérons le soutien d'un certain quota q d'électeurs, nous nous référons à la règle de "q-approval fallback bargaining" [Brams and Kilgour, 2001].

Toutes ces SCRs imposent aux électeurs une volonté de compromis, mais nous soutenons qu'elles ne garantissent pas efficacement un résultat où les agents ont effectivement fait des compromis. Cet exemple motive notre point de vue:

Exemple 8.1. *Considérons le profil de préférence suivant avec $n = 100$:*

49	51
c	a
b	b
a	c

Lorsque $q \in \llbracket 1, \frac{n}{2} + 1 \rrbracket$, tous les compromis BK choisissent a et c tandis que les compromis BK révisés choisissent a . Ces résultats n'apparaissent pas comme un compromis car près de la moitié des votants obtiennent leur meilleur choix tandis que l'autre moitié doit se contenter de leur pire choix. Observez que b reçoit un soutien unanime lorsque chaque électeur recule d'un pas par rapport à son point idéal.

When $q \in \llbracket 1, \frac{n}{2} + 1 \rrbracket$, all BK compromises pick a and c while the revised BK compromises select a . These outcomes do not appear as a compromise as almost half of the voters obtain their best choice while the remaining half have to be contented with their worst one. Observe that b receives unanimous support when each voter falls back one step from her ideal point.

Nous définissons qu'une règle est *Egalitarian Compromise* (EC) (resp., *Egalitarian Compromise Compatible* (ECC)), si tous (resp., quelques) les gagnants sont parmi les alternatives avec les pertes les plus également distribuées. Évidemment, EC est une sous-classe de ECC.

De plus, nous définissons qu'une règle est *Paretian Compromise* (PC) (resp., *Paretian Compromise Compatible* (PCC)), si tous (resp., quelques) les gagnants sont parmi les alternatives Pareto optimales avec les pertes les plus également distribuées. Évidemment, PC est une sous-classe de PCC.

Conclusions. Nous prouvons que : aucune procédure de Condorcet n'est ECC ou PCC (Theorem 7.3), aucune règle de score n'est ECC (Theorem 7.4) ou PCC à l'exception de la règle d'antipluralité (Theorems 7.5 and 7.6), les compromis BK ne sont ni ECC ni PCC (Theorems 7.8 and 7.9) à l'exception de fallback bargaining qui est PC (Theorem 7.7).

Nos résultats d'impossibilité sont énoncés pour l'ensemble des mesures d'écart Σ , mais ils prévalent pour tout sous-ensemble de Σ . Comme Σ est le plus grand ensemble de mesures d'écart sensibles, ils sont valables pour tout concept spécifique d'équité que l'on pourrait choisir. D'autre part, nos résultats de possibilité sur le fallback bargaining et l'antipluralité ne se propagent pas aux sous-ensembles de Σ . En fait, dès qu'une restriction raisonnablement légère sur Σ est imposée, les deux règles ne sont plus PCC (Theorem 7.10). Avec deux individus et une restriction similaire sur Σ , tous les SCR à deux personnes bien connus de la littérature, à savoir le fallback bargaining, les règles de fallback bargaining, Pareto and veto rules, short listing et veto rank, ne parviennent pas à trouver des compromis ex-post (Theorem 7.11).

L'exclusion du principe d'égalité des pertes par la quasi-totalité de la littérature conduit à se demander si le principe ne présente pas d'intérêt dans un contexte de choix social discret. Cela semble être le cas pour les situations de vote où le nombre d'électeurs dépasse le nombre de candidats et où, en général, chaque candidat est classé dernier par au moins un électeur. Dans ces cas, la principale préoccupation concerne le soutien des alternatives plutôt que l'égalité. D'autre part, les problèmes de choix collectifs à deux personnes sont généralement interprétés comme des situations d'arbitrage ou de négociation où le consentement mutuel est un élément essentiel pour parvenir à une solution. Ainsi, le principe de perte égale semble être valide pour les problèmes de choix collectifs à deux personnes et notre analyse soulève la question de la conception de nouvelles règles d'arbitrage discrètes compatibles avec le principe de perte égale.

8.5 Conclusions et Perspectives

Tout au long de ce manuscrit, nous avons abordé plusieurs problèmes liés au domaine du choix social. En particulier, nous nous sommes concentrés sur deux scénarios. En considérant un modèle classique dans lequel les préférences d'un ensemble de votant sont connues, nous avons défini deux classes de règles de vote capables de refléter une notion de compromis dans laquelle l'égalitarisme, au sens de concéder également, est une préoccupation majeure. En outre, nous avons pris de la distance par rapport à cette perspective standard dans laquelle les préférences sont supposées être connues dès le départ et nous avons étudié le problème de l'élicitation des préférences dans différents contextes.

