

Multi-Winner Voting: Representation, Manipulation, Deliberation

DISSERTATION

submitted in partial fulfillment of the requirements for the degree of

Doktor der Technischen Wissenschaften

by

Oliviero Nardi, M.Sc

Registration Number 12229255

to the Faculty of Informatics
at the TU Wien

Advisor: Prof. Dr. Stefan Woltran
Second advisor: Assoc. Prof. Dr. Robert Ganian
Third advisor: Privatdoz. Dr. Martin Lackner
Mitwirkung: Dr. Jan Maly

The dissertation has been reviewed by:

Markus Brill

Davide Grossi

Vienna, September 14, 2026

Oliviero Nardi

Declaration of Authorship

Oliviero Nardi, M.Sc

I hereby declare that I have written this thesis independently, that I have completely specified the utilized sources and resources and that I have definitely marked all parts of the work—including tables, maps and figures—which belong to other works or to the internet, literally or extracted, by referencing the source as borrowed.

I further declare that I have used generative AI tools only as an aid, and that my own intellectual and creative efforts predominate in this work. In the appendix “Overview of Generative AI Tools Used” I have listed all generative AI tools that were used in the creation of this work, and indicated where in the work they were used. If whole passages of text were used without substantial changes, I have indicated the input (prompts) I formulated and the IT application used with its product name and version number/date.

Vienna, September 14, 2026

Oliviero Nardi

Acknowledgements

I just *play* all the time,
and am fortunate enough
to get paid for it.

— Martin Gardner [1998]

First and foremost, I would like to acknowledge my advisors: Stefan Woltran, Martin Lackner, and Jan Maly. I am grateful to have had the three of you guiding me along the way. Stefan, thank you for entrusting me with this position, for all the opportunities you provided, for allowing me to carve my own path, and for always being there when needed. Martin, thank you for bringing me back to social choice theory and for helping me give direction to my PhD back when we wrote our first paper together. I always appreciated your supportive and cheerful energy, and I always left our meetings recharged and inspired. Jan, first, thank you for steering me towards doing a PhD with Stefan in Vienna when I was a research assistant in Amsterdam. Most of all, thank you for your close guidance and support after you moved back to TU Wien; you are a dedicated and insightful advisor, and I enjoyed working with you. I am also thankful for the opportunity to work in your group these last few months at WU Wien.

I would like to thank my reviewers, Markus Brill and Davide Grossi, for carefully reading this thesis and for their valuable feedback, which helped me improve the manuscript significantly. I really appreciate the time you devoted to it and am grateful to have had two leading figures in the field as reviewers. I am grateful to Ezio Bartocci and Laura Kovács for agreeing to be part of my defence committee. I am also grateful to Dominik Peters and Jérôme Lang for hosting me at LAMSADE in Paris in the autumn of 2024. Next, I would like to thank Ulle Endriss for introducing me to social choice theory during my master's through his inspiring lectures. I also feel very fortunate to have worked with so many wonderful co-authors: Anton Baychkov, Michael Bernreiter, Arthur Boixel, Robert Bredereck, Ronald de Haan, Théo Delemazure, Ulle Endriss, Bailey Flanigan, Umberto Grandi, Ayumi Igarashi, Martin Lackner, Jan Maly, Nick Mattei, Nicolas Maudet, Arianna Novaro, Jannik Peters, Simon Rey, Felicia Schmidt, Stanisław Szufa, Zoi Terzopoulou, and Stefan Woltran.

I would also like to thank Juliane Auerböck and Beatrix Buhl for all the support, and Clarissa Schmid for her help during the (stressful) final stretch of the PhD.

When I moved to Vienna, my decision was mostly influenced by career considerations: Where should I do my PhD? I did not suspect that Vienna would bring me some of the people I now cherish the most. I would now like to turn to them. (I will put you in several neatly separated boxes, but that is just for expository convenience.)

First, my Italian friend group: Alessia, Angelo, Davide S., Elena, Eleonora, Ennio, Federica, Ferdinando, Francesco, Giovanni, Giulia, Ludovica, Luigi, Luisa, Marta, and Michele. I am grateful to have found you. Having you here meant the world to me, and it is so special how close we got over the years. You went from “my friends from Vienna” to just friends *tout court*. How sad I am to leave you only shows how lovely it was being here with you, and I know our friendship will last longer than a PhD. I am looking forward to seeing you regularly.

I think my office mates deserve an extra mention: Federica, Davide S., and Giovanni. I cannot imagine this PhD without you, and I am grateful to have had you there from (essentially) the start to (essentially) the end, through the good and the inevitable bad that a PhD entails, through countless lunches, breaks, and iced cappuccini. I could not have asked for better office mates. Federica deserves a further word, for having been there from the *very* first day of my PhD until the end, both as an office mate and as a friend. You are the first person I met in Vienna, and you know I looked up to you as a PhD role model.

I have been lucky enough to be in the LogiCS@TUWien doctoral school. It was a warm and friendly cohort, and I loved being part of it. I want to mention here some people from this group in particular. Hai, you are one of the most hilarious people I know, and the trip to Vancouver with you and Federica is one of the best memories I have of the PhD. Tianwei, I always loved our conversations, from the deepest to the silliest topics. Tobias, you are generous, kind, and a great friend, and I always enjoyed having drinks with you. LogiCS@TUWien would not have been the same without the warm presence of Anna Prianichnikova, our coordinator, who always encouraged us to be friends first. It is sad to finish the PhD without her, and I am sure I speak for all of us when I say that she is dearly missed.

Outside of LogiCS@TUWien, I had the pleasure of sharing my PhD years with many others. Let me start with my “original co-advisees”, Anna, Matthias, and Michael, and, alongside them, Wolfgang. It was great to share part of this path with you, and I always looked forward to our cake meetings. I think this is also a good place to mention the many great people on my floor, in particular Alessandro, Anouk, Davide D. S., Davide L., and Ida; thank you for making my work environment such a nice place to be. Next, my fellow social choice theory students in Vienna: Felicia, Hanna, Paul, and Stefan F.; I am happy you came into the picture, and I am happy to have had such friendly people working on similar topics in the second leg of my PhD, and going on a ride in the Prater was definitely the best way to take a break during the workday.

While my time in Vienna was great in itself, I also need to acknowledge the role that people back home played in this. Being an expat is an exhilarating experience, but it also brings many difficulties. Knowing that I could always return home to a lovely

bunch helped a lot. First, I want to thank my lifelong friends, Alessio, Giacomo Belfiore, and Pietro. We basically share the same mind by now, and I cannot imagine my life without you in it. I always miss you, but I find it reassuring that, when we meet, it is as if no time has gone by. I would also like to acknowledge all the other friends from Chiampo I gathered over the years, such as Filippo, Gianluca, and Giulio, for always making my trips home an amazing time, and my hometown a place worth missing. As for Giacomo Bertizzolo, my fellow expat, I had no doubts that distance would not harm our friendship, and I am happy to think that wherever life takes us, we will always share a connection. Alessandro C., Alessandro S., Edoardo, Giovanni, Giulio, Jacopo, Luigi, Leonardo, Max, Nicola, my nonvicentini friends: at this point, speaking to you has become such an integral part of my everyday life that I rarely stop to think about how great it is to have you. Thank you for listening to me when I needed to vent, for your advice on my future, and for simply being there. Some of my happiest memories during the past four years have been due to you.

These acknowledgements would not be complete without mentioning my big and wonderful family—Gina, Claudio, Benedetta, Carolina, Vittoria, Ettore, Luca, Giordano, Giulia, Mariasole, Armando, Augusto, Arturo, Ruggero, Lilli, and Gianni. This family is such a fundamental part of who I am, and it shapes me through and through. Thank you for supporting me materially during my earlier studies, and thank you for making me the person that I am. Missing you is hard, but our bond is indissoluble, distance notwithstanding, and coming home to you will never cease to be special.

I saved the most important for last. Sometimes, when I considered whether to continue my academic career, I wondered: If I don't, what was the point of all this? Moving to Vienna, leaving it all behind, starting over. If I am not even going to continue, what was the point? I now realise this is a pointless question: Josephine, the point was you. Above all, I am grateful that life brought us together, and I am grateful for having you to count on. You can always count on me, too.

Kurzfassung

Während die meisten Menschen die zentrale Rolle anerkennen, die Wahlen in Demokratien spielen, machen sich nur wenige über die Feinheiten der Gestaltung von Wahlregeln und deren Auswirkungen Gedanken. Die Sozialwahltheorie, die mathematische Untersuchung kollektiver Entscheidungen, hat sich zum Ziel gesetzt, zu klären, welche Wahlregeln wünschenswert sind. Ein aktueller Trend in der Forschung konzentriert sich auf Wahlen, bei denen das Ergebnis nicht aus einer einzelnen Person besteht (wie bei Präsidentschaftswahlen), sondern vielmehr aus einer Gruppe von Personen (zum Beispiel Parlamente und Ausschüsse). Die vorliegende Arbeit trägt in mehreren Richtungen zu diesem Gebiet der Mehrfachwahl bei: von vertrauten Anwendungen (Wahl von Parlamenten und Ausschüssen) über weniger traditionelle (im Laufe der Zeit wiederholte Wahlen) bis hin zu nicht-wahlbezogenen Anwendungen (Auswahl einer Reihe von Standpunkten aus einer Diskussion, um diese zusammenzufassen). Wir konzentrieren uns auf die Zustimmungswahl, bei der es den Wählenden freisteht, so viele Alternativen zu unterstützen, wie sie möchten, anstatt eine einzige Präferenz angeben zu müssen (sei es für eine Partei oder für Kandidaten/eine Kandidatin).

Der erste Beitrag befasst sich mit Ausschusswahlen. Diese haben zum Ziel, aus einem größeren Kandidatenpool eine festgelegte Anzahl von Gewinnerinnen und Gewinnern (den Ausschuss) auszuwählen. Dies ist einer der in den letzten Jahren am intensivsten untersuchten Bereiche der Sozialwahltheorie. Ein bedeutender Teil der Forschungsarbeiten befasst sich mit der sogenannten Proportionalität, das heißt der Gewährleistung, dass der Ausschuss die Präferenzen der Wählerschaft proportional widerspiegelt. Der Großteil der Bemühungen konzentrierte sich bislang jedoch auf eine bestimmte Ausprägung davon: die Verhinderung von Unterrepräsentation. Konkret bedeutet dies: Wenn ein bestimmter Anteil der Bevölkerung ähnliche Präferenzen hat, sollte mindestens derselbe Anteil des Ausschusses diese Präferenzen widerspiegeln. Im Gegensatz dazu hat die komplementäre Idee, d.h. eine Überrepräsentation zu verhindern, praktisch keine Beachtung gefunden, obwohl sie im Prinzip ebenso wünschenswert ist und konkrete Anwendungsmöglichkeiten bietet. Intuitiv tritt eine Überrepräsentation dann auf, wenn eine Gruppe von Wählenden einen unverhältnismäßig großen Teil des Ausschusses bestimmt. Wir formulieren hierfür ein starkes und ansprechendes Axiom – also eine mathematische Definition dessen, was „Überrepräsentation vermeiden“ bedeutet –, das wir als *gerechtfertigte Höchstquote* (justified upper quota) bezeichnen. Wir führen eine Verallgemeinerung der Thiele-Regeln (einer gut untersuchten Familie von Abstimm-

mungsregeln für Ausschüsse) ein, die wir als *zusammengesetzte Thiele-Regeln* (composite Thiele rules) bezeichnen, und charakterisieren die einzige Regel in dieser Klasse, die unser Axiom erfüllt. Diese Regel, die eine natürliche Erweiterung einer von John Quincy Adams vorgeschlagenen Sitzverteilungsregel darstellt, wurde bisher nicht untersucht. Darüber hinaus führen wir den Begriff der *gerechtfertigten Nähe zur Quote* (justified near quota) ein, ein Axiom, das Unterrepräsentation und Überrepräsentation ausgleicht, und gelangen zu analogen Ergebnissen: Wir zeigen, dass die einzige Thiele-Regel, die dieses Axiom erfüllt, die natürliche Erweiterung einer anderen Sitzverteilungsmethode ist, nämlich der Sainte-Laguë-Methode. Schließlich analysieren wir die Vereinbarkeit unserer Axiome mit etablierten Proportionalitätskonzepten wie EJR+, einem prominenten Axiom zur Vermeidung von Unterrepräsentation.

Der nächste Beitrag wendet sich einem anderen Szenario zu: Einzelwahlen, die im Laufe der Zeit wiederholt stattfinden. Hier bilden die ausgewählten Personen kein Gremium, sondern stellen eine zeitliche Abfolge von Entscheidungen dar. Diese Perspektive ermöglicht es, Wahlregeln so zu entwickeln, dass alle Wählenden zufriedengestellt sind, wenn schon nicht mit einer einzigen Entscheidung, so doch zumindest im Laufe der Zeit. Solche Fairnessüberlegungen könnten jedoch einen Anreiz für Wählende schaffen, ihre Präferenzen falsch darzustellen, da sie eine besonders einfache Form des strategischen Wahlverhaltens ermöglichen: *Trittbrettfahren* (free-riding). Trittbrettfahren tritt auf, wenn Wählende ihre Präferenzen bei einer Wahl falsch angeben, indem sie einen Kandidaten ablehnen, den sie eigentlich befürworten. Dabei erwecken sie den Eindruck, dass sie unterrepräsentiert sind, und erhöhen dadurch ihre Chance, bei nachfolgenden Entscheidungen stärker berücksichtigt zu werden. Das Wechselspiel zwischen Fairness und strategischer Manipulation ist ein wiederkehrendes Thema in der Sozialwahltheorie, und das Trittbrettfahren erscheint aufgrund seiner Einfachheit und Unmittelbarkeit besonders problematisch. Wir zeigen einerseits, dass selbst schwache Repräsentationsgarantien die Möglichkeit des Trittbrettfahrens implizieren. Andererseits untersuchen wir das Trittbrettfahren mittels Komplexitätsanalyse und Simulationen und kommen zu dem Schluss, dass diese Strategie mit einem nicht zu vernachlässigenden Risiko verbunden ist. Mit anderen Worten: Es kann dazu führen, dass Wählende schlechter gestellt sind, als sie es ohne Trittbrettfahren gewesen wären. Somit ist diese Form der Manipulation weniger attraktiv, als man es intuitiv annehmen könnte. Wir ergänzen unsere Analyse durch die Betrachtung von Trittbrettfahren in einem verwandten Szenario, in dem die Wahlen parallel stattfinden, und kommen zu ähnlichen Schlussfolgerungen.

Unser abschließender Beitrag befasst sich nicht mit einem Wahlkontext an sich, auch wenn er potenzielle Auswirkungen auf die demokratische Praxis hat. Wir veranschaulichen die Vielseitigkeit des Rahmens für Mehrfachwahlverfahren, indem wir ihn auf das Problem der Diskursanalyse anwenden. Im Sinne der deliberativen Demokratie besteht ein wachsendes Interesse an Plattformen, die einen groß angelegten Meinungsaustausch unter den Bürgerinnen und Bürgern ermöglichen sollen. Die Interpretation dieser potenziell umfangreichen Diskussionen bleibt jedoch eine offene Forschungsfrage. Wir schlagen eine neuartige und interpretierbare Methode vor, um diese Debatten so zusammenzufassen, dass die geäußerten Positionen proportional zu ihrer Unterstützung

dargestellt werden. Wir betrachten verschiedene Standpunkte in einer Diskussion als die „Kandidaten“ unserer Wahl, die Unterstützung, die sie erhalten, als Stimmen, die zu ihren Gunsten abgegeben wurden, und die Zusammenfassung als eine Menge gewinnender Standpunkte. Um die Kandidatenstandpunkte zu identifizieren, nutzen wir die abstrakte Argumentationstheorie (abstract argumentation theory), einen gängigen Formalismus zur Darstellung von Debatten und ihrer Semantik. Wir stellen hierfür mehrere Algorithmen vor und vergleichen sie sowohl theoretisch als auch mittels Simulationen.

Abstract

While most people would recognise the pivotal role elections play in democracies, few pause to consider the subtleties of voting rule design and their impact. Social choice theory, the mathematical study of collective decisions, sets out to clarify which voting rules are desirable. A recent trend in research focuses on elections in which the outcome is not a single individual (as in presidential elections) but rather a collection of winners (for example, parliaments and committees). This thesis contributes to the field of multi-winner voting in several directions and under a range of scenarios: from familiar settings (election of parliaments and committees) to less traditional ones (elections repeated over time), and even to non-electoral applications (selecting a slate of viewpoints from a discussion in order to summarise it). We focus on approval voting, an emerging format in which voters are free to “approve” as many alternatives as they like, rather than indicating a single preference (be it for a party or a candidate).

The first contribution concerns committee elections, in which the goal is to select a fixed number of winners (the committee) from a larger pool of candidates. This is one of the most studied settings in social choice theory in recent years. A significant body of work addresses proportionality, that is, ensuring that the committee reflects the preferences of the electorate proportionally. However, most of the effort has concentrated on a specific formulation of this: preventing underrepresentation. Specifically, if a certain fraction of the population has similar preferences, then at least the same fraction of the committee should reflect these preferences. By contrast, the complementary idea of preventing *overrepresentation* has received virtually no attention, despite being, in principle, equally desirable and having concrete applications. Intuitively, overrepresentation occurs when a group of voters determines a disproportionately large part of the committee. We formulate a strong and appealing axiom for this—that is, a mathematical definition of what “avoiding overrepresentation” means—which we call *justified upper quota*. We introduce a generalisation of Thiele rules, a well-studied family of committee voting rules, called *composite Thiele rules*, and characterise the unique rule in this class satisfying our axiom. This rule, which is a natural extension of an apportionment rule proposed by John Quincy Adams, has not been studied before. Furthermore, we introduce *justified near quota*, an axiom that balances underrepresentation and overrepresentation, and obtain analogous results: we show that the only Thiele rule that satisfies it is the natural extension of another apportionment method, Sainte-Laguë’s method. Finally, we analyse the compatibility of our axioms with established notions of proportionality, such as EJR+, a popular axiom for avoiding underrepresentation.


The next contribution turns to another setting: single-winner elections held repeatedly over time. Here, the “multiple winners” are not a committee, but rather the temporal sequence of winners. This perspective allows voting rule designers to satisfy all voters, if not with a single decision, at least over time. Such fairness considerations, however, might incentivise voters to misrepresent their preferences, enabling a particularly straightforward form of strategic voting: *free-riding*. Free-riding occurs when a voter misreports her preferences in an election by opposing a candidate she actually approves of, with the goal of appearing underrepresented, thereby increasing the chance of receiving greater consideration in subsequent decisions. The interplay between fairness and strategic manipulation is a recurrent theme in social choice theory, and free-riding appears particularly problematic due to its simplicity and immediacy. On the negative side, we show that even weak representation guarantees imply the possibility of free-riding. On the positive side, we study free-riding through complexity analysis and simulations and conclude that it carries a non-negligible risk. In other words, it might leave a voter worse off than she would have been had she voted earnestly. Thus, free-riding is less appealing than one could intuitively assume. We complement our analysis by examining free-riding in a related setting in which elections are held simultaneously rather than over time, and reach similar conclusions.