Il existe de nombreuses directions intéressantes pour les travaux futurs pour chacun des problèmes considérés. En ce qui concerne l'élicitation simultanée des règles de vote et des préférences des agents, on pourrait développer davantage de stratégies avec différentes heuristiques pour les comparer à celles que nous avons proposées. Par exemple, des stratégies moins coûteuses en calcul permettraient de tester des ensembles de données plus importants, avec plus de votants et plus d'alternatives. Il serait également intéressant d'étendre l'élicitation des règles de vote autres que les règles de position. En outre, nous pensons que la transformation des questions en profils d'exemple est un concept très intéressant qui n'est pas beaucoup exploré dans la littérature et qui pourrait être appliqué dans différents contextes.

Le chapitre consacré à l'analyse de la procédure d'élicitation à l'aide de MJ n'est pas encore tout à fait terminé, de sorte qu'il existe de nombreuses idées d'extensions possibles. On pourrait, par exemple, étudier la possibilité de manipuler le processus d'élicitation lui-même. En particulier, comment la partie chargée d'interroger les électeurs peut manipuler le résultat en dirigeant les questions vers des électeurs spécifiques. Une autre piste peut concerner l'explicabilité de ce processus. Les électeurs peuvent avoir du mal à croire que le fait de juger une seule personne donne une bonne approximation du résultat qui serait obtenu en demandant le profil complet.

Enfin, lorsque l'on considère le concept de compromis, différentes notions de compromis peuvent être conçues. Une approche pourrait consister à choisir des alternatives qui ne sont ni les meilleures ni les pires d'un électeur. Tout dépend de l'idée de justice que l'on veut représenter et des priorités du scénario considéré. En outre, le compromis entre équité et efficacité peut être exploré dans d'autres contextes, plus riches. Pour conclure, nous aimons penser que chacune de nos contributions n'est que le début d'une longue série de développements futurs.

RÉSUMÉ

Les règles de vote représentent des moyens formels d'agréger les préférences d'un groupe d'électeurs en une décision collective. Mais quelle règle choisir dépend du contexte et de la conception de la justice de la société concernée. Bien que toutes les règles de vote puissent être considérées comme un compromis entre individus, certaines le sont plus que d'autres. Dans ce manuscrit, nous proposons deux définitions de compromis basées sur une notion de «equal-loss», favorisant un résultat où chaque électeur concède le plus également possible. Nous étudions plusieurs règles qui reflètent un certain concept de compromis et nous montrons comment, bien qu'elles imposent une volonté de compromis aux individus, elles ne garantissent pas un résultat où chacun a effectivement concédé. De plus, nous remettons en question l'hypothèse, souvent faite dans le cadre classique, que la règle de vote est fixée à l'avance et qu'une information complète sur les préférences des individus est fournie. En nous concentrant sur les règles de vote positionnelle, nous proposons des procédures d'élicitation capables de poser des questions aux votants et à l'entité responsable de la conception de la règle de vote, afin d'obtenir rapidement les informations les plus pertinentes. En considérant la règle de Jugement Majoritaire, nous analysons une procédure d'élicitation des préférences utilisée dans des scénarios de vote réels. Sur la base d'une information partielle sur les préférences des individus, nous analysons la probabilité de sélectionner un gagnant différent de celui obtenu en considérant le profil complet et nous étudions le nombre minimum de questions à poser aux électeurs pour que cette probabilité soit faible.

MOTS CLÉS

Intelligence Artificielle, Computational Social Choice, Théorie de la décision, Élicitation de préférence, Modélisation des préférences.

ABSTRACT

Voting rules are formal means of aggregating preferences of a group of voters over a set of alternatives into a collective decision. But which rule to use is a difficult question and the answer depends on the conception of justice of the concerned society or committee. Although all voting rules can be seen as a compromise between individuals, some are more so than others. In this manuscript we provide two notions of compromise based on an equal-loss perspective, favoring an outcome where every voter concedes as equally as possible. We study several rules that reflect some idea of compromise and show how, although imposing a willingness to compromise on individuals, they fail to ensure an outcome where everyone has effectively compromised. Moreover, we question the assumption, often made in the classical setting, that the voting rule is fixed beforehand and full information on the preferences of the individuals is provided. Focusing on positional scoring rules, we propose elicitation procedures based on minimax regret, able to interleave questions to the voters and to the entity responsible for designing the voting rule in order to quickly obtain the most relevant information. Considering Majority Judgment rule, we analyze a particular preference elicitation procedure used in real voting scenarios. Given a partial information on individuals preferences, we analyze the probability of selecting a different winner from the one obtained by considering the full profile and study the minimum number of questions to ask voters in order for this probability to be low.

KEYWORDS

Artificial Intelligence, Computational Social Choice, Decision Theory, Preference Elicitation, Preference Modeling.