Our final contribution is not about an electoral setting *per se*—though it still has potential consequences for democratic practice. We exemplify the versatility of the multi-winner voting framework by applying it to the problem of summarising discussions. Indeed, in the spirit of deliberative democracy, there has been a growing interest in platforms that enable large-scale deliberation among citizens. Facilitating the interpretation of these potentially extensive discussions remains an open challenge. We propose a novel and interpretable way to summarise them that represents the positions expressed in proportion to their support. We cast different viewpoints in a discussion as the “candidates” of our election, the support they receive as votes expressed in their favour, and the summary as a set of winning viewpoints. To identify the candidate viewpoints, we leverage abstract argumentation theory, a well-understood framework for representing debates and their semantics. We introduce several algorithms for this and compare them both theoretically and via simulations.

Contents

Kurzfassung	ix
Abstract	xiii
Contents	xv
1 Prologue	1
1.1 Scope, Methodological, and Normative Commitments	1
1.2 Why This Thesis?	2
1.3 Chapter Summaries and Related Publications	4
1.4 Further Details	6
2 Background	9
2.1 Historical Background	9
2.2 Technical Background	20
3 Representation	33
3.1 Introduction	33
3.2 Preliminaries	37
3.3 Composite Thiele Rules	43
3.4 Avoiding Overrepresentation	46
3.5 Balancing Underrepresentation and Overrepresentation	62
3.6 Combining Lower, Near, and Upper Quota Axioms	67
3.7 Conclusion	75
4 Manipulation	77
4.1 Introduction	77
4.2 Related Work	79
4.3 Preliminaries	82
4.4 Possibility and Risk of Free-Riding	90
4.5 Computational Complexity	100
4.6 Numerical Simulations	117
4.7 Beyond the Sequential Case	127
4.8 Conclusion	146
	xv

5	Deliberation	149
5.1	Introduction	149
5.2	Related Work	153
5.3	Preliminaries	154
5.4	Perfect Representation	162
5.5	Optimising Representation	169
5.6	Justified Representation	175
5.7	Numerical Simulations	180
5.8	Conclusion	188
6	Epilogue	191
A	Additional Results on Representation	193
A.1	Justified Upper Quota Swaps and Might Break EJR+	193
A.2	Counterexamples for Other Voting Rules	197
B	Additional Results on Deliberation	203
B.1	Approximation Ratio	203
B.2	Varying the Number of Selected Extensions	204
B.3	Representation Metrics	205
	Overview of Generative AI Tools Used	211
	Acronyms	213
	Bibliography	215





Prologue

Forming a parliament; choosing a committee; appointing, every term, a trustee. These are all elections where the winner is not a single individual, but a complex body—be it a multi-member group or a succession of officeholders. These elections will be our focus.

1.1 Scope, Methodological, and Normative Commitments

This is a thesis on multi-winner voting, broadly understood. Under this umbrella, we consider elections in which the outcome is not a single winner but rather a *collection* of winners. The example that most readers will recognise is the election of a parliament. Another common example is the selection of a committee, for instance, an advisory panel in a university. A less familiar example involves elections that take place repeatedly over time; in such cases, the multi-winner outcome is given by the temporal sequence of selected candidates. For example, if a board periodically appoints a representative, then the sequence of appointed representatives constitutes the outcome. While these three examples cover a diverse range of settings, they are united by the fact that different members of the selected slate can represent different groups within the electorate—either simultaneously or across successive periods. This is our area of interest; in the following paragraphs, we will clarify it further and further.

Although we consider a diverse range of elections—at least when it comes to the type of outcome that is sought—we fix the ballot format, that is, the input to our voting process. Specifically, we consider only *approval elections* [Brams and Fishburn, 1983], in which voters may approve any number of the available candidates. This differs from the commonplace ballots used in most elections, where voters are required to express a single preference (be it for a specific candidate or a predetermined party). We focus on approval voting because of its many appealing qualities; in fact, it is one of the most studied formats for multi-winner elections [Lackner and Skowron, 2023]. Notably, it balances expressiveness and simplicity, in contrast to single-preference ballots (which

elicit limited information) and ballots in which voters rank, or assign finer grades to, all options (which require complex judgments). This balance benefits both the voter and the theoretician: the former can express herself while facing a lighter cognitive load; the latter can work with informative preferences which are relatively straightforward to interpret. The volume edited by [Laslier and Sanver \[2010a\]](#) and the book by [Lackner and Skowron \[2023, Chapter 1\]](#) contain several arguments and references on the merits of this format.

Having now restricted the object of our interest, we also restrict our analytical angle. Elections are too pervasive to be studied in full generality: they arise at every stage of collectivity, from small social units to entire nations, whenever differences cannot be deliberated away. We therefore do not attempt to address all facets of voting. Instead, we focus on one technical, yet *essential*, aspect of it: the voting *rule*. In other words, we study how votes are aggregated and translated into a set of winners, setting aside other institutional questions, such as how candidates enter the election. We conceptualise voting rules as formal objects, *algorithms*, and study them in the language of mathematics, seeking precise definitions of intuitive concepts such as “fairness”. This places us in the field of social choice theory [[Arrow et al., 2002](#)], the formal study of voting rules, in particular, multi-winner and combinatorial voting [[Lang and Xia, 2016](#), [Faliszewski et al., 2017](#), [Lackner and Skowron, 2023](#)].

Finally, if approval-based multi-winner elections are our object and formal rule analysis our method, proportional representation will be our normative commitment. Proportionality prescribes that the composition of the collection of winners reflects the composition of the electorate—for example, seats should be distributed proportionally to the support of each party. The notion of proportionality takes different meanings across the settings we study, but the guiding idea remains the same: *outcomes should represent the electorate*. We commit to this principle throughout this thesis. The question of whether proportionality is an appropriate normative standard is not addressed here: such discussions belong primarily to political science and empirical analysis [[Mill, 1865](#), [Lijphart, 1985](#), [Norris, 1997](#), [Lijphart, 1999](#), [Powell, 2000](#), [Gallagher and Mitchell, 2005](#), [Herron et al., 2018](#)]. Our concern here is with its formal definition, its implementation, and its implications within social choice theory.

1.2 Why This Thesis?

The formal study of approval-based multi-winner voting, under the normative lens of proportionality, is a developed and active field [[Lackner and Skowron, 2023](#)]. Readers unfamiliar with this area, looking for a clearer sense of why this matters or of the concrete questions addressed, may refer to Chapter 2, where we sketch the history of proportional multi-winner voting and introduce its main concepts. Here, we do not introduce the field itself, but rather explain how the contributions of this thesis fit *within* the field.

Social choice theory is a model-based discipline, where by model we mean a stylised description of some target system [[Frigg and Hartmann, 2025](#)]. Researchers take an election format (such as parliamentary elections, committee elections, or repeated elections),

formalise it mathematically, and then study the resulting model. Under a simplified view, one can contribute to this discipline along three dimensions. The first is to advance our understanding of an established model by analysing its properties, resolving its open questions, and possibly examining it from fresh perspectives. The second is to introduce new models, in particular by focusing on electoral processes that have not yet been formalised in the literature, thereby initiating their systematic study. The third hinges on the abstract nature of the developed models and applies them to concrete problems, possibly beyond their original electoral interpretation. These approaches are mutually reinforcing: the first deepens our understanding of well-studied models, the second proposes new settings to study, while the third clarifies why a model is meaningful by demonstrating its applicability or generality.

We aim to contribute along all three dimensions. Chapter 3 (“Representation”) follows the first approach. Here, we focus on the main model studied in approval-based multi-winner elections, [Approval-Based Committee \(ABC\) voting](#) [[Lackner and Skowron, 2023](#)]. This models elections where the goal is to select a fixed number of winners from a larger pool of candidates. Within this formalism, we address a fundamental question: *what does it mean to achieve proportional representation?* This has received considerable attention, and many definitions of proportionality for [ABC](#) voting exist. We propose a novel perspective based on avoiding the overrepresentation of coalitions of voters, thereby limiting the power of such cohesive groups to disproportionately influence the outcome. In formulating proportionality in these terms, we address an open question explicitly posed by [Lackner and Skowron \[2023, Chapter 7, Q8\]](#) in their book *Multi-Winner Voting with Approval Preferences*, a central reference in this field.

Chapter 4 (“Manipulation”) follows the second approach. We study a type of election in which multiple decisions have to be made, either simultaneously (for instance, several referenda held at the same time) or sequentially (for instance, a board that regularly votes on new issues). The model we propose generalises models that have received attention in the literature, such as multiple referenda [[Amanatidis et al., 2015](#)] and perpetual voting [[Lackner, 2020](#)], which correspond to the two examples above. In this framework, we study one of the main recurring themes in social choice theory: how rules that embody some form of fairness consideration—in our case: representation guarantees—incentivise voters to misreport their preferences.

Lastly, Chapter 5 (“Deliberation”) is situated in the third dimension; we show how the theory of proportional representation in [ABC](#) voting can be applied even outside the purely electoral setting—though the framework we develop may still be useful in democratic practice. In recent years, platforms for democratic participation have attracted growing interest from both academics and practitioners [[OECD, 2020](#)], in the spirit of deliberative democracy [[Elster, 1998](#)]. On such platforms, citizens can engage directly in large-scale discussions of political issues by expressing their views on proposed themes or by supporting statements posted by other citizens. We study how multi-winner voting techniques can be applied to summarise these discussions so that different viewpoints are represented in proportion to the support they receive.

To sum up and answer the question in the title of this section,

we contribute to the established field of approval-based multi-winner elections

- (i) by proposing a novel perspective on proportional representation for its most studied model, [ABC](#) voting;*
- (ii) by studying how representation interacts with strategic manipulation when multiple decisions have to be made;*
- (iii) by showing how the theory of multi-winner voting can be applied in concrete applications, such as large-scale deliberation.*

Note that, throughout the thesis, we distinguish between the informal notion of approval-based multi-winner voting in general (encompassing different election formats) and the specific models we study, including [ABC](#) voting.

1.3 Chapter Summaries and Related Publications

We now switch to a more technical register and briefly summarise the contents of each chapter. We further indicate the peer-reviewed publications on which Chapters 3 to 5 are based. All listed publications are joint work, with authors listed in alphabetical order.

- Chapter 2 provides the necessary background and notation. It consists of two parts. The first (Section 2.1) sketches the historical development of approval-based multi-winner voting, situating this thesis within the broader literature on the topic. We discuss both social choice theory in general—our methodological angle—and specifically proportional representation—our normative angle. The second part (Section 2.2) introduces the mathematical framework used throughout the thesis. Furthermore, it introduces our centrepiece model, [ABC](#) voting, together with some concrete voting rules and formal definitions of proportionality. This model will not be our sole focus, but it plays a chief role in this work (and, in general, in the literature). Hence, it is worth explaining in detail. It will also serve as a warm-up for the reader to the formal study of voting.
- Chapter 3 is based on a paper published in the proceedings of the 27th ACM Conference on Economics and Computation, together with Anton Baychkov, Martin Lackner, Jan Maly, and Jannik Peters [[Baychkov et al., 2026](#)]. This paper was, in turn, based on earlier work with Martin Lackner, presented at the 16th Multidisciplinary Workshop on Advances in Preference Handling [[Lackner and Nardi, 2025](#)]. We study the standard [ABC](#) voting model and propose a new perspective on proportional representation that focuses on preventing the *overrepresentation* of cohesive groups of voters, rather than their underrepresentation, as virtually all previous work on [ABC](#) voting does. We additionally consider another perspective, in which we aim to balance overrepresentation and underrepresentation. We study this phenomenon within the class of composite Thiele rules, a family of voting rules we introduce, which generalises the class famously proposed by the Danish mathematician Thorvald N. Thiele [[Thiele, 1895](#)]. For both approaches, we give a

formal definition of proportionality under that view; we call the resulting *axioms* (that is, formal definitions of voting rule properties) [Justified Upper Quota \(JUQ\)](#) and [Justified Near Quota \(JNQ\)](#). We show that, among composite Thiele rules, there is exactly one voting rule satisfying each axiom. For JUQ, this is a rule that generalises a method for parliamentary elections proposed by John Quincy Adams, the sixth president of the United States [[Balinski and Young, 2001](#)]. For JNQ, it is a rule known as “Sainte-Laguë approval voting” [[Lackner and Skowron, 2023](#)], which likewise extends a method for parliamentary elections by [Sainte-Laguë \[1910\]](#). In fact, our whole approach is rooted in the theory of party-list parliamentary elections, which are—mathematically—a fragment of the [ABC](#) voting model, and where the issue of overrepresentation is well understood.

- Chapter 4 is based on a paper published in the proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems with Jan Maly and Martin Lackner [[Lackner et al., 2023a](#)]. It was later extended and published in the Journal of Artificial Intelligence Research by the same group of authors [[Lackner et al., 2026](#)]. We consider elections involving several decisions. We prioritise the case where decisions are taken sequentially, over time, but we also consider the simultaneous case. We study two families of voting rules: one that naturally corresponds to the rules introduced by Thiele (mentioned above) for [ABC](#) voting, and the other based on a class of mathematical aggregators, the [Ordered Weighted Averaging \(OWA\)](#) operators [[Yager, 1988](#)]. However, some of our results hold for *all* voting rules. We argue that proportional representation in this context gives rise to a simple way of manipulating the election, which we call free-riding. Free-riding describes a strategy by which a voter misreports her preferences in an early decision by opposing a candidate she actually approves in order to appear underserved, thereby increasing the chance of receiving favourable treatment in later decisions. We show that even minimal forms of representation imply this possibility; however, we also argue that this form of strategic voting involves some risk, contrary to initial intuition, limiting its appeal. We approach this both theoretically (i.e., with theorems and proofs) and via numerical simulations.
- Chapter 5 is based on a paper published in the proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems with Michael Bernreiter, Jan Maly, and Stefan Woltran [[Bernreiter et al., 2024a](#)]. Here, we interpret discussions—sets of statements and their support—as multi-winner elections, where the outcome is a set of viewpoints, each viewpoint being a set of comments expressing a certain position. The goal is to produce a “representative summary” of such discussions, in which different positions are represented according to the support they receive. Technically, we represent discussions as abstract argumentation frameworks [[Dung, 1995](#), [Baroni et al., 2018](#)], a logical formalism that models relations between statements and identifies sets of statements that can be jointly accepted. We use this toolbox to identify the viewpoints expressed in the discussion, and then apply techniques from [ABC](#) voting to select a representative

sample of them, casting viewpoints as candidates, users' support for statements as votes, and the resulting summary as the winning committee. We propose concrete algorithms for this setting, in particular based on the aforementioned [OWA](#) operators. Our method is fully explainable, in the sense that users can be given reasons why each viewpoint appears in the summary (or not). Moreover, the resulting summaries are extractive, that is, they consist of sets of comments taken from the original discussion rather than text generated *ex novo* by a language model. As in Chapter 4, we approach this theoretically, through axioms and theorems, and via simulations, assessing the performance of the proposed methods. We provide recommendations on which rule to use depending on the desired properties of the summary.

- Chapter 6 concludes with a summary and some research directions.

1.4 Further Details

Other publications were not included in this thesis, though all of them are more or less related to this setting. We briefly mention them in chronological order. In one paper [[Igarashi et al., 2024](#)], we study another type of repeated collective decisions, namely, the fair distribution of indivisible resources and tasks over time. Another work, presented at a workshop [[Maly et al., 2025](#)], deals with strategic manipulation under the assumption that voters misrepresent their preferences for the sake of the group rather than their own interests. We explore the possible paradoxical consequences of this behaviour in the context of single-winner and multi-winner elections. In a demonstration paper, we showcase an application based on [ABC](#) voting for selecting a proportionally representative set of comments on social media [[Maly et al., 2026](#)]. This goes in a similar direction to what is discussed in Chapter 5—albeit with slightly different theoretical underpinnings (that is, we do not use formal argumentation theory). Finally, in a position paper [[Grandi et al., 2026](#)], we outline a research agenda that combines social choice theory with modern text-processing algorithms to design and analyse collective decision-making procedures.

The image on page [xvii](#) is intended as an evocative illustration of proportional representation. It consists of eight blocks. To determine the number of pixels of each colour in each block, we use Sainte-Laguë's rule for party-list parliamentary elections [[Sainte-Laguë, 1910](#), [Balinski and Young, 2001](#)] which plays a prominent role in Chapter 3. In particular, we apply it to the vote totals of the ten largest parties in the 2022 Italian general election; similar colours correspond roughly to similar parties. The largest block has as many seats as it has pixels. The other blocks have progressively fewer seats: the second-largest block has seven seats, the third-largest has six, and so on down to the smallest block, which has only one seat. For these blocks—in which the number of pixels does not match the number of seats—we scale the resulting colour distributions to fill them. Finally, the arrangement of pixels within each block is chosen purely for visual

appeal and carries no additional meaning, though the same placement algorithm is used throughout all blocks.

Background

In this chapter, we will introduce the background material. It is structured around two main parts. In the first part, we present a brief history of approval-based multi-winner voting in general. In the second part, we present our mathematical notation and a concrete model: [Approval-Based Committee \(ABC\)](#) voting.

2.1 Historical Background

In this section, we will trace the history of approval-based multi-winner voting. We will follow two separate branches: the history of social choice theory, which is the mathematical study of voting, and the history of proportional representation. Both informed the development of the field.

2.1.1 The Classical Story: Electing a Single Winner

Social choice theory [[Arrow et al., 2002](#)] can be defined as the study of preference aggregation, or, more broadly, as the study of collective decision-making. Social choice theory focuses on the *mechanical core* of preference aggregation—the voting rule—and conceptualises voting rules as mathematical functions or, more recently, algorithms. Of course, voting involves much more than the rule itself, including behavioural and institutional dimensions, but these aspects lie outside the traditional focus of social choice theory and are instead studied in other branches of political science (see, e.g., the volumes edited by [Gallagher and Mitchell \[2005\]](#) and [Herron et al. \[2018\]](#)).

The idea of voting as a means of resolving disputes within groups is antique, but the scientific inquiry into how voting should be done is far more recent. One way to frame this history is the following question: “How should we vote when there are more than two alternatives to choose from?” Indeed, deciding between two alternatives is straightforward (we pick the alternative supported by the majority of voters), but the problem is not as trivial when there are three or more. In particular, the naive solution

of electing the alternative nominated by the majority of voters (the “plurality rule”) is seriously flawed—it might select an alternative that a majority of the population believes to be the worst. This was already observed in antiquity by Roman senator Pliny the Younger, in AD 105 [McLean et al., 1995].

The earliest positive work on voting theory, and the above question in particular, can be traced back to Ramon Llull (1232–1316), Catalan philosopher from Mallorca, whose ideas on voting are discussed by Colomer [2013]. Llull was concerned with elections within the Catholic Church, such as those for bishops and abbots. He proposed reducing the problem of choosing from three or more candidates to the two-candidate case. In particular, he proposed a system based on exhaustive pairwise majority votes among all pairs of candidates—with the candidate winning the most “majority contests” being elected. As we shall see, this predates some of the most influential ideas in the history of social choice, and Llull’s voting system was independently reinvented in modern times by Copeland [1951] (although the two methods differ slightly in how ties are handled). However, Llull’s work on voting remained largely unknown until it was rediscovered in the 20th century. As a side note, Llull is also considered by some as a precursor of computer science [Bonner, 2007, Bonner et al., 2025], showing that the modern algorithmic focus of social choice theory (which we will discuss later) has ancient origins.

Given Llull’s obscurity until recent times, the origins of social choice theory are typically linked to the French Revolution (as argued, e.g., by Suzumura [2010]). Enlightenment thinkers, concerned with human rights and guided by a strong faith in Reason, turned their attention to the study of voting. In particular, mathematician Jean-Charles, chevalier de Borda (1733–1799) and philosopher Marie Jean Antoine Nicolas de Caritat, Marquis de Condorcet (1743–1794) disagreed on the proper way to vote when there are more than two alternatives [Suzumura, 2010, McLean, 2019].

Borda proposed a method based on a system of points associated with each position in a preference ranking [Borda, 1781]. Condorcet, in his *Essai sur l’application de l’analyse à la probabilité des décisions, rendues à la pluralité des voix* [Condorcet, 1785], advocated for pairwise majority voting—a system similar to Llull’s. Condorcet realised that the system he proposed was incompatible with Borda’s, meaning that Borda’s positional voting system might fail to elect a pairwise majority winner [McLean, 2019]. Yet Condorcet also recognised a flaw in majority voting. In particular, majority preferences can sometimes be cyclical, in a scenario now known as *Condorcet’s paradox*. Consider the following table of preference rankings, where three voters called Arrow, Black, and Carroll rank three alternatives from best to worst in their respective columns:

Arrow	Black	Carroll
<i>a</i>	<i>b</i>	<i>c</i>
<i>b</i>	<i>c</i>	<i>a</i>
<i>c</i>	<i>a</i>	<i>b</i>

Here, a majority of voters (Arrow and Carroll) prefer *a* to *b*; another majority of voters (Arrow and Black) prefer *b* to *c*; and another majority (Black and Carroll) prefer *c* to

a , leading to a cycle. Condorcet proposed a voting method that still elects a winner when such a paradox arises. However, how to deal with this shortcoming of majority voting remains debated. Indeed, scholars continue to propose alternative ways to extend pairwise majority voting to situations involving cyclical majorities. Recent examples include the popular Schulze method [Schulze, 2011], which is used by many organisations [Schulze, 2025], the Split Cycle method, proposed by Holliday and Pacuit [2023], and the River method, proposed as recently as this year by Döring et al. [2026].

Interestingly, the opposition between Borda’s and Condorcet’s conceptions predates them. In fact, Borda’s system was discovered a few centuries earlier by Nicholas of Cusa (1401–1464), though this was unknown until the last century. Nicholas of Cusa knew of Llull’s work on voting, but instead argued for the same system later proposed by Borda [McLean, 2019]. It is interesting to observe the “Borda-Condorcet” (or Cusa-Llull) conflict appearing and reappearing independently, suggesting the natural character of these incompatible methods.

After Condorcet and Borda, some intellectuals began to consider the design of voting rules. A notable example is Charles Dodgson a.k.a. Lewis Carroll (1832–1898), the writer of *Alice in Wonderland*. He rediscovered Condorcet’s paradox (and introduced the modern term of “cycle” to describe the phenomenon [McLean, 2019]) and proposed a rule to elect a candidate when such a paradox occurs [Dodgson, 1876, Caragiannis et al., 2016]. Another milestone is the work of Scottish economist Duncan Black (1908–1991). Black [1948] explicitly sets out to bring the issue of voting to the attention of economists, and identifies natural assumptions on the preferences of voters that are sufficient to exclude the occurrence of Condorcet’s paradox (the case of “single-peaked” preferences, which may arise, for example, when voters must agree on a numerical quantity). Although Black himself seemed unaware of Condorcet’s and Dodgson’s work at the time of his 1948 article (see the preface of Black’s later book on the theory of elections [Black, 1958]), he is credited by Campbell and Tullock [1965] with popularising these ideas to modern audiences:

The investigations in this field of his principal predecessors, Condorcet and Lewis Carroll, had made no impact on the intellectual community of their day and had been completely forgotten. Their work is known today only because Black, after discovering the phenomenon himself, discovered his predecessors.

Shortly after Black’s work, modern social choice theory was born. The precise moment in time can be traced to the publication of *Social Choice and Individual Values* [Arrow, 1951], by Kenneth J. Arrow (1921–2017), an American economist. “Modern” here means that, while older thinkers were concerned mostly with the design and critique of *specific* voting rules, Arrow tried to think of voting rules in general and systematically analyse them. He formalised voting rules mathematically as functions that map vectors of strict orders (the voters’ preferences over some set of alternatives) to a single strict order (the collective order). Such a function is called a *social welfare function*. Moreover, Arrow also

formalised normatively appealing properties of voting rules (for example, “all voters should be treated equally”, or “if everyone prefers *this* to *that*, then *that* should not win”) as mathematical expressions called *axioms*.¹ Arrow’s most important contribution, his *general possibility theorem* (now known as Arrow’s impossibility theorem), shows that no voting rule satisfies a small set of axioms that appear, at first sight, uncontroversial and mild. This theorem is one of the most well-known results in social choice theory, if not the most well-known, and shows that voting rule design is necessarily a matter of compromise. Interestingly, we can think of Arrow’s theorem as strengthening and generalising Condorcet’s observation regarding the paradoxical nature of majority voting. Indeed, Arrow’s most controversial and demanding axiom, the *independence of irrelevant alternatives*, states that the social preference between two alternatives *a* and *b* can depend only on the voters’ preferences for these two alternatives. This generalises majority voting, which holds that *a* should be socially preferred over *b* whenever a majority of voters prefer *a* to *b*. Arrow’s theorem shows that this weaker formulation of Condorcet’s (and Lull’s) idea is already paradoxical.

In preparing the grounds for Arrow’s theorem, we focused on the history of voting theory *per se*—Lull, Borda, Condorcet, Dodgson, and Black. However, Arrow’s framework did not come out of a vacuum, and had important precursors in welfare economics, in a tradition initiated by the utilitarianism of Bentham [1789] (a contemporary of Borda and Condorcet) and culminating in the study of ordinal social welfare functions [Bergson, 1938, Samuelson, 1947]. See the chapter by Suzumura [2010] for details on the welfare economics perspective.

Since Arrow’s theorem, several fundamental (mostly negative) results have been proven within his axiomatic framework. To cite one influential result, Gibbard [1973] and Satterthwaite [1975] have independently shown that no voting rule can be immune to manipulation whilst simultaneously satisfying a few natural axioms. Manipulation here means misrepresenting one’s preferences to achieve a better outcome—strategic voting. Manipulation and coping with the Gibbard–Satterthwaite theorem remain major themes in social choice theory [Conitzer and Walsh, 2016, Meir, 2018].

Note that the setting where ballots are rankings is the dominant one in classical social choice theory, but other formats have been studied. The main alternative is our focal format, *approval voting*, where voters can approve or not approve each candidate; that is, they submit a subset of candidates they approve. In a context where a single winner is to be elected, typically the candidate approved most often is chosen. This voting method has a long history: for example, it was employed in ancient Sparta [Girard, 2010], in the elections of the *Doge* in the *Serenissima* Republic of Venice [Lines, 1986], in eighteenth-century Russia for local elections [Aleskerov, 2005], and in papal elections [Colomer and McLean, 1998]. However, in modern times, its rediscovery is attributed to Steven J. Brams and Peter C. Fishburn [Brams and Fishburn, 1978, Brams, 1980, Brams and Fishburn, 1983], though Robert J. Weber is credited with coining the term

¹Notably, Arrow studied mathematical logic with Alfred Tarski (see the biographical sketch by Starr [2018]), and perhaps his approach on formalising voting was influenced by that.

“approval voting” [Weber, 1978, 1995]. Approval voting is curious in that, despite the substantial scholarly attention it has received and its adoption as a voting rule in many scientific societies, it has been largely ignored in real-world political elections [Brams and Fishburn, 2010]. We already mentioned in Chapter 1 that approval voting combines simplicity and expressivity. To give another example of its theoretical appeal, we note that Brams and Fishburn [1978] argue that this rule is resistant to manipulation and conducive to sincere voting behaviour, in contrast to what the Gibbard-Satterthwaite theorem states for ranked ballots. However, this conclusion was part of a lively dispute in the 1980s sparked by Niemi [1984] (see the reconstruction by Laslier and Sanver [2010b]).

This takes us almost to the end of our short story, the “computational turn”. So far, voting rules have been mostly conceptualised as functions. However, they can also be seen (and arguably, more naturally) as algorithms—indeed, they must be *computable*. However, computability might not be enough; one might want to restrict attention to those rules that are *efficiently* computable. Starting in the late 1980s, a number of theoretical computer scientists have begun to study the computational complexity of voting rules. In a diptych that is traditionally cited among the first works in this direction, Bartholdi III et al. show that there are voting rules that are hard to compute [Bartholdi III et al., 1989b] and rules that are easy to compute but hard to manipulate (which has been called a “computational barrier” to manipulation) [Bartholdi III et al., 1989a]. Interestingly, the two rules are, respectively, Dodgson’s rule and a variant of Copeland’s (or Lull’s) rule, both mentioned above. Since then, there has been extensive research at the interface between social choice theory and computer science, in a subfield called computational social choice, a comprehensive overview of which can be found in the handbook by Brandt et al. [2016].

It is in this algorithmic setting that the study of multi-winner voting, especially under approval ballots, has risen to the forefront of social choice research [Faliszewski et al., 2017]. Indeed, multi-winner elections inherently have a combinatorial flavour, since one must choose a subset of winners from a larger set of candidates. Many fundamental questions in this setting concern how such subsets are selected, compared, or optimised, which often leads to computationally challenging problems. However, before delving into this further, there is another branch of history that pours in approval-based multi-winner voting: the history of proportional representation. We will focus first on that.

2.1.2 The Other Story: Proportionality and Approval-Based Multi-Winner Elections

The prototypical example of a multi-winner election is the election of a representative body—for example, parliamentary elections or municipal boards. The goal is to select a number of representatives from a larger pool of candidates. According to UNESCO [2012], the earliest documented parliamentary use—at least in the European tradition—of representative bodies was in 1188, when Alfonso IX, King of León, gathered members

of the nobility to gain political support, and the noblemen acted as representatives of their subjects [Dalaqua, 2025]. How can one go about selecting such a body?

One way to make this selection is called block (or bloc) voting. According to Janson [2018], it has been a popular way of electing multi-member bodies in history, at least from the Middle Ages (for example, in British elections [Gallagher and Mitchell, 2005, page 158]) to the 19th century, and is in some countries still in use [ACE Electoral Knowledge Network, 2026]. Say that k candidates have to be elected. We allow voters to freely nominate a block of at most k candidates, and the k candidates that the most voters nominate are elected. This method has the advantage of being straightforward to explain, easy to execute, and of allowing voters to express themselves freely without having to choose a party. However, it is exceedingly majoritarian—any minority is completely ignored by block voting.

An alternative approach would be to distribute the seats among the candidates in a way that matches the distribution of opinions in the voting population—that is, *proportionally*. Proportional representation is a special kind of political representation where, in the words of Black [1958], “the microcosm is to be a true reflexion of the macrocosm”. The history of this idea is vast, so it cannot be told in full here. Two sources are the books by Hoag and Hallett [1926, Chapter IX] and Carstairs [2013]. The book by Balinski and Young [2001] (originally published in 1982) focuses on the equivalent problem of allocating seats in the United States House of Representatives, where states are assigned a number of representatives in proportion to their populations. We will next focus on some salient moments of this history, which partly interpolates with the history of social choice theory told above.

Thinkers began reflecting on proportional representation around the end of the 18th century in France, England, and the United States. One early and striking formulation of the principle of proportionality is due to the French intellectual Mirabeau (1749–1791), a key figure of the French Revolution (quoted from the book of Hoag and Hallett [1926]):

A representative body is to the nation what a chart is for the physical configuration of its soil: in all its parts, and as a whole, the representative body should at all times present a reduced picture of the people—their opinions, aspirations, and wishes, and that presentation should bear the relative proportion to the original precisely as a map brings before us mountains and dales, rivers and lakes, forests and plains, cities and towns. The finer should not be crushed out by the more massive substance, and the latter not be excluded; the value of each element is dependent upon its importance to the whole and for the whole.

How to exactly achieve this ideal, however, is not a trivial problem. Again in France, Condorcet himself proposed in 1793 a semi-proportional system for the election of district boards, in which each citizen could vote for two candidates (and the boards counted 18 winners). In the same year, Louis Antoine de Saint-Just (1767–1794) proposed treating all of France as a single district in parliamentary elections and allowing each voter to nominate a candidate (no parties involved): an early example of a truly proportional sys-

tem. However, Saint-Just’s proposal was quickly shut down by the “violent opposition” of Robespierre [Hoag and Hallett, 1926, Chapter IX].

In Britain, in the early 19th century, Thomas Wright Hill proposed a system now known as the [Single Transferable Vote \(STV\)](#), where each candidate requires a certain quota to be elected, and surplus votes (to be selected randomly) are returned to their casters, who can now vote for their second preference. Variations of this idea (in which, among other differences, instead of returning ballots to their owners, voters directly rank candidates on their ballots) are still used around the world (for example, in Ireland and Malta [ACE Electoral Knowledge Network, 2026]). A particularly famous version was proposed in 1855 by Carl Christoffer Georg Andr   for Danish elections, and later popularised by the Englishman Thomas Hare (1806–1891).² Here, the aforementioned quota is computed as the number of cast ballots divided by the number of members to be elected—a quantity known as the “Hare quota”. While [STV](#), similarly to block voting, has the attractive feature of allowing voters to freely express their preferences for individual candidates, it is not a perfect system. An early critic of [STV](#) was someone we already met, Dodgson, who considered it to be fundamentally flawed, because it could elect the “wrong” set of candidates [McLean, 2019].³ Indeed, it is not clear how the surplus votes should be divided.

An alternative approach that circumvents the problems of [STV](#) is party-list elections, where each voter supports one of the predetermined parties available. Seats are then distributed to parties according to some formula—the voting rule. According to the [ACE Electoral Knowledge Network](#) [2026], these are the dominant type of proportional-representation elections today. Two prominent ways of resolving such elections were introduced by Belgian mathematician Victor D’Hondt [D’Hondt, 1878] and by the French mathematician Andr   Sainte-Lagu   [Sainte-Lagu  , 1910].⁴ For clarity, we will present these methods in a modern language, not how they were first introduced. Both belong to a family of methods known as divisor methods, which work by assigning seats to parties sequentially. At each step, a weight is assigned to every party, and the party with the highest weight wins the seat. The weight of a party i is computed as the ratio between the number of votes cast in favour of this party and a function d of the seats

²Hare was praised by the philosopher John Stuart Mill, an influential advocate of representative government, for the invention of this system, which he himself endorsed [Mill, 1865]. Hare independently rediscovered this voting method shortly after Andr   [Hare, 1857].

³It is interesting to note that, besides single-winner majority voting, Dodgson was also interested in the mathematics of proportional representation. McLean [2019] claims that Dodgson’s work on proportional representation [Carroll, 1884] is “the earliest known work to discuss both the assignment of seats to each of a number of multi-member districts (the apportionment problem) and the assignment of seats within each district to the parties (the proportional representation problem)”, two distinct problems whose congruence was not fully understood until the second part of the 20th century. Dodgson’s work on proportionality was also discussed by Duncan Black [McLean et al., 2012].

⁴It is to be noted, however, that the same methods were first introduced by Thomas Jefferson and Daniel Webster, respectively, in the context of apportionment of seats in the United States House of Representatives [Balinski and Young, 2001].

assigned to this party so far:

$$\frac{\text{votes for party } i}{d(\text{seats given to } i \text{ so far})}$$

D’Hondt’s method sets $d(\ell) = \ell + 1$, while Sainte-Laguë’s sets $d(\ell) = 2\ell + 1$. Another method in this class is Adams’ method, proposed by the sixth president of the United States, John Quincy Adams, which sets $d(\ell) = \ell$; this method is not adopted as widely as D’Hondt’s and Sainte-Laguë’s, but will play a crucial role in Chapter 3. Note that divisor methods have been systematically studied through the axiomatic approach by [Balinski and Young \[2001\]](#) and, more recently, [Pukelsheim \[2017\]](#).

These three methods—D’Hondt’s, Sainte-Laguë’s, and Adams’—differ in technical details. Still, they can all be seen as distinct concrete ideas for achieving the abstract notion of proportional representation. However, they all force voters to vote for all the members of a specific party, and no one else. On the other hand, block voting allows voters to express themselves freely, but it is not proportional. For example, block voting was used in Sweden in the late 19th century, and as political parties became organised, it became clear that block voting concentrated all the power in the largest party [[Janson, 2018](#)]. In fact, Sweden (and other northern European countries) underwent a period of parliamentary reform in this period due to these issues [[Carstairs, 2013](#), [Janson, 2018](#), [Camps et al., 2024](#)]. It is in this context that we find three important figures in the history of multi-winner voting, the Swedes Lars Edvard Phragmén (1863–1937) and Gustaf H. Eneström (1852–1923) and the Dane Thorvald N. Thiele (1838–1910). All were mathematicians, and between 1894 and 1896, they proposed proportional methods for parliamentary elections in which voters could freely nominate (approve of) any set of candidates without having to pick a specific party [[Phragmén, 1894, 1895](#), [Thiele, 1895](#), [Eneström, 1896](#)]. These are among the earliest examples of approval-based multi-winner voting as studied today.

Eneström and Phragmén proposed an approval-based analogue to STV.⁵ Phragmén later proposed another method, which was inspired by D’Hondt’s method for party-list elections (Phragmén called his method a generalisation of D’Hondt’s), but with the expressive freedom of block voting [[Janson, 2018](#)]. This method is still rather popular in the social choice theory literature [[Lackner and Skowron, 2023](#), [Brill et al., 2024](#)]. Thiele, whose ideas will be central to this thesis, wrote his contribution as an improvement of Phragmén’s latter method. Notably, Thiele introduced a family of methods parametrised by a weight function, similarly to the family of divisor methods outlined above (in fact, as we shall discuss in Chapter 3, the family of Thiele methods can be seen as a generalisation of divisor methods [[Lackner and Skowron, 2021](#)]). Interestingly, resolving the election under (essentially) all of these methods entails solving a hard computational problem [[Aziz et al., 2015](#), [Skowron et al., 2016](#)]. Thiele himself was aware of this and introduced two ways of computing an approximate result (one of which we will also discuss). This prefigures the centrality of computational issues in multi-winner voting.

⁵It is unclear, among Phragmén and Eneström, who came up with this method first [[Janson, 2018](#), [Camps et al., 2024](#)].

In 1918, a different (and in some ways, pioneering) perspective on elections with multiple winners and proportionality was offered by Gottlob Frege (1848–1925), the famous logician, in a manuscript that has been recently rediscovered (published in a collection edited by [Gabriel and Dathe \[2000\]](#)). Frege proposed an electoral system for repeated single-winner elections aimed at guaranteeing proportionality over time. One can interpret this as a kind of multi-winner election, where the “multiple winners” are not a board of representatives or a committee, but the sequence of elected winners. This temporal perspective is now an emerging field in computational social choice (see, e.g., the work of [Lackner \[2020\]](#), or the survey by [Elkind et al. \[2024a\]](#)), and will be featured in this thesis. For a mathematical analysis of Frege’s voting rule, see the work of [Harrenstein et al. \[2020\]](#).

The next big leap in the theory of representative multi-winner voting came in the 1980s, after the Arrovian revolution (and after the popularisation of approval voting by Brams and Fishburn). While up to this moment several thinkers had proposed interesting voting rules (particularly for parliamentary elections), from this point onwards, scholars had the tools to think about the problem more systematically. The watershed contribution was that of [Chamberlin and Courant \[1983\]](#). They remarked on how “social choice theorists have paid almost no attention to methods of selecting representative bodies”.⁶ [Chamberlin and Courant](#) cast this problem as that of electing a representative committee—an interpretation that is still dominant today [[Lackner and Skowron, 2023](#)]. They proposed and studied a rule aimed at maximising “representativeness” of the selected committee. The rule was defined as a variant of the Borda method for single-winner elections and assumed that voters have linearly ordered preferences over the candidates (however, as we will see, Thiele already proposed an analogous rule in the context of approval voting).⁷ Later, [Monroe \[1995\]](#) defined a similar rule to the one proposed by [Chamberlin and Courant](#). [Monroe](#) similarly mainly worked with linear orders, but he also considered approval ballots, which he considered to be “the system [he] most strongly recommend for consideration as a serious alternative by electoral reformers”. He claimed this because of several reasons, including the purported resistance of approval balloting to manipulation. Significantly, this was among the first papers in modern social choice theory to consider the problem of representation in committee elections with approval ballots. Moreover, it is to be noted that [Monroe \[1995\]](#) did not give a precise explanation about how to implement his proposal; in an early example of an algorithmic paper dealing with multi-winner voting, [Potthoff and Brams \[1998\]](#) showed how to use integer programming [[Wolsey, 2020](#)] to compute the outcome of this rule as well as the rule proposed by [Chamberlin and Courant \[1983\]](#).

⁶As exceptions, they cite the work of [Black \[1958, Chapter 11\]](#), [Fishburn \[1973, Chapter 3\]](#), and [Good and Tideman \[1976\]](#).

⁷Around the same time, Peter C. Fishburn also worked on committee elections for linearly ordered [[Fishburn, 1981b](#)] and approval [[Fishburn, 1981a](#)] preferences. We do not focus on his contributions, as [Fishburn](#) concentrated on a Condorcet-like (that is, majoritarian) perspective, ignoring the issue of representation. Observe that, despite being published earlier than the work of [Chamberlin and Courant \[1983\]](#), the work of [Fishburn \[1981b\]](#) cites an earlier version of the paper by [Chamberlin and Courant](#).

In the 1990s and early 2000s, following these seminal papers, more authors began studying approval-based multi-winner elections. The chapters by Kilgour [2010] and Laffond and Lainé [2010], both in the *Handbook on Approval Voting* [Laslier and Sanver, 2010a], contain some references to these early works. Soon after, the emergent field of computational social choice took interest in multi-winner voting—as mentioned above, this is natural, as multi-winner elections are inherently combinatorial. Let us mention some of the earliest papers in this direction. One example is the work of LeGrand [2004] and its follow-up by LeGrand et al. [2007], which analysed the complexity of an approval-based multi-winner rule known as the *minimax procedure* [Brams et al., 2007]. We briefly revisit this rule later in the chapter. Another example is due to Procaccia et al. [2007] and its follow-ups and extensions [Meir et al., 2008a,b], which studied the computational complexity of winner determination and manipulation across many formats of multi-winner elections, including approval voting. A further early example of the computational study of multi-winner voting is the analysis by Procaccia et al. [2008] of the rules proposed by Chamberlin and Courant [1983] and Monroe [1995] mentioned above. From these early examples to today, multi-winner voting has received lots of attention and has become one of the main research areas in computational social choice. Many references to this early work can be found in the book chapter by Faliszewski et al. [2017], which covers both approval-based and non-approval-based elections.

We do not single out the contributions of computer scientists arbitrarily. Our interest lies in the fact that, in recent years, some of the most important advances in proportionality for multi-winner voting have come from that community. Most of these concern committee elections with approval ballots (i.e., *Approval-Based Committee (ABC) voting*), the *de facto* standard multi-winner voting model in computational social choice. A prime example would be the influential paper by Aziz et al. [2017]. This work, whose influence will be felt also in this thesis, introduced a precise definition of what it means for an outcome to be proportionally representative in the context of ABC elections. Indeed, in this setting, the notion of proportionality is far less clear to define than in parliamentary elections (we will discuss this in Section 2.2.8). Aziz et al. [2017] formalised proportionality through an axiom called *Justified Representation (JR)*, requiring that if a group of voters is large enough to control a seat in the outcome and their preferences are sufficiently similar, at least one voter in this group should approve some member of the selected committee. Aziz et al. [2017] also introduced a strengthening of JR, *Extended Justified Representation (EJR)*, and showed that among the rules introduced by Thiele, only one rule, *Proportional Approval Voting (PAV)*, satisfies EJR. Interestingly, PAV corresponds to D’Hondt’s method in scenarios where candidates are organised into parties, and, for this reason, was recommended by Thiele for elections where proportionality is desirable.

The axioms proposed by Aziz et al. [2017] have shaped the research on proportionality in multi-winner voting within computational social choice. Their work has sparked a large body of subsequent research—though as we argue in Chapter 3, this is only one of many possible ways of conceptualising proportional representation in ABC voting. Since then, several other notions of proportionality have been proposed, and our

understanding of which rules satisfy which axioms has progressed. A comprehensive overview of this line of work is offered by [Lackner and Skowron \[2023\]](#). In particular, Chapter 4 of their book presents several proportionality axioms beyond JR and EJR, with varying degrees of strength and different conceptual underpinnings, and highlights their logical relations. The book also collects other strands of literature concerning the ABC model: foundational axiomatic work (including, for instance, the study of voting paradoxes and undesirable behaviour), algorithmic issues (such as the design of fair voting rules that are easy to compute or the approximation of rules that are hard to compute), and strategic considerations (capturing the tension between manipulability and fairness).

This brings us to the present day. We conclude with an outlook. Though significant problems remain open [[Lackner et al., 2023b](#), Chapter 7], proportionality in multi-winner elections, particularly in ABC voting, is by now relatively well understood. More recently, however, ABC voting has been viewed not only as a theory of collective choice, but as a general-purpose formalism for modelling proportional representation problems. Recent research in computational social choice has focused on adapting ideas from ABC voting to other decision-making settings, most prominently participatory budgeting, where citizens vote on how to distribute a city’s budget among some projects [[Wampler et al., 2021](#)]. In the standard model [[Rey et al., 2025](#)], participatory budgeting is a generalisation of ABC voting, and many notions of proportionality in this setting build directly on those developed for ABC voting (including EJR). Other examples include proportionally clustering data points [[Kellerhals and Peters, 2024](#)], selecting a representative set of matchings among individuals [[Boehmer et al., 2025a](#)], and proportionally recommending comments in social media [[Revel et al., 2025](#)]. Finally, [Masařík et al. \[2024\]](#) recently introduced a “generalised theory” of proportionality that extends the theory of ABC voting. Their model, essentially a generalisation of the ABC voting framework that accommodates feasibility constraints, captures several scenarios. These include committee elections with diversity constraints (e.g., quotas), repeated single-winner elections, referenda consisting of multiple sub-questions, and many others. This line of work provides a unified view of proportionality across all these settings.

This concludes our short history of proportional representation in approval-based multi-winner elections.

2.1.3 Summary and Connections to the Thesis

We have presented the history of approval-based multi-winner voting from two perspectives—the history of social choice theory in general and the history of proportional representation specifically. Some important features that emerged, and that are relevant to our analysis, are:

- the conceptualisation of a voting rule as a mathematical function and the use of the axiomatic method;
- the emphasis on strategic and algorithmic issues;

- the importance of proportionality and representation as fundamental desiderata in multi-winner elections, in all their embodiments: parliamentary elections, committee elections, but also repeated single-winner elections *à la* Frege;
- the use of approval voting, a simple, expressive, and “sincere” ballot format;
- the recent use of ABC voting as a building block to obtain a theory of proportionality in other settings, such as participatory budgeting.

Let us briefly connect this to the thesis. We will draw on the theory of proportional representation in parliamentary elections to define new proportionality axioms for ABC voting. Then, we will study a model analogous to Frege’s, in which representation is understood over time, and focus on the strategic and computational issues specific to that model. The last contribution leans on the generality of the formal model of multi-winner elections and the recent interest in the connection between social choice and computer science. We will show how one can use the toolbox of ABC voting in a practical application, namely, obtaining an extractive summary of a discussion that is representative of the opinions of its participants.

This concludes the historical overview. The next part will focus on the formal details of the thesis.

2.2 Technical Background

In the previous section, we have presented the history of multi-winner voting. We have touched upon the fact that there are many possible types of elections that embody this—parliamentary elections, committee elections, repeated elections. In this section, we present the main model used currently in social choice theory to formalise approval-based multi-winner voting: ABC voting.

ABC voting is not the only model we will use, but due to its centrality, and due to its formal and conceptual connections to the other models we will study, it is worth introducing in depth.

Before doing so, we will take a detour, and present further technical background used throughout the thesis.

2.2.1 Basic Mathematical Conventions

We typically indicate sets as capital letters, scalars as lowercase letters, and vectors in boldface. Furthermore, we indicate the entries of a vector by the same letter, i.e., $\mathbf{a} = (a_1, a_2, \dots)$.

The sets $\mathbb{R}$, $\mathbb{Q}$, and $\mathbb{Z}$ indicate the real numbers, rational numbers, and integers, respectively. We indicate with $\mathbb{N}$ the set of natural numbers excluding zero, i.e., $\mathbb{N} = \{n \in \mathbb{Z} : n > 0\}$. Furthermore, we define $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$, $\mathbb{R}_{\geq 0} = \{\alpha \in \mathbb{R} : \alpha \geq 0\}$, and $\mathbb{R}_{>0} = \mathbb{R}_{\geq 0} \setminus \{0\}$. Moreover, given $\ell \in \mathbb{N}$, we define $\ell\mathbb{N} = \{\ell \cdot s : s \in \mathbb{N}\}$ (e.g., $2\mathbb{N}$ are the even numbers). Given a set S , we indicate its powerset as 2^S .

As usual, given a number $\ell \in \mathbb{N}$, we define $[\ell] = \{s \in \mathbb{N} : s \leq \ell\}$, while given two numbers $\alpha, \beta \in \mathbb{R}$ with $\alpha \leq \beta$, we define the interval $[\alpha, \beta] = \{\gamma \in \mathbb{R} : \alpha \leq \gamma \leq \beta\}$. Moreover, given a real number $\alpha \in \mathbb{R}$, we define $\lceil \alpha \rceil$ as α rounded up, that is, $\lceil \alpha \rceil = \min \{\ell \in \mathbb{N} : \ell \geq \alpha\}$. Analogously, we define $\lfloor \alpha \rfloor$ as α rounded down.

Furthermore, $\mathbb{1}[\cdot]$ is the indicator function: for a boolean expression φ ,

$$\mathbb{1}[\varphi] = \begin{cases} 1 & \text{if } \varphi \text{ is true,} \\ 0 & \text{otherwise.} \end{cases}$$

Given some number $\alpha \in \mathbb{R}$, we define $\mathbf{a}$ as the vector $(\alpha, \dots, \alpha)$. For instance, $\mathbf{0} = (0, \dots, 0)$. The length of this vector will be stated explicitly if not clear from context. Additionally, for two vectors $\mathbf{a}$ and $\mathbf{b}$ of length ℓ , their dot product is indicated as

$$\mathbf{a} \cdot \mathbf{b} = \sum_{i \in [\ell]} a_i b_i.$$

Next, we sometimes write $a_i \in \mathbf{a}$ to indicate that some entry in $\mathbf{a}$ equals a_i . Finally, a permutation of a vector $(a_1, \dots, a_\ell)$ is a vector $(a_{\rho(1)}, \dots, a_{\rho(\ell)})$ where $\rho : [\ell] \rightarrow [\ell]$ is a bijection.

2.2.2 Propositional Logic and Complexity Theory

We will use notions from propositional logic and complexity theory. We will sketch here some of the main notions, but we assume that the reader is somewhat familiar with these topics. If not, the books by [Van Dalen \[2004\]](#) and [Arora and Barak \[2009\]](#) are standard introductions to propositional logic and complexity theory, respectively.

2.2.2.1 Propositional Logic

We assume familiarity with basic propositional logic: variables, formulae, clauses, interpretations, models, [Conjunctive Normal Form \(CNF\)](#), and so forth. We will state some notational conventions in order to aid the reader. A propositional variable (typically indicated as x) is either true or false. A literal is either a variable or its negation, i.e., x or $\neg x$. A clause is a disjunction of literals, typically indicated as γ ; a clause is true if at least one of its literals is true. Sometimes we treat clauses as sets, i.e., we write $\neg x \in \gamma$. A [CNF](#) formula is a conjunction of clauses; a conjunction of clauses is true if all of its clauses are. We also sometimes write such formulae as set of clauses. An interpretation I is a set of variables, and I is a model of formula φ if by setting all $x \in I$ as true and every other variable as false, then φ is true.

2.2.2.2 Complexity Theory

The notions of (decision) problem, problem instance, algorithm, and runtime are taken as primitive.

P is the class of problems solvable in polynomial time, whereas NP is the class of problems whose solutions we can verify to be correct in polynomial time. It is immediate to see that $P \subseteq NP$, but it is unknown whether $P \neq NP$ holds, though it is generally thought to be true.

A problem is NP -hard if it is at least as hard as any other problem in NP . Formally, this is captured by the notion of *reducibility*. Reducing one problem to another means that we can transform any instance of the former to an instance of the latter, such that yes-instances are mapped to yes-instances (and *vice versa*), and this transformation can be done in polynomial time. NP -hard problems are those to which we can reduce any problem in NP .

A problem is NP -complete if it is NP -hard and it belongs to NP . The prototypical NP -complete problem would be SAT, the problem of deciding whether a propositional formula has at least one model [Cook, 1971, Levin, 1973]. Note that, by transitivity, to show NP -hardness of some problem $PROB$, it is enough to show that one can reduce *one* known NP -hard problem—such as SAT—to $PROB$. Hardness and completeness can be defined analogously for other classes of problems.

Generally, we consider problems in P to be tractable and NP -hard problems to be intractable. In other words, for our purposes, a proof of NP -hardness is strong evidence of intractability.

To discuss the runtime of algorithms, we sometimes use the so called big-Oh notation. Intuitively, we write $O(f(\ell))$ to denote the set of functions that are asymptotically upper bounded by $f(\ell)$ up to a constant factor. For example, $g(\ell) = 2\ell$ belongs to $O(f(\ell))$ with $f(\ell) = \ell$. This helps us classifying algorithms that have comparable runtime (i.e., linear, quadratic, exponential) while abstracting over constants. A lesser known concept we will use is the $O^*(\cdot)$ notation, which suppresses factors polynomial in the input size. For example, a running time of $O(\ell^2 \cdot 2^\ell)$ is written as $O^*(2^\ell)$.

Perhaps the most “obscure” complexity-theoretic notion we will employ is Θ_2^P , the class of decision problems solvable in polynomial time when having access to $O(\log n)$ calls to an NP -oracle [Wagner, 1990]. Here, an oracle is a subroutine that is able to immediately solve an instance of a problem belonging to NP . Conceptually, problems that are hard for this class are believed to be strictly harder than NP -complete problems.⁸ We will recall the definition of this class when we will actually employ it.

Finally, we will use the notion of **Fixed-Parameter Tractability (FPT)**, referring to algorithms running in time $f(k) \cdot \ell^{O(1)}$ for some computable function f of a parameter k (which depends on the input) and input size ℓ . The idea is that super-polynomial complexity is confined to the parameter k , while the dependence on the input size ℓ remains polynomial. In other words, if the parameter k is expected to be small, an **FPT** result yields algorithms that are efficient in practice despite potentially large input size. See the textbook by Cygan et al. [2015] for more details.

⁸Coincidentally, Hemaspaandra et al. [1997] have shown that the problem of determining the winner under the voting system proposed by Dodgson (mentioned above) is hard for Θ_2^P .

2.2.3 Sampling and Simulations

In Chapters 4 and 5, we complement our theoretical analysis with numerical simulations. We assume familiarity with elementary descriptive statistics, including means and standard deviations, and with standard sampling operations, such as using a computer program to sample from probability distributions like the normal distribution. We use these sampling techniques to generate synthetic elections. For a general guide to numerical experiments in computational social choice, see the work of [Boehmer et al. \[2024b\]](#); for a focused discussion of sampling approval elections, see the work of [Szufa et al. \[2022\]](#).

2.2.4 Main Model: Approval-Based Committee Voting

We are now ready to introduce the [ABC](#) voting model. We will give an example at the end of the section.

We indicate as $C = \{a, b, \dots\}$ the finite set of *candidates* with $|C| = m$. These are the alternatives to choose from.⁹ Furthermore, we assume that there are $n \in \mathbb{N}$ *voters*, and that the set of voters is $N = [n]$. We are also given a target *committee size*, indicated as $k \in [m]$. We use variables i and j to refer to an arbitrary voters in N . We will use letters such as a, b, c , and so on, to indicate either an arbitrary or a specific candidate. This will always be clear from context.

The preferences of the voters are given by the so-called *profile*

$$A: N \rightarrow 2^C \setminus \{\emptyset\}.$$

This is a function that maps each voter to her nonempty *approval ballot*. In other words, for every $i \in N$, we have that $\emptyset \subset A(i) \subseteq C$. That is, voters express their preferences by approving a subset of candidates. We remark that we do not ask for voters to approve at most k candidates (as done, for instance, in block voting, discussed in the historical background). Voters can potentially approve all candidates. Moreover, for the sake of generality, we connect each voter's identity to her ballot. This is standard in social choice theory; however, none of the voting methods we will study use this information, hence one could equivalently formalise ballots as anonymous (for example, with a multiset of subsets). We typically denote A with a vector notation, that is, we write $A = (A(1), \dots, A(n))$. Given a candidate $c \in C$, we define $N(c) = \{i \in N: c \in A(i)\}$ as the set of its supporters. We define as $n_c = |N(c)|$ the size of this set. Finally, we define

$$\mathcal{A} = \left(2^C \setminus \{\emptyset\}\right)^N$$

as the set of all profiles.

⁹We adopt the term “candidate” rather than the more neutral “alternative” or “option”, as we mostly work under the interpretation of electing a committee of representatives, but this model is rather general (see Example 2.1). [Lackner and Skowron \[2023\]](#) also mainly use the term “candidate”.

The above discussion covers the input of a voting process, that is, the preferences of the voters. Let us discuss the outcome. A *committee* is a set $W \subseteq C$ of size at most k . If $|W| = k$, we say that W is an *exhaustive* committee. We denote as $\mathcal{W} = \{W \subseteq C: |W| \leq k\}$ the set of all committees and as $\mathcal{W}_k = \{W \subseteq C: |W| = k\}$ the set of all exhaustive committees. We remark that we consider inexhaustive committees for technical reasons; however, an implicit goal of any reasonable voting method is to return an exhaustive committee.

We now formally define voting rules; as discussed in the historical overview, a voting rule is a function whose domain is $\mathcal{A}$, the set of all possible preference profiles. Given that many profiles naturally lead to ties, we formalise a voting rule as a function that returns a nonempty set of committees, the set of (tied) winning committees. Note that, ideally, a voting rule should return a single committee when possible, and that how to resolve ties is not part of our model (for example, one could do random tiebreaking, or use an externally defined order over the candidates). Formally, then, a *voting rule* is a function

$$F: \mathcal{A} \rightarrow 2^{\mathcal{W}} \setminus \{\emptyset\}.$$

Given two voting rules F and G , we say that F refines G if, for every $A \in \mathcal{A}$, it holds that $F(A) \subseteq G(A)$.

We now give an example to clarify the model.

Example 2.1. Let us now give an example on how one might formalise a multi-winner election within our model. To exemplify the generality of it, let us consider a scenario where the committee to be elected is not a committee in the literal sense (that is, a board of human beings). Imagine that a list of 8 books has been shortlisted for a prize. The prize consists of a feature in a literary festival in Vienna, for which there are three available spots. Hence, three books will win the prize. We can model this by letting $k = 3$ and by abstractly representing the 8 shortlisted books as the set of candidates

$$C = \{a, b, c, d, e, f, g, h\}.$$

Suppose that the final selection will be made by a panel of six experts, which we identify with the set $N = [6]$. One of the experts is a *connoisseur* of social choice theory, and hence proposes to use approval voting to express the preferences. The

ballots are given by the approval profile A with

$$\begin{aligned} A(1) &= \{a, b, c\}, \\ A(2) &= \{a, d\}, \\ A(3) &= \{b, c, d, f\}, \\ A(4) &= \{d, e, g\}, \\ A(5) &= \{f\}, \\ A(6) &= \{a, h\}. \end{aligned}$$

This means, e.g., that the voter labelled with “3” votes for (approves) books b , c and f , and nothing else. Which three books win the election? We can use any rule among the (astronomically-many but finite) functions of form $F: \mathcal{A} \rightarrow 2^{\mathcal{W}} \setminus \{\emptyset\}$ to settle the election. For example, one such rule would be

$$F(A) = \left\{ W \subseteq \bigcup_{i \in N} A(i) : |W| = k \right\}$$

which returns all sets of size k of books that have been nominated at least once. This seems like a bad rule (in particular, one needs some serious tiebreaking), but it is nonetheless a possible rule. Another (bad) rule is the constant rule $F(A) = \{\{a, b, c\}\}$, which requires no tiebreaking but ignores the input entirely. A more reasonable rule could be “satisfaction approval voting” [Brams and Kilgour, 2015]:

$$F(A) = \arg \max_{W \in \mathcal{W}_k} \left(\sum_{i \in N} \sum_{c \in W} \frac{\mathbb{1}[c \in A(i)]}{|A(i)|} \right).$$

Here, each voter has one point to give, which she divides equally among all books she approves of. The sets of three books with highest scores are selected (again, tiebreaking might be necessary). For example, voter 1 gives $1/3$ points to a , b , and c (each). In total, alternative a receives $7/6$ points (from voters 1, 2, and 6). Here, the only winning committee is $\{a, d, f\}$.

In the next sections, we will present some concrete voting rules. Note that in the remainder of the chapter, and occasionally throughout the rest of the thesis, we will pair our examples with links to Pref.Tools,¹⁰ an online platform developed by Dominik Peters that allows users to enter an ABC election and compute its outcome. These links allow readers to quickly verify the outcomes presented in our examples and, if desired, experiment with the profiles under different voting rules. The website supports several popular voting rules, as implemented in the ABCVOTING Python package [Lackner et al., 2023b]. The links we provide already encode the input (A and k) and readers may select the voting rules whose outcomes they wish to check by clicking on “Choose rules”. Users can also click on the output of a voting rule to see further details (e.g., check for

¹⁰<https://pref.tools/abcvoting> (accessed: 2026-04-22).

ties). Note that the website labels candidates by the integers $0, 1, \dots, m-1$. We adopt the convention that our candidate a corresponds to 0, candidate b to 1, and so on.

2.2.5 Thiele Rules

We now formally introduce the family of voting rules introduced by Thiele [1895]. As discussed in the historical background, this was one of the earliest contribution on ABC voting; however, these rules continue to attract significant interest [Lackner and Skowron, 2023]. Intuitively, each Thiele rule is defined by a weight function which determines the score that a voter assigns to a committee given the number of candidates she approves in that committee. We define these functions in terms of the marginal contribution of each additional approved candidate to the voter’s score.

Definition 2.1. A *weight function* is a polynomial-time computable function $\omega: \mathbb{N} \rightarrow \mathbb{R}_{\geq 0}$ such that $\omega(1) = 1$ and, for every $\ell \in \mathbb{N}$, $\omega(\ell) \geq \omega(\ell + 1)$. We sometimes write weight functions as vectors.

We assume polynomial-time computability to simplify the complexity analysis in Chapters 4 and 5. This is anyway a reasonable assumption: we are unaware of any weight function requiring super-polynomial time in the literature. We now define Thiele rules formally.

Definition 2.2. Given $A \in \mathcal{A}$, $i \in N$ and $W \in \mathcal{W}_k$, we define

$$\text{score}_\omega(A(i), W) = \sum_{\ell=1}^{|A(i) \cap W|} \omega(\ell) \quad \text{and} \quad \text{score}_\omega(A, W) = \sum_{j \in N} \text{score}_\omega(A(j), W).$$

The *Thiele rule* corresponding to ω is then

$$\text{THL}_\omega(A) = \arg \max_{W \in \mathcal{W}_k} \left(\text{score}_\omega(A, W) \right).$$

When computing the score a voter assigns to a committee, we look at the number of candidates the voter approves in that committee. The first approved candidate contributes $\omega(1)$ to the score; the second, $\omega(2)$; and so on.

To illustrate the definition further, we present some standard Thiele voting rules [Janson, 2018, Lackner and Skowron, 2023].

- **Approval Voting (AV)** is the Thiele rule given by $\omega = (1, 1, 1, \dots)$. Each voter i assigns to committee W a score of $|A(i) \cap W|$. This rule selects the k candidates that are approved most often and is arguably the most natural rule one can think of. Thiele called this the “strong” method and noted its lack of proportionality.

- **Chamberlin-Courant Approval Voting (CC)** is the Thiele rule given by $\omega = (1, 0, 0, \dots)$. Each voter i assigns to committee W a score that is 1 if i is covered by W (i.e., approves at least one candidate in it), and 0 otherwise (the score of a committee does not increase if a voter approves more than one candidate). This rule selects the committees that cover as many voters as possible. The name of this rule comes from the two scholars who popularised it [Chamberlin and Courant, 1983], whom we covered in the historical overview. However, Thiele already discussed this as the “weak” method.
- **Proportional Approval Voting (PAV)** is the Thiele rule given by $\omega = (1, 1/2, 1/3, \dots)$. Each voter i assigns to committee W a score corresponding to the $|A(i) \cap W|$ -th harmonic number. Thus, any additional candidate a voter approves in the committee will yield diminishing returns. Forest Simmons gave this rule its name in 2001, when he rediscovered the method [Kilgour, 2010]. Thiele noted that this method was particularly suitable for elections requiring proportional outcomes, as its outcome coincides with D’Hondt’s method when candidates are organised into parties.

As already observed by Thiele, one can think of **AV** and **CC** as lying on two different ends of the spectrum of Thiele rules: while **AV** is entirely majoritarian, **CC** tries to represent all voters whenever possible. **PAV** lies, in some sense, in between these two, trying to represent groups of voters proportionally to their size.

Example 2.2. Let us apply **AV**, **CC** and **PAV** to the preferences expressed in Example 2.1. The two most approved candidates are a and d with 3 approvals, so any winning committee under **AV** must include a and d . The last spot is reserved for either b , c , or f , each of which gets 2 approvals. Hence, the (tied) winning committees are $\{a, b, d\}$, $\{a, c, d\}$, and $\{a, d, f\}$. Under **CC**, the winning committees are $\{a, d, f\}$, $\{a, e, f\}$, and $\{a, f, g\}$, as one can see that these are the only committees that cover all voters. For example, surely we must include f as there is one voter (voter 5) that only approves this candidate. Finally, **PAV** selects only $\{a, d, g\}$.¹¹

2.2.6 Ordered Weighted Averaging Rules

The second family of rules we study has a much shorter history and has seen much less attention in the literature. It is based on the concept of a **Ordered Weighted Averaging (OWA)** operator, introduced by Yager [1988]. These are numeric aggregating functions that generalise standard aggregators such as the mean, the maximum, or the minimum. In short, we use such operators to aggregate the “satisfaction” of the voters with a committee (i.e., the number of candidates they approve in it), and pick the committees yielding highest aggregate satisfaction.

¹¹These results can be verified at: <https://pref.tools/abcvoting/?3&11100000&10010000&01110100&00011010&00000100&10000001> (accessed: 2026-10-05).

Seemingly, this family has not been studied in the [ABC](#) voting model, but only in the related models of committee elections with linear preferences [[Elkind and Ismaili, 2015](#)] and multiple referenda [[Amanatidis et al., 2015](#)]. The latter model is similar to the model we study in Chapter 4.¹²

In the following definition, we will refer to *families* of weight functions—that is, for each number $n \in \mathbb{N}$ of voters, we will have a distinct weight function $\omega^{(n)}$ (Definition 2.1). The reason for this may become clear after the definition but, intuitively, in [OWA](#) methods we assign a weight to each voter depending on her satisfaction relative to that of other voters. This parametric formulation of weight functions allows us to express, for example, that half of the population has weight 1 and the other half 0. Again, we assume one can compute $\omega^{(n)}(\ell)$ in polynomial time for every $n, \ell \in \mathbb{N}$.

Definition 2.3. Given a profile $A \in \mathcal{A}$ and $W \in \mathcal{W}_k$, we define the sorted score vector $s(A, W)$ as the vector $(|A(1) \cap W|, \dots, |A(n) \cap W|)$ sorted in nondecreasing order. Given a family of weight functions $\omega^{(n)}$, the [OWA rule](#) corresponding to $\omega^{(n)}$ is then

$$\text{OWA}_{\omega^{(n)}}(A) = \arg \max_{W \in \mathcal{W}_k} \left(s(A, W) \cdot \left(\omega^{(n)}(1), \dots, \omega^{(n)}(n) \right) \right).$$

In other words, when evaluating a candidate solution W , we consider the satisfaction with W of the voters in nondecreasing order. If there are n voters, the least satisfied voter has weight $\omega^{(n)}(1)$, the second-least satisfied voter $\omega^{(n)}(2)$, and so on. We compute the score of W by summing over the weighted satisfaction of the voters. The exhaustive committees with maximal score are selected.

Note that, while the Thiele and [OWA](#) rules defined by $\omega^{(n)} = (1, 1, 1, \dots)$ are equivalent (they both are [AV](#)), this needs not hold in general.

Example 2.3. Let $k = 3$ and consider profile

$$A = (\{a, b\}, \{a, c\}, \{a, b, d\}, \{a, b, d\}).$$

Recall that this vector notation has to be interpreted as “the i -th entry is the preference of voter i ”. For example, here, both voter 3 and 4 approve a, b , and d .

Consider weight function $\omega = (1, 1/2, 1/3, \dots)$ (this is independent of n). Here, one can check that THL_ω (i.e., [PAV](#)) selects $W = \{a, b, d\}$ (and only this committee).¹³ However, OWA_ω here selects $W' = \{a, b, c\}$. Indeed, while W has sorted satisfaction vector $(1, 2, 3, 3)$, W' has sorted satisfaction vector $(2, 2, 2, 2)$. The latter is preferred by OWA_ω .

¹²Observe that in some papers some authors refer to the family of Thiele rules as [OWA](#) rules [[Skowron et al., 2016](#)]. This is different from our definition.

¹³One can verify this here: <https://pref.tools/abcvoting/?3&1100&1010&1101&1101> (accessed: 2026-10-05).

2.2.7 Sequential Rules

The rules introduced above require optimising an objective function over a combinatorial space. In fact, Thiele rules are in general NP-hard to compute [Skowron et al., 2016]. For OWA rules, one can reasonably expect a similar complexity, but to the best of our knowledge a formal hardness result has not been proven, except for the minimax rule [LeGrand et al., 2007], which is the OWA rule defined by $\omega = (1, 0, 0, \dots)$. Still, large subclasses of OWA rules are hard to compute in the related models of multiple referenda [Amanatidis et al., 2015] and multi-issue voting (which we cover in Chapter 4).

We thus introduce a general way to approximate optimisation-based rules in polynomial time—sequential rules. Interestingly, Thiele already saw the computational problems inherent in the family of rules he proposed, and he already introduced two ways of approximating them [Janson, 2018]. One of the two methods was indeed the sequential version we now introduce.

Definition 2.4. Consider a rule F optimising some score $f : \mathcal{A} \times 2^C \rightarrow \mathbb{R}$, i.e.,

$$F(A) = \arg \max_{W \in \mathcal{W}_k} (f(A, W)).$$

The sequential formulation $\text{seq-}F$ builds a committee by adding candidates one at a time. Starting from $W_0 = \emptyset$, at each step $\ell \in [k]$ it nondeterministically selects

$$c_\ell \in \arg \max_{c \in C \setminus W_{\ell-1}} (f(A, W_{\ell-1} \cup \{c\}))$$

and sets $W_\ell = W_{\ell-1} \cup \{c_\ell\}$. Finally, $\text{seq-}F(A)$ is the set of all committees W_k reachable by some sequence of such choices.¹⁴

All Thiele and OWA rules fit the above schema, and hence their sequential formulations are straightforwardly defined. AV and seq-AV coincide: intuitively, one can observe that picking the k most approved candidates can be reduced to picking the most approved candidate, then the second most approved candidate, and so on until k candidates have been selected. However, this is not true in general.

Example 2.4. Let $k = 2$ and $A = (\{a, b\}, \{a, c\}, \{b\}, \{c\})$. Here, CC simply picks $\{b, c\}$, as this covers all voters. However, seq-CC works differently. In the first step, the candidate that covers most voters is a . Hence, this candidate is selected. In the next step, we need to find the candidate that, together with a , leads to the committee that covers as many candidates as possible. Thus, the winning committees are $\{a, b\}$

¹⁴This approach, in which we return all committees that could be constructed by *some* way of handling the ties, is known as the “parallel universe model” [Conitzer et al., 2009, Freeman et al., 2015, Faliszewski et al., 2017]. We will often use this approach to model sequential rules.

and $\{a, c\}$. This differs from the result of [CC](#), and additionally both tied committees fail to cover all voters.

The previous example shows that, indeed, sequential rules might fail to uphold the normative ideal embodied by the rule they approximate (in this case, covering as many voters as possible). Nonetheless, they are a standard way of approximating computationally hard rules [[Lackner and Skowron, 2023](#)]. Note that for Thiele rules, a sequential rule approximates its optimisation-based counterpart with a ratio of $1 - 1/e \approx 0.63$ [[Skowron et al., 2016](#)].

2.2.8 Proportionality Axioms

We conclude with a specific class of axioms—or properties—for committees and voting rules. The properties we are concerned with are mathematical encodings of what it means for an outcome (or a rule) to be “proportional”. Typically, proportionality is understood as the fact that, if a group of voters is large enough to control a certain fraction of the committee, then it should do so, provided that the voters in this group have similar preferences. That is, we should never underrepresent groups. One of the main themes of [Chapter 3](#) is that this is but one of the ways one can conceptualise proportionality; for instance, we might ask to never overrepresent groups instead. Nonetheless, this is the standard approach.

First of all, we begin with the following definition. It encodes what a “group” deserving of representation even is.

Definition 2.5. For every $\ell \in \mathbb{N}$, a group $N' \subseteq N$ is ℓ -cohesive if $|N'| \geq \ell \cdot n/k$ and $|\bigcap_{i \in N'} A(i)| \geq \ell$.¹⁵

Intuitively, a group of n' voters is ℓ -cohesive if its size is large enough to warrant at least ℓ candidates in the committee (expressed by the relation $n'/n \geq \ell/k$) and if the voters can agree on a common pool of ℓ candidates. One might therefore argue, in the spirit of proportional representation, that all members of such a group should be guaranteed to approve at least ℓ candidates in the committee. However, this cannot be ensured. This has been shown by [Aziz et al. \[2017\]](#), and the impossibility holds even if we were to confine our attention to 1-cohesive groups only. The reason is that cohesive groups may overlap; the argument goes as follows.

Example 2.5 ([Aziz et al. \[2017\]](#)). Consider the 9-voter profile

$$A = \{\{a\}, \{a\}, \{a, b\}, \{b\}, \{b, c\}, \{c\}, \{c, d\}, \{d\}, \{d\}\}$$

¹⁵Note that the quantity n/k is the Hare quota mentioned in the historical background. One could define an alternative notion using, for instance, Droop’s quota, introduced by [Droop \[1881\]](#), defined as $n/(k+1)$. [Casey and Elkind \[2025\]](#) recently explored this direction.

and let $k = 3$. For any of the four candidates, there is a 1-cohesive group of $n/k = 3$ voters that all approve this candidate, and where at least one voter approves *only* this candidate. For example, if we consider candidate a , the first three voters all approve a , and voters 1 and 2 approve only a . It follows that, if we want all members of all 1-cohesive groups to approve at least one candidate in the output committee, we need to select all four candidates, which exceeds our target size.

In light of this impossibility, we need to scale back our requirements for cohesive groups. With this in mind, we present two satisfiable axioms—meaning that for every possible input there exists a committee satisfying them—originally introduced by Aziz et al. [2017]. Before doing so, we fix the following convention. In axiom definitions, we refer to a committee W and profile A . A voting rule F satisfies an axiom if, for all values of k, m and n , for all $A \in \mathcal{A}$, and for all $W \in F(A)$, committee W and profile A satisfy the axiom.

The first axiom requires that at least one voter in every 1-cohesive group approves at least one candidate in the committee.

Definition 2.6. Committee W satisfies *Justified Representation (JR)* if for every 1-cohesive group $N' \subseteq N$ there exists some $i \in N'$ with $|W \cap A(i)| \geq 1$.

The following strengthens JR and generalises it to ℓ -cohesive groups with $\ell \geq 1$.

Definition 2.7. Committee W satisfies *Extended Justified Representation (EJR)* if for every ℓ -cohesive group $N' \subseteq N$ there exists some $i \in N'$ with $|W \cap A(i)| \geq \ell$.

These two axioms are enough to distinguish between the three Thiele methods we mentioned. Indeed, AV is fully majoritarian and fails even JR (let $k = 3$ and think of two voters approving $\{a, b, c\}$ and one voter approving $\{d, e, f\}$ —AV selects $\{a, b, c\}$, ignoring the third voter). On the other hand, CC satisfies JR but fails EJR. Finally, PAV satisfies EJR, and is in fact the only Thiele method with this property.¹⁶ Note that, by contrast, seq-PAV fails even JR, showing that indeed sequential approximations of rules might behave very differently than their optimisation counterparts. All these results have been proven by Aziz et al. [2017].

Strengthening JR into EJR introduces another layer of complexity—quite literally. Indeed, while one can decide in polynomial time whether a committee satisfies JR [Lackner and Skowron, 2023], deciding the same for EJR is NP-hard [Aziz et al., 2017]. Brill and Peters [2023] proposed a strengthening of EJR called EJR+ that, in contrast to the original axiom, is verifiable in polynomial time. This strengthening is quite agreeable, as many natural rules that satisfy EJR also satisfy EJR+. For example, PAV still satisfies EJR+ [Brill and Peters, 2023].

¹⁶A weaker version of EJR, called proportional justified representation, is already enough to characterise PAV [Sánchez-Fernández et al., 2026]. However, we do not consider this axiom in this thesis.

Definition 2.8. Committee W satisfies **EJR+** if there is no candidate $c \in C \setminus W$, group $N' \subseteq N$, and $\ell \in \mathbb{N}$ with $|N'| \geq \ell \cdot n/k$ such that $c \in \bigcap_{i \in N'} A(i)$ and, for all $i \in N'$, $|A(i) \cap W| < \ell$.

Essentially, to verify that a committee satisfies **EJR+**, we check whether there exists a group of voters that is large enough to deserve ℓ seats, unanimously approves at least one candidate not in the committee, and yet contains no voter who approves ℓ candidates in the committee. If the answer is no, then the committee satisfies **EJR+**.

We conclude this chapter with some examples for these axioms.

Example 2.6. Let $k = 5$ and consider profile

$$A = (\{a, b, c, d, e\} \times 2, \{f, g, h, i, j\} \times 3).$$

This notation means that there are, respectively, 2 copies of the first ballot and 3 copies of the second ballot. Hence, this is a 5-voter profile.

Here, the first two voters constitute a 2-cohesive group. Indeed, we have $2 \geq \ell \cdot n/k = 2$, and the two voters agree on at least two candidates (five in total). On the other hand, the same group is *not* a 3-cohesive group. The last three voters, on the other hand, are 3-cohesive.

In this input, **AV** fails **EJR**, as it selects $\{f, g, h, i, j\}$, leaving voters 1 and 2 wholly unsatisfied. On the other hand, **PAV** here selects all committees picking two candidates from the first two voters and three from the last three voters, such as $\{a, b, f, g, h\}$, all of which satisfy **EJR**.¹⁷

Lastly, consider profile $A = (\{a, b\}, \{a, b\}, \{a, c\})$ and let $k = 2$. For this input, **CC** selects, among others, the dubious committee $\{b, c\}$ (it, indeed, covers all voters). This committee fails **EJR+**: indeed, let $N' = N$ and set $\ell = 2$. Then, candidate a is approved by all voters in this set, and we have $|N'| \geq \ell \cdot n/k$. However, all voters in N' have a satisfaction lower than 2. Interestingly, this committee satisfies **EJR**: indeed, the group N' is not 2-cohesive (by the definition used in **EJR**) since there are no two candidates that are approved by all voters. Hence, this group does not deserve to control two candidates under **EJR**.

¹⁷These outputs can be verified at: <https://pref.tools/abcvoting/?5&1111100000&1111100000&0000011111&0000011111&0000011111> (accessed: 2026-10-05).

CHAPTER 3

Representation

It should be in miniature, an exact portrait of the people at large. It should think, feel, reason, and act like them.

— John Adams [1776]¹

3.1 Introduction

In this chapter, we will study proportional representation in multi-winner voting. We will provide a novel perspective on the matter. Rather than focusing on avoiding the underrepresentation of groups of voters, as virtually all work in this direction does, we will focus on avoiding their *over*representation. We will be studying this issue in *ABC* voting [Lackner and Skowron, 2023], the standard model for approval-based multi-winner elections.

In particular, we have seen in the previous chapter that *EJR* (and related axioms) were the dominant approach to formalising proportionality in the *ABC* voting model. We have seen that *PAV* is the only rule in the family proposed by Thiele [1895] that satisfies *EJR* [Aziz et al., 2017]. This rule, in turn, generalises D’Hondt’s method for party-list parliamentary elections [Janson, 2018, Brill et al., 2022]. We want to argue that the common view of proportionality expressed in these ideas—avoiding underrepresentation—is but one of many possible views.

¹Not to be confused with John Quincy Adams (1767–1848), the sixth president of the United States, whose apportionment method is studied in this chapter. John Adams (1735–1826) was the *second* president of the United States.

To explain this, it is useful to start from party-list parliamentary elections, which are a special case of ABC elections, where the problem appears in its purest form. The issue can be thus summarised: in the abstract, proportionality prescribes that the fraction of votes a party receives should match the fraction of seats it controls. However, in practice, a party might deserve fractionally many seats—we thus need to approximate “ideal proportionality” by rounding these seat counts. Different rounding strategies will embody different views on proportionality by determining which aspects of the “ideal” notion should be prioritised in its approximation.

To make things concrete, consider the fictional country of PUNYLAND. This country has only 100 citizens, and, accordingly, a small parliament with only four seats—some sort of democratic quadrumvirate. Four parties are running in this election, aptly named LARGE, MEDIUM, TINY, and TINIER, which receive, respectively, 50, 33, 9, and 8 votes. The ideal notion of proportionality would say that, then, the LARGE party should receive 50% of the seats, the MEDIUM party 33%, and so on. These quantities—computed as the fraction of votes received times the number of available seats—are known as the “quotas” of the parties. One can visualise this in the following table:

Party	Votes	Share of Votes	Quota
LARGE	50	50%	2.00
MEDIUM	33	33%	1.32
TINY	9	9%	0.36
TINIER	8	8%	0.32
Total	100	100%	4

Here, some parties deserve a fractional amount of seats, and thus we must round these fractions and deviate from “ideal” proportionality. We want a general way to do so—what is known as an *apportionment method*. Clearly, ordinary mathematical rounding (i.e., rounding up only for quotas with a fractional part larger than or equal to 0.5) does not work here—it would give a total of 3 seats, which is less than what we want. So, how to proceed?²

One first, and arguably natural, idea is the following. We give to each party the integer part of its quota (for example, a party having a quota of 3.62 has an integer part of 3). Then, we distribute the remaining seats giving priority to the parties with the largest fractional parts (for example, a party with a quota of 20.42 has a fractional part of 0.42). Under this method, the apportionment of seats would be, in order of party size, 2, 1, 1, and 0. This method is known as Hamilton’s method (or the Largest Remainders method). It has the nice property of always being “within quota”.³ That is, if a party has a quota of 37.71, it will either receive 37 or 38 seats—not fewer, not more.

²The interested reader can verify the results of the apportionment methods in the following discussion via Pref.Tools at <https://pref.tools/apportionment/> (accessed: 2026-05-17). The interface is similar to the case of the ABC voting interface.

³Pukelsheim [2017] uses different terminology, and refers to this as staying within the “ideal frame”.

Nonetheless, Hamilton’s method has procedural drawbacks. In particular, it fails a property known as house monotonicity [Balinski and Young, 2001]: increasing the size of the parliament might lead to a party receiving *fewer* seats under the same votes. For instance, if PUNYLAND suddenly decided to add a fifth seat, thereby turning its ruling council into a democratic quinquvirate, then the Hamilton apportionment would be 3, 2, 0, and 0. The TINY party lost one seat! This problem is known as the *Alabama Paradox*. As the name suggests, this actually occurred in history. Balinski and Young [2001] recount this story in the context of the apportionment of seats in the United States House of Representatives. They mention how in 1881, chief clerk of the Census Office C. W. Seaton—an appropriate name in a discussion about seat apportionment—observed in a letter to the Congress of the United States that “this is to me conclusive proof that the process employed in obtaining it is defective, and that it does not in fact apportion Representatives among the States according to their respective numbers” and that “an entirely different process should be employed”.

D’Hondt’s method, which we mentioned in Section 2.1 and will formally introduce later, sidesteps the Alabama paradox. This is one of its virtues, and an immediate consequence of the fact that seats are allocated sequentially—adding an extra seat does not change the previous computations. In this example, D’Hondt’s method would allocate 3, 1, 0, and 0 seats to the four parties (respectively). We can see that this is not within quota: the LARGE party is receiving 3 seats, whilst its quota is exactly 2. In what sense is this proportional, then? Precisely, because no party is *underrepresented*. In particular, every party receives at least as many seats as the integral part of their quota. This is a general property of D’Hondt’s method [Balinski and Young, 2001]: every party always receives their *lower quota*—what we have so far called the “integral part”, defined as $\lfloor q \rfloor$ where q is the quota.

This property is, in some sense, the defining property of D’Hondt’s method—its view on proportionality. This carries over to the ABC voting model: PAV and EJR embody the same idea. Looking back at the definition of EJR (Section 2.2.8), we can see how we defined this axiom in terms of groups of voters deserving “at least this many candidates”, and how the outcome should give them at least that—they should not be underrepresented. PAV satisfies EJR, and thus guarantees this form of protection against underrepresentation.

Still, this is not the only way to avoid the Alabama paradox. Other “paradox-free” methods might have different, legitimate views on proportionality. For example, in Section 2.1, we have introduced Sainte-Laguë’s and Adams’ methods. How would, say, Adams’ method behave in our toy example? It would give exactly one seat to each party (Adams’ method always gives at least one seat to every party, if possible). Notably, this underrepresents the LARGE party—their lower quota is 2, but they get one seat. On the flip side, this does not *overrepresent* any party: no party receives more than their “upper quota”, defined as $\lceil q \rceil$ for a party with quota q . One can show that Adams’ method *never* overrepresents parties [Balinski and Young, 2001], making it, in some sense, the “dual” of D’Hondt’s method. Note that D’Hondt’s method here overrepresents the LARGE party—its upper quota is 2, but it receives 3 seats.

Sainte-Laguë’s method occupies an intermediate position: conceptually, it aims to stay as near to quotas as possible, balancing under- and overrepresentation, while avoiding procedural issues such as the Alabama paradox. In this example, it would allocate 2, 2, 0, and 0 seats, respectively. Here, it does not under- or overrepresent any party, but in general it can [Balinski and Young, 2001]. The following table summarises the results:

Party	Votes	Quota	Hamilton	D’Hondt	Adams	Sainte-Laguë
LARGE	50	2.00	2	3	1	2
MEDIUM	33	1.32	1	1	1	2
TINY	9	0.36	1	0	1	0
TINIER	8	0.32	0	0	1	0

Adams’ and Sainte-Laguë’s views—preventing overrepresentation and staying close to quota—have received little to no attention in the ABC voting model. In particular, in contrast to underrepresentation and EJR, there is no axiom that extends these ideas to ABC voting. What is more, most well-known proportional rules can overrepresent groups in certain inputs—this is the case for all methods that coincide with D’Hondt’s in party-list parliamentary elections, as the PUNYLAND example indicates. These rules include PAV, as well as Phragmén’s second method (mentioned in the historical background) [Brill et al., 2024], and the method of equal shares, a popular ABC voting rule proposed by Peters and Skowron [2020].

Despite this lack of interest, one might argue that overrepresentation is as serious an issue as underrepresentation. Indeed, some authors have noted the importance of this matter. Lackner and Skowron [2023, Chapter 7] noted how most proportionality notions apply to rules that try to avoid underrepresentation, and listed as an open question the development of proportionality notions that go beyond this, and mentioned how Sainte-Laguë’s method might be superior to D’Hondt’s in certain scenarios. Peters [2025] explicitly identifies the study of overrepresentation and Sainte-Laguë’s ideas within ABC voting as an open question. Going in a slightly technical direction, Cevallos and Stewart [2021] and later Boehmer et al. [2024a] highlight the connection between overrepresentation and security in distributed systems, particularly in the Polkadot network [Burdges et al., 2020], a blockchain with an internal mechanism based on ABC voting.⁴ Beyond approval-based multi-winner elections, Jerrett and Anshelevich [2025] studied the problem of selecting a fixed-size subset of a given set of points embedded in a metric space (e.g., the Euclidean plane). Under this setting, which can be thought to model a type of multi-winner election with cardinal (i.e., numerical) preferences,

⁴However, note that the rules considered by Cevallos and Stewart [2021] and Boehmer et al. [2024a], *maximin-support* and *Phragmén-maximin-support*, do not prevent overrepresentation even when applied to party-list parliamentary elections. Indeed, maximin-support was originally introduced by Sánchez-Fernández et al. [2024] as a rule that generalises D’Hondt’s method of apportionment, which guarantees the lower quota but not the upper quota.

they introduce and study a notion for preventing overrepresentation.⁵ Beyond these exceptions, little work has been conducted in this direction.

In this chapter, we aim to fill this gap. First, we will introduce an axiom that formalises what it means to avoid overrepresentation: [Justified Upper Quota \(JUQ\)](#). We show that this axiom pinpoints a unique rule among the class of composite Thiele rules, a generalisation of Thiele rules which we introduce. This rule, which we call [Adams Approval Voting \(Adams-AV\)](#), is the natural Thiele counterpart of Adams' apportionment method. This mirrors how [EJR](#), an underrepresentation-avoidance axiom, pinpoints [PAV](#), which extends D'Hondt's method. To the best of our knowledge, although [Lackner and Skowron \[2021\]](#) briefly mentioned [Adams-AV](#), we are the first to systematically study this rule. We obtain similar results for Sainte-Laguë's method. We consider the near quota axiom from the apportionment literature, which formalises what it means for an apportionment method to balance under- and overrepresentation, and which Sainte-Laguë's method satisfies [[Balinski and Young, 2001](#)]. We extend it to the [ABC](#) voting setting via an axiom we call [Justified Near Quota \(JNQ\)](#), and we show that [Sainte-Laguë Approval Voting \(SL\)](#),⁶ a Thiele method based on Sainte-Laguë's method of apportionment, satisfies it. Moreover, it is the only composite Thiele rule with this property. Finally, we investigate the compatibility of [JUQ](#), [JNQ](#), and [EJR+](#), that is, whether one can satisfy them simultaneously. We find that, if we allow committees to be smaller than the target size, this is possible. For exhaustive committees, this remains open; we show that some natural avenues for proving compatibility do not work.

3.2 Preliminaries

We have already introduced [ABC](#) voting, as well as some rules and proportionality axioms, in Section 2.2. In this section, we will introduce additional notation to speak about quotas in the context of [ABC](#) voting. We will also discuss a special case of this model: apportionment.

We warn the reader of the following. The notion of “quota”, as described in the introduction, naturally arises in apportionment, and is easiest understood there. It would be thus most natural to start from apportionment and then generalise it to [ABC](#) voting. However, apportionment is a special case of [ABC](#) voting, and hence it is more parsimonious from a modelling perspective to define quotas in the context of [ABC](#) voting, since all the definitions carry over to its special cases. We will thus follow this somewhat unintuitive ordering.

Furthermore, we will use the term “upper quota axioms” to mean axioms for preventing overrepresentation, and similarly for lower quota and underrepresentation.

⁵Interestingly, the work of [Jerrett and Anshelevich \[2025\]](#) appeared almost simultaneously with the original version of the paper on which this chapter is based [[Lackner and Nardi, 2025](#)], in October 2025.

⁶This rule is also known as SLAV [[Lackner and Skowron, 2023](#)], and it can be found on Pref.Tools under the latter name. That was also the name used in the paper on which this chapter is based [[Baychkov et al., 2026](#)].

3.2.1 Quotas in Approval-Based Committee Voting

For a set of voters $N' \subseteq N$, we define its quota as

$$q(N') = \frac{k \cdot |N'|}{|N|}.$$

The quota of a group is the (possibly fractional) number of seats that represents the same proportion of the committee as the group represents of the electorate. Intuitively, under a proportional view, N' would be entitled to fill this many seats if its members can agree on sufficiently many candidates to do so. We call $\lceil q(N') \rceil$ the upper quota and $\lfloor q(N') \rfloor$ the lower quota of N' . We finally define $q_c = q(N(c))$ for a candidate $c \in C$.

In the next subsection, we present an example illustrating quotas (Example 3.2). For now, recall EJR^+ , introduced in Section 2.2.8. Observe that, differently from how we expressed it in Definition 2.8 (which is how Brill and Peters [2023] originally introduced the axiom), EJR^+ can be equivalently expressed in terms of lower quotas:

Definition 3.1. A committee $W \subseteq C$ satisfies EJR^+ if there is no candidate $c \in C \setminus W$ and nonempty group of voters $N' \subseteq N(c)$ such that

$$|A(i) \cap W| < \lfloor q(N') \rfloor \text{ for all } i \in N'.$$

That is, there should be no group N' of voters approving a common unselected candidate c such that everyone in this group approves fewer committee members than the group's lower quota. Otherwise, this group would deserve c to be included in the committee. One can verify that Definitions 2.8 and 3.1 are indeed equivalent by setting ℓ from Definition 2.8 to $\lfloor q(N') \rfloor$.

3.2.2 Apportionment

The apportionment setting (that is, party-list parliamentary elections) is a special case of the ABC voting model. Intuitively, here candidates are partitioned into parties and each voter approves all candidates of exactly one party; in this sense, voters are partitioned into parties as well. For this reason, we will typically think of a party as the voters that support a given set of candidates, rather than the set of candidates itself. Formally:

Definition 3.2. A profile A is an *apportionment profile* if for every $i, j \in N$ we have either $A(i) = A(j)$ or $A(i) \cap A(j) = \emptyset$. Furthermore, for every $i \in N$ we must have $|A(i)| = k$, i.e., we assume that each party has enough candidates to fill the parliament. An *apportionment method* is an exhaustive voting rule whose domain is restricted to apportionment profiles. We say that a voting rule F *extends* an apportionment method G if $F(A) = G(A)$ for all apportionment profiles A .

We define additional notation for apportionment. Let p be the number of unique ballots featured in A , i.e.,

$$p = |\{A(i) : i \in N\}|.$$

We can partition N into p equivalence classes $N_1, \dots, N_p$ such that, for every $\ell \in [p]$, we have that $i, j \in N_\ell$ if and only if $A(i) = A(j)$. Each N_ℓ is called a *party*, and p is the number of parties. We indicate the size of a party $i \in [p]$ as $v_\ell = |N_\ell|$. Note that $v_\ell > 0$ must hold for every $\ell \in [p]$.

With an abuse of notation, we mostly write apportionment profiles as

$$A = (v_1, \dots, v_p),$$

assuming some ordering of the parties. Note that, given the structure of apportionment profiles, this is all the information we need to compute the results, assuming that all parties are to be treated equally (and all the methods we consider do so). Given a party $\ell \in [p]$, we write $q(v_\ell)$ to indicate its quota, i.e., $q(v_\ell) = q(N_\ell)$. We will typically index parties with variable $\ell \in [p]$, and sometimes with variables $i, j \in [p]$.

Similarly, we can consider every pair of committees W and W' with

$$|A(i) \cap W| = |A(i) \cap W'| \quad \text{for every } i \in N$$

as equivalent, assuming—as we do—that we shall judge the desirability of any $W \subseteq C$ only by the vector of voter satisfaction $(|A(1) \cap W|, \dots, |A(n) \cap W|)$ it induces. Hence, we will indicate the result W of an apportionment method via a *seat allocation*

$$a(W) = (a_1, \dots, a_p)$$

such that, for every party $\ell \in [p]$, we have $a_\ell = |A(i) \cap W|$ for an arbitrary $i \in N_\ell$. Typically, we will abstract away from the underlying committees returned by F and just write its outputs as a nonempty set of seats allocations $\{a, \dots\}$. We note that an apportionment method might also return multiple tied seat allocations.

Example 3.1. Consider the 8-voter preference profile

$$A = (\{c_1, c_2, c_3, c_4\} \times 4, \{c_5, c_6, c_7, c_8\} \times 2, \\ \{c_9, c_{10}, c_{11}, c_{12}\}, \{c_{13}, c_{14}, c_{15}, c_{16}\})$$

and let $k = 4$. This is an apportionment profile with $p = 4$ parties. It can be represented as

$$(v_1, v_2, v_3, v_4) = (4, 2, 1, 1).$$

Here, **PAV** would select all committees W containing three candidates from the set of candidates representing the first party (i.e., $\{c_1, c_2, c_3, c_4\}$) and one candidate from the set representing the second party (i.e., $\{c_5, c_6, c_7, c_8\}$). In the apportionment

language, **PAV** selects only one seat allocation, namely

$$\mathbf{a} = (3, 1, 0, 0).$$

This is the same result as D’Hondt’s (as **PAV** extends D’Hondt’s method [Brill et al., 2018]).⁷

We now introduce notions of under- and overrepresentation avoidance in the context of apportionment. We will consider one additional axiom—near quota—that aims to balance under- and overrepresentation. Intuitively, near quota requires that no seat can be reallocated between two parties such that both parties become better represented.

Definition 3.3. An apportionment method satisfies *Apportionment Upper Quota (AUQ)* (respectively, *Apportionment Lower Quota (ALQ)*) if, for every $\mathbf{a}$ selected by this rule and party $\ell \in [p]$, we have that $a_\ell \leq \lceil q(v_\ell) \rceil$ (respectively, $a_\ell \geq \lfloor q(v_\ell) \rfloor$). Moreover, an apportionment method satisfies *Apportionment Near Quota (ANQ)* if for every $\mathbf{a}$ selected by this rule there exist no two parties $i, j \in [p]$ such that

$$|q(v_i) - (a_i - 1)| < |q(v_i) - a_i| \text{ and } |q(v_j) - (a_j + 1)| < |q(v_j) - a_j|.$$

In other words, **ANQ** requires that no seat can be reallocated from party i to party j whilst simultaneously reducing both parties’ distances from their respective quotas.

These definitions also apply to **ABC** voting rules: a rule satisfies these axioms if its restriction to apportionment profiles does. We refer to the book of Balinski and Young [2001] for an extended discussion of lower, upper, and near quota in apportionment.

It is easy to see that, for apportionment profiles, **EJR+** and **ALQ** coincide; we will prove this claim after the following illustrative example.

Example 3.2. To illustrate the three different quota concepts, fix $k = 11$ and consider the following table, which shows an apportionment profile with $n = 110$ voters and $p = 4$ parties along with four possible allocations $\mathbf{a}^{(1)}$, $\mathbf{a}^{(2)}$, $\mathbf{a}^{(3)}$, and $\mathbf{a}^{(4)}$:

⁷The result for **PAV** can be checked at <https://pref.tools/abcvoting/?4&1111000000000000&1111000000000000&1111000000000000&1111000000000000&0000111100000000&0000111100000000&0000000011110000&0000000000001111> (accessed: 2026-05-11). To do so, the user has to click on the rule name (**PAV**) and then on “compute all (may be very slow)”. Recall that on the website candidates are zero-indexed, i.e., candidate 0 corresponds to our candidate c_1 , and so on.

ℓ	1	2	3	4
v_ℓ	24	36	35	15
$q(v_\ell)$	2.4	3.6	3.5	1.5
$\lfloor q(v_\ell) \rfloor$	2	3	3	1
$\lceil q(v_\ell) \rceil$	3	4	4	2
$a_\ell^{(1)}$	3	3	3	2
$a_\ell^{(2)}$	1	4	4	2
$a_\ell^{(3)}$	2	5	3	1
$a_\ell^{(4)}$	2	4	4	1

First, allocation $a^{(1)}$ satisfies **ALQ** and **AUQ**. It does not, however, satisfy **ANQ**, as reallocating one seat from party 1 to party 2 would decrease their respective distances to their quotas by 0.2 each. Allocation $a^{(2)}$ satisfies **AUQ** and **ANQ**. It does not, however, satisfy **ALQ** due to party 1. Allocation $a^{(3)}$ satisfies **ALQ** and **ANQ**, but not **AUQ**, due to party 2. Finally, allocation $a^{(4)}$ satisfies all three notions.

We now can make the following observation.

Observation 3.1. For apportionment profiles, **EJR+** is equivalent to **ALQ**.

Proof. Let A be an apportionment profile with parties $(v_1, \dots, v_p)$ and let W be a committee inducing allocation $a(W) = (a_1, \dots, a_p)$. First, assume that W violates **ALQ**. Then there exists a party $\ell \in [p]$ such that $a_\ell < \lfloor q(v_\ell) \rfloor$. Let i be a voter belonging to party ℓ . Since $|A(i)| = k$ and $a_\ell < k$, there exists a candidate $c \in A(i) \setminus W$ that is approved by all the voters in party ℓ . Since $|A(i) \cap W| = a_\ell < \lfloor q_c \rfloor$ for all such voters i , this also witnesses a violation of **EJR+**.

Next, assume that W violates **EJR+**. Then there exist $c \in C \setminus W$ and a nonempty set $N' \subseteq N(c)$ such that $|A(i) \cap W| < \lfloor q(N') \rfloor$ for all $i \in N'$. Let ℓ be the party of voters approving c . Then $a_\ell < \lfloor q(N') \rfloor \leq \lfloor q(v_\ell) \rfloor$. Hence, W violates **ALQ**. $\square$

As a final ingredient of the apportionment setting, we introduce divisor methods [Balinski and Young, 2001]. Divisor methods are a normatively appealing class of apportionment methods, which includes D'Hondt's, Sainte-Laguë's, and Adams' methods. In particular, divisor methods are the only apportionment methods that are population monotone. Population monotonicity captures the idea that if v_i increases while v_j decreases, then a_i should not decrease while a_j increases. (See the book of Balinski and Young [2001] for a formal definition and a proof of this claim.) Additionally, they all avoid the Alabama paradox, by virtue of being sequential methods.

Definition 3.4. A *divisor method* is defined by a function $d: \mathbb{N}_0 \rightarrow \mathbb{R}_{\geq 0}$ with $d(\ell) \leq d(\ell + 1)$ for every $\ell \in \mathbb{N}_0$. We sometimes write d as a vector. This method proceeds sequentially, starting with $\mathbf{a} = \mathbf{0}$ and iteratively allocating the next seat to a party $i \in [p]$ maximising $v_i/d(a_i)$ (possibly by breaking ties). This is repeated until $\sum_{i \in [p]} a_i = k$. The divisor method defined by d returns all allocations $\mathbf{a}$ that can be obtained by the above procedure plus *some* way of breaking ties.⁸

We adopt the convention that $i/0 > j$ for all $i, j \in \mathbb{N}$. Moreover, $i/0 > j/0$ whenever $i > j$. We can now formally define our three methods of interest:

- D'Hondt's method is given by $d_H = (1, 2, 3, \dots)$;
- Adams' method is given by $d_A = (0, 1, 2, \dots)$;
- Sainte-Laguë's method is given by $d_S = (1, 3, 5, \dots)$.

D'Hondt's and Sainte-Laguë's method are also known as Jefferson's and Webster's methods [Balinski and Young, 2001].

These three methods are notable as they are the only divisor methods that satisfy, respectively, lower, upper, and near quota in apportionment [Balinski and Young, 2001]. Note that, since $d_A(0) = 0$ and given our assumptions about divisions by zero, Adams' method starts by assigning one seat to each party ordered by their size (until all parties are assigned at least one seat or the committee size k is reached) and then proceeds similarly to D'Hondt's method.

Example 3.3. A convenient way to compute the results of a divisor method is as follows. Consider an apportionment profile with four parties and party sizes (22, 10, 5, 3). Let $k = 6$. We can visualise the computation of Sainte-Laguë and Adams' methods in the following tables, respectively:

v_i	1	3	5	7	9	11
22	<u>22</u>	<u>7.33</u>	<u>4.40</u>	3.14	2.44	2
10	<u>10</u>	<u>3.33</u>	2	1.43	1.11	0.91
5	<u>5</u>	1.67	1	0.71	0.56	0.45
3	3	1	0.60	0.43	0.33	0.27

v_i	0	1	2	3	4	5
22	<u>∞</u>	<u>22</u>	<u>11</u>	7.33	5.50	4.40
10	<u>∞</u>	10	5	3.33	2.50	2
5	<u>∞</u>	5	2.50	1.67	1.25	1
3	<u>∞</u>	3	1.50	1	0.75	0.60

In the first column, we have the party sizes. On the following column headers, we have the divisor sequences that characterise the two methods. Excluding the first column, we have in each cell the size of a party divided by a certain weight. For example, the second cell of the first row corresponds to the weight given to the party with 22 members before having received any seats, while the third cell of the same row corresponds to its weight after having received *one* seat. For Adams' method,

⁸Note that Balinski and Young [2001] define divisor methods in a slightly different but equivalent way (see Proposition 3.3 of their book).

we indicate division by zero as ∞ .

Each party starts with zero seats. At each step, the party with the highest weight wins the seat. Whenever we assign a new seat to a party, we decrease its weight according to the table. Underlined entries indicate the maximum weight at each step. The results are, respectively, (3, 2, 1, 0) and (3, 1, 1, 1).⁹

3.3 Composite Thiele Rules

There is a close connection between divisor methods and Thiele rules: the divisor method given by d is extended uniquely among Thiele rules by the one with $\omega(\ell) = d(0)/d(\ell-1)$ [Brill et al., 2018, Lackner and Skowron, 2021]. Note that this is not well-defined for divisor methods with $d(0) = 0$, such as Adams' method.

Among Thiele rules, PAV uniquely extends D'Hondt's method. It is also the only Thiele rule that satisfies EJR+ [Brill and Peters, 2023]. Since we are interested in upper and near quota notions, it is then natural to look at the Thiele rules that extend Adams' and Sainte-Laguë's methods. For Sainte-Laguë, the corresponding Thiele rule is:

Definition 3.5. *Sainte-Laguë Approval Voting (SL)* is the Thiele rule given by $\omega = (1, 1/3, 1/5, \dots)$.

By contrast, Adams' method does not have a corresponding Thiele method, as mentioned above. Its behaviour, however, can be captured in the class of *composite* Thiele rules, which we now introduce.

Intuitively, composite Thiele rules are defined through a sequence of Thiele rules and work in stages. First, we compute the result of the first Thiele rule. Then, at each subsequent stage, we apply the next rule to break the ties of the previous stage. This is analogous to the class of composite social choice scoring functions in single-winner voting with ranked preferences, a class famously studied by Young [1975]. To the best of our knowledge, composite Thiele rules have not been considered before in the literature.

Definition 3.6. Recall the definition of weight functions and Thiele rules (Definitions 2.1 and 2.2). Recall that the Thiele rule THL_ω picks all committees W maximising

$$\text{score}_\omega(A, W) = \sum_{i \in N} \sum_{\ell=1}^{|A(i) \cap W|} \omega(\ell).$$

⁹Recall that one can verify these results at <https://pref.tools/apportionment/> (accessed: 2026-05-17).

Given two weight functions ω_1 and ω_2 , we define the composition

$$(\text{THL}_{\omega_2} \circ \text{THL}_{\omega_1})(A) = \arg \max_{W \in \text{THL}_{\omega_1}(A, k)} \text{score}_{\omega_2}(A, W).$$

For any sequence of weight functions $\omega = (\omega_1, \omega_2, \dots)$ of length $\ell \in \mathbb{N} \cup \{\infty\}$, the corresponding ℓ -composite Thiele rule is

$$\text{THL}_{\omega}(A) = \bigcap_{i \in [\ell]} (\text{THL}_{\omega_i} \circ \dots \circ \text{THL}_{\omega_1})(A).$$

Note that if $\omega = (\omega_1, \dots, \omega_{\ell})$ is finite then the above collapses to

$$\text{THL}_{\omega}(A) = (\text{THL}_{\omega_{\ell}} \circ \dots \circ \text{THL}_{\omega_1})(A)$$

and that this is well-defined even for sequences of length $\ell = \infty$ as there are only finitely many committees.

Note that 1-composite Thiele rules correspond to the usual class of Thiele rules as described in Definition 2.2. We refer to this subclass as *simple Thiele rules*. Furthermore, it follows immediately from the definition that, given a composite Thiele rule THL_{ω} with $\omega = (\omega_1, \dots, \omega_{\ell})$, every rule defined by $(\omega_1, \dots, \omega_{\ell}, \omega_{\ell+1}, \dots)$ refines THL_{ω} .

We allow for ∞ -composite Thiele rules to also capture rules such as the one given by the infinite sequence

$$\omega = \left((1, 0, 0, \dots), (1, 1, 0, 0, \dots), (1, 1, 1, 0, 0, \dots), \dots \right).$$

This is an analogue to the leximin social welfare function defined by Rawls [1971] (see also the work of Sen [2017]).¹⁰

With this, we can now define [Adams Approval Voting \(Adams-AV\)](#). To the best of our knowledge, this rule has not been studied previously in the literature, though it has been mentioned by Lackner and Skowron [2021] (albeit not within the machinery of composite Thiele rules).

Definition 3.7. *Adams-AV* is the 2-composite Thiele rule given by $\omega_1 = (1, 0, 0, \dots)$ and $\omega_2 = (1, 1, 1/2, 1/3, \dots)$.

Observe that *Adams-AV* first maximises the number of covered voters, that is, voters approving at least one candidate in the committee (just as *CC* does). Then, it proceeds with a *PAV*-like weight function. This resembles exactly the behaviour of Adams' method and indeed, we can verify that *Adams-AV* extends it. Before proving this, we give an example.

¹⁰Observe that, if we assume n to be fixed, we do not technically need infinite sequences to capture this rule.

Example 3.4. Consider profile

$$A = (\{a\}, \{a, b\}, \{b, c\}, \{b, c\}, \{b, c\}, \{b, c\})$$

and let $k = 2$. Here, both **AV** and **PAV** pick only $\{b, c\}$, the two most approved candidates. **SL**, on the other hand, is less majoritarian in this example, selecting both $\{a, b\}$ and $\{b, c\}$. Next, **CC** selects $\{a, b\}$ and $\{a, c\}$, as we must select a to cover the first voter.¹¹ Finally, among the **CC**-optimal committees, **Adams-AV** selects only $\{a, b\}$.

We now prove that **Adams-AV** extends Adams' apportionment method. To do so, we need the following technical result.

Lemma 3.2. On instances with n voters and target committee size k , **Adams-AV** is equivalent to the Thiele rule defined by

$$\omega = \left(1, \frac{1}{nk}, \frac{1}{2nk}, \frac{1}{3nk}, \dots\right).^{12}$$

Proof. Let ω be the weight function defined in the statement and define

$$\text{cov}(A, W) = |\{i \in N : A(i) \cap W \neq \emptyset\}|,$$

the number of voters covered by W in A . Fix n, k and a profile A . Let $\omega_1 = (1, 0, 0, \dots)$ and $\omega_2 = (1, 1, 1/2, 1/3, \dots)$. Observe that $\text{THL}_{(\omega_1, \omega_2)}$ is **Adams-AV**. Assume that $\text{THL}_\omega(A)$ contains W . We show that this implies that W is ω_1 -optimal. Were this not the case, we could have another committee W' that covers strictly more voters. But this would imply

$$\begin{aligned} \text{score}_\omega(A, W') &\geq \text{cov}(A, W') \geq \text{cov}(A, W) + 1 > \\ \text{cov}(A, W) \cdot \left(1 + \frac{k-1}{nk}\right) &\geq \text{cov}(A, W) \cdot \left(1 + \sum_{\ell=2}^k \omega(\ell)\right) \geq \text{score}_\omega(A, W), \end{aligned}$$

where the first strict inequality follows from $\text{cov}(A, W) \leq n$. This contradicts the ω -optimality of W : hence, W must be ω_1 -optimal. Next, observe that, among ω_1 -optimal committees, W must also be ω_2 -optimal. Indeed, ω and ω_2 are identical from the second entry onward up to a rescaling, and thus for committees that cover the same amount of voters ω - and ω_2 -optimality are equivalent.

This shows that $\text{THL}_\omega(A) \subseteq \text{THL}_{(\omega_1, \omega_2)}(A)$. The other direction is analogous. $\square$

¹¹These results can be verified at <https://pref.tools/abcvoting/?2&100&110&011&011&011> (accessed: 2026-09-06). Recall that **SL** is named SLAV in Pref.Tools.

¹²Note that in the standard definition (see our Definition 2.2 or the definition by Lackner and Skowron [2023, Chapter 2]), the weights of Thiele rules may not depend on n or k . Hence, this lemma is *not* showing that **Adams-AV** belongs to the class of simple Thiele rules.

We now show the following.

Proposition 3.3. *Adams-AV* extends Adams' method.

Proof. By Theorem 1 of Brill et al. [2018] and our Lemma 3.2, we know that, when n and k are fixed, *Adams-AV* extends the divisor method F defined by $d = (1, nk, 2nk, 3nk, \dots)$. It remains to be shown that F is equivalent to Adams' method on apportionment elections with n voters and target committee size k .

Recall that Adams' method starts by assigning a seat to every party (in order of party size) and continues until all seats are assigned or until all parties have received at least one seat. Then, consider a partial allocation a obtained during the execution of F . Let $i \in [p]$ be the largest party such that $a_i = 0$. For every party $j \in [p]$ with $a_j = 0$ we have $v_i/d(0) = v_i \geq v_j = v_j/d(0)$. Furthermore, for every party $j \in [p]$ with $a_j > 0$,

$$\frac{v_i}{d(0)} = v_i \geq 1 > \frac{v_j}{nk} \geq \frac{v_j}{d(a_j)}.$$

Thus, as long as there are parties that have received no seats, F will assign the next seat to the largest party among them (and break the ties arbitrarily). This coincides with the behaviour of Adams' method. Finally, once $a_\ell > 0$ for all parties $\ell \in [p]$, it is straightforward to see that F behaves like Adams' method, since their d -sequences are identical from the second entry onward up to rescaling. $\square$

3.4 Avoiding Overrepresentation

In the previous section we introduced *Adams-AV*, an *ABC* voting rule that extends Adams' apportionment method. The latter never overrepresents parties; however, it is not immediately clear whether *Adams-AV* always avoids overrepresentation in *ABC* voting, or even how we should *define* overrepresentation in that setting. In this section we answer both of these questions by generalising the notion of upper quota to groups of voters in *ABC* elections.

3.4.1 Some Difficulties with Upper Quota Axioms

A natural starting point is to take inspiration from appealing lower quota axioms for *ABC* voting. In particular, we begin by defining a natural upper quota analogue of *EJR+*.

We will interpret groups of form $N(c)$ for some $c \in C$ as (possibly intersecting) parties, and we will require that there is no selected candidate whose supporters all receive more than this candidate's upper quota.

Definition 3.8. A set $W \subseteq C$ of candidates satisfies *Upper Quota (UQ)* if there is no selected candidate $c \in W$ with $|A(i) \cap W| > \lceil q_c \rceil$ for all $i \in N(c)$.

Note that, in contrast to [EJR+](#), we only consider the full set of supporters $N(c)$ and not its subsets, as there could exist a small subset $N' \subset N(c)$ that is overrepresented while the entire set $N(c)$ is not. For example if we consider the full set of voters N , this group is never overrepresented (it proportionally deserves the whole committee); however, some of its subsets (in particular, singletons) may appear overrepresented according to the above condition.

Moreover, observe that to be overrepresented *everyone* needs to be over the upper quota, similarly to how [EJR](#)-style axioms require all members of a particular voter set to be below lower quota for underrepresentation to occur.

[UQ](#) is trivially satisfiable by inexhaustive committees, in particular by the empty committee;¹³ however, it cannot always be satisfied for exhaustive committees.

Example 3.5. Consider the profile $A = (\{a, b, c\}, \{d, e, f\}, \{g\})$ and let $k = 6$. This is *almost* an apportionment profile, but here voters approve fewer than k candidates. Any exhaustive committee must contain either three candidates approved by voter 1 or three candidates approved by voter 2 (or both). Suppose, without loss of generality, that $\{a, b, c\} \subseteq W$. Then we have that $|A(1) \cap W| = 3 > \lceil q_c \rceil = 2$, violating [UQ](#). A similar result holds for any other exhaustive committee.

Due to the closeness of the above example to the apportionment setting, it seems hard to argue with the interpretation that voter 1 and 2 each form a party (or group) and that, if we select three candidates they approve, they are overrepresented.

Example 3.5 illustrates that defining upper quota axioms for [ABC](#) voting poses some additional challenges compared to lower quota axioms, even when preferences do not intersect: it is unclear how to determine whether a group is “receiving too much” compared to its size when the only other option is to wastefully select less than k candidates.

3.4.2 Main Upper Quota Axiom: Justified Upper Quota

Example 3.5 shows that in some cases overrepresentation is unavoidable for exhaustive committees. Our solution is to permit overrepresentation only when it is—in some sense—necessary. We interpret this necessity as there being no other group who could receive a candidate without any of the voters breaching their upper quota. We call such violations of upper quota “justified”. We can now introduce our axiom.

Definition 3.9. Committee W satisfies *Justified Upper Quota (JUQ)* if, whenever there is a $c \in W$ such that

$$|A(i) \cap W| > \lceil q_c \rceil \quad \text{for all } i \in N(c),$$

¹³Note that q_c is defined in terms of the target committee size k , not the actual size of W .

then there exists no $d \in C \setminus W$ and nonempty $N' \subseteq N(d)$ such that

$$|A(i) \cap ((W \cup \{d\}) \setminus \{c\})| \leq \lceil q(N') \rceil \quad \text{for all } i \in N'.$$

Note that we allow for arbitrary subsets N' of $N(d)$ in our definition: for us it is sufficient that there exists some nonempty subgroup of the approvers of d who could be awarded an extra candidate without breaching their upper quota. Observe that **UQ** implies **JUQ**, and that we require $N' \neq \emptyset$ as otherwise the second condition would be vacuously true for any unselected candidate.

Even when every exhaustive committee overrepresents some group according to **UQ**, the weaker **JUQ** can remain meaningfully satisfiable.

Example 3.6. Let us revisit the scenario of Example 3.5, where $A = (\{a, b, c\}, \{d, e, f\}, \{g\})$ and $k = 6$. First, consider the committee $W_1 = \{a, b, c, d, e, g\}$. This committee satisfies **JUQ**. Indeed, consider, for example, $c \in W_1$. We have that $\lceil q_c \rceil = 2$, but voter 1 approves 3 candidates on the committee. However, the only candidate not in W_1 is f , and we have that, if we were to replace c with f in W_1 , then voter 2 would get a satisfaction of $3 > 2 = \lceil q_f \rceil$. On the other hand, the committee $W_2 = \{a, b, c, d, e, f\}$ violates **JUQ**. Indeed, again c is a witness for an upper quota violation. However, now we should can swap c with g , as $\lceil q_g \rceil = 2 \geq |A(3) \cap ((W_2 \cup \{g\}) \setminus \{c\})|$.

JUQ only disallows overrepresentation in the most egregious cases: if *all* members in a group $N(c)$ are above their upper quota, and there is another group $N' \subseteq N(d)$ where *all* members would be at most at their upper quota after replacing c with d , then **JUQ** is violated. Note the asymmetry: it would be perhaps more natural to require only *some* members of N' to be at most at their upper quota to warrant the replacement of c by d . However, there are reasons to believe that this “symmetric variant” of **JUQ** is too strong: no composite Thiele satisfies this axiom.

Example 3.7. Consider the 8-voter profile

$$A = (\{a\}, \{a\}, \{a, d\}, \{a, b, c\}, \{a, b, c\}, \{a, b, c\}, \{a, b, c, d\}, \{a, b, c, d\})$$

and let $k = 3$. **Adams-AV** here picks only $\{a, b, c\}$. However, $N(c)$ is overrepresented as all its members have satisfaction $3 > 2 = \lceil q_c \rceil$. If we switch c with d , then not all voters in $N(d)$ would be overrepresented as the voter with ballot $\{a, d\}$ would have satisfaction $2 = \lceil q_d \rceil$. Hence, **Adams-AV** and all of its refinements violate the “symmetric variant” of **JUQ**.

As we shall see (Theorem 3.6), the same holds for every other composite Thiele rule,

motivating our focus on **JUQ**.¹⁴ We will therefore leave a more in-depth analysis of this stronger notion to future work.

Next, similarly to the relationship between **EJR+** and **ALQ**, we can show that **JUQ** and **AUQ** are equivalent on apportionment profiles.

Proposition 3.4. For apportionment profiles, **JUQ** is equivalent to **AUQ**.

Proof. Consider an apportionment profile A with parties $(v_1, \dots, v_p)$ and consider any committee W with $\mathbf{a}(W) = (a_1, \dots, a_p)$.

First, assume that W violates **JUQ**. Then there exists a candidate $c \in W$ (approved by some party $\ell \in [p]$) such that $\lceil q(v_\ell) \rceil = \lceil q_c \rceil < |A(i) \cap W| = a_\ell$ for all voters i in party ℓ and hence W also violates **AUQ**.

For the other direction, we observe that by definition of the quota it holds that

$$\sum_{\ell \in [p]} a_\ell \leq k = \sum_{\ell \in [p]} k \cdot \frac{v_\ell}{n} = \sum_{\ell \in [p]} q(v_\ell).$$

Hence, if there exists a party $i \in [p]$ with $a_i > q(v_i)$ there must exist a second party $j \in [p]$ with $a_j < q(v_j)$ and thus, since $a_j \in \mathbb{N}$, we get $a_j + 1 \leq \lceil q(v_j) \rceil$. Hence, if **AUQ** is violated, we can take any selected candidate approved by the members of party i as c , any unselected candidate approved by the members of party j as d , and the set $N(d)$ as witnesses that **JUQ** is also violated. $\square$

3.4.3 Characterisation of Adams Approval Voting

In the previous section, we introduced **JUQ** as our upper quota axiom. It is easy to see that we can always satisfy **JUQ** via a inexhaustive committee: for instance, the empty committee trivially satisfies it. However, similar to **UQ**, it is not obvious that there always exists a committee of size k satisfying **JUQ**. We show that this is always the case. In particular, we prove that every committee selected by **Adams-AV** satisfies **JUQ**.

Theorem 3.5. Any refinement of **Adams-AV** satisfies **JUQ**.

Proof. Let $\omega_1 = (1, 0, 0, \dots)$ and $\omega_2 = (1, 1, 1/2, 1/3, \dots)$, the weight functions defining **Adams-AV**. Fix a profile $A \in \mathcal{A}$ and a committee $W \in \mathcal{W}_k$ that fails **JUQ**. Let c, d , and $N' \subseteq N(d)$ be the witnesses of the violation of **JUQ** as described in Definition 3.9. We will show that committee $W' = (W \setminus \{c\}) \cup \{d\}$ has either a higher ω_1 -score than W (i.e., it covers more voters) or has the same ω_1 -score but higher ω_2 -score. This is enough to show that any refinement of **Adams-AV** satisfies **JUQ**: were this not the case, we could

¹⁴The same holds for all popular **ABC** voting rules we are aware of, since—as we discuss in the next section—they all fail **JUQ**. However, a rule we introduce later on (Section 3.4.5), the **Upper Quota Elimination Rule (UQER)**, would satisfy the stronger symmetric variant of **JUQ**, as a close inspection of the proof of Theorem 3.10 shows.

improve the [Adams-AV](#) score of the winning committee. In the following we assume $n_c > 0$, as otherwise the proof is trivial.

First, note that $\text{score}_{\omega_1}(A, W') \geq \text{score}_{\omega_1}(A, W)$. This holds because $n_c > 0$ implies $\lceil q_c \rceil \geq 1$ and thus, by definition of [JUQ](#), $|A(i) \cap W| \geq 2$ for every $i \in N(c)$. Hence, removing c cannot make any previously covered voter uncovered. We can then assume that $\text{score}_{\omega_1}(A, W') = \text{score}_{\omega_1}(A, W)$ as otherwise our claim follows.

We prove two useful facts. First, $|N' \setminus N(c)| > 0$. Otherwise, we would have $N' \subseteq N(c)$. This implies $\lceil q_c \rceil \geq \lceil q(N') \rceil$ and that some $i^* \in N' \cap N(c)$ exists. Together with the definition of [JUQ](#), this means that $\lceil q_c \rceil \geq \lceil q(N') \rceil \geq |A(i^*) \cap W'| = |A(i^*) \cap W| > \lceil q_c \rceil$, a contradiction. Hence, $|N' \setminus N(c)| > 0$. Moreover, since $\text{score}_{\omega_1}(A, W') = \text{score}_{\omega_1}(A, W)$, any $i \in N' \setminus N(c)$ is covered and thus, by [JUQ](#), we get $\lceil q(N') \rceil > 1$.

Next, define $\Delta = \text{score}_{\omega_2}(A, W) - \text{score}_{\omega_2}(A, W')$. We will show $\Delta < 0$, which is enough to conclude the proof.

Since we only swap c and d , the satisfaction of every $i \in N$ with $\{c, d\} \subseteq A(i)$ or $\{c, d\} \cap A(i) = \emptyset$ is the same for both W and W' . This, together with the fact that $\text{score}_{\omega_1}(A, W') = \text{score}_{\omega_1}(A, W)$, gives

$$\begin{aligned} \Delta = & \sum_{i \in N(c) \setminus N(d)} \left(\sum_{\ell=2}^{|A(i) \cap W|} \frac{1}{\ell-1} \right) + \sum_{i \in N(d) \setminus N(c)} \left(\sum_{\ell=2}^{|A(i) \cap W|} \frac{1}{\ell-1} \right) \\ & - \sum_{i \in N(c) \setminus N(d)} \left(\sum_{\ell=2}^{|A(i) \cap W'|} \frac{1}{\ell-1} \right) - \sum_{i \in N(d) \setminus N(c)} \left(\sum_{\ell=2}^{|A(i) \cap W'|} \frac{1}{\ell-1} \right). \end{aligned}$$

Given that $|A(i) \cap W'| = |A(i) \cap W| - 1$ for $i \in N(c) \setminus N(d)$ and $|A(i) \cap W'| = |A(i) \cap W| + 1$ for $i \in N(d) \setminus N(c)$, we get

$$\begin{aligned} \Delta = & \sum_{i \in N(c) \setminus N(d)} \frac{1}{|A(i) \cap W| - 1} - \sum_{i \in N(d) \setminus N(c)} \frac{1}{|A(i) \cap W|} \leq \\ & \sum_{i \in N(c) \setminus N'} \frac{1}{|A(i) \cap W| - 1} - \sum_{i \in N' \setminus N(c)} \frac{1}{|A(i) \cap W|}. \end{aligned}$$

By definition of [JUQ](#), any $i \in N(c)$ approves at least $\lceil q_c \rceil + 1$ candidates in W , while any $i \in N' \setminus N(c)$ at most $\lceil q(N') \rceil - 1$. Recall that $N' \setminus N(c)$ is nonempty and that $\lceil q(N') \rceil > 1$. We get

$$\Delta \leq \frac{|N(c) \setminus N'|}{\lceil q_c \rceil} - \frac{|N' \setminus N(c)|}{\lceil q(N') \rceil - 1}.$$

Let $|N'| = n'$ and $|N(c) \cap N'| = n^\dagger$. We can derive

$$\Delta \leq \frac{n_c - n^\dagger}{\lceil q_c \rceil} - \frac{n' - n^\dagger}{\lceil q(N') \rceil - 1} < \frac{n(n_c - n^\dagger)}{k \cdot n_c} - \frac{n(n' - n^\dagger)}{k \cdot n'} = \frac{n}{k} \left(\frac{n^\dagger}{n'} - \frac{n^\dagger}{n_c} \right).$$

If $n^\dagger = 0$ or $n_c \leq n'$, we obtain $\Delta < 0$ which concludes the proof. Hence, assume $n^\dagger > 0$ and $n_c > n'$. The latter two facts imply, respectively, that $\lceil q_c \rceil > \lceil q(N') \rceil$ and that there

is some $i^* \in N(c) \cap N'$. We have already argued above that this implies a contradiction. Thus $\Delta < 0$, completing the proof. $\square$

We complement this by a characterisation of [Adams-AV](#) among Thiele rules: a composite Thiele rule satisfies [AUQ](#) if and only if it is a refinement of [Adams-AV](#).

Theorem 3.6. A composite Thiele rule satisfies [AUQ](#) if and only if it is a refinement of [Adams-AV](#).

Note that the negative part of Theorem 3.6 can in part be derived from the theory of apportionment, at least when it comes to a class of simple Thiele rules with strictly positive weights. Indeed, we know that every such simple Thiele rule extends a unique divisor method [Brill et al., 2018], and that all divisor methods except Adams' violate [AUQ](#) [Balinski and Young, 2001], at least in the class studied by Balinski and Young [2001] which forces $d(i) \in [i, i + 1]$ (with appropriate normalisation) for any divisor method defined by d . However, for non-simple Thiele rules and those not fitting the definition of Balinski and Young, this does not follow. We thus give a direct proof of the result.

Proof of Theorem 3.6. First, observe that any refinement of [Adams-AV](#) satisfies [AUQ](#) (Theorem 3.5 and Proposition 3.4). Fix ω and suppose THL_ω satisfies [AUQ](#). We will show that THL_ω refines [Adams-AV](#). We will work in five steps, using five apportionment profiles.

1. First, we show that $\omega_1 = (1, 0, \dots)$, i.e., THL_ω refines [CC](#).
2. Second, we argue that ω has at least two entries, i.e., $\omega = (w_1, w_2, \dots)$.
3. Third, we prove that $\omega_2(\ell) > 0$ for all $\ell \in \mathbb{N}$.
4. Fourth, we demonstrate that $\omega_2(\ell) \geq 1/(\ell-1)$ for all $\ell \in \mathbb{N} \setminus \{1\}$.
5. Finally, we conclude that $\omega_2 = (1, 1, 1/2, 1/3, \dots)$.

The above is sufficient to show that THL_ω refines [Adams-AV](#). In the following, we will only argue about apportionment profiles. Recall that we denote those via a vector of party sizes $A = (v_1, \dots, v_p)$.

Step 1. Fix $k = 2$ and consider $n \in 2\mathbb{N}$ with $n > 2$. Consider profile

$$A = (\underbrace{1, \dots, 1}_{n/2 \text{ many}}, n/2).$$

Observe that $p = n/2 + 1$. Since $\omega_1(1) = 1 \geq \omega_1(2)$, we cannot have $a_p = 0$, as we could increase $\text{score}_{\omega_1}(A)$ by removing one seat from some party $\ell \in [p - 1]$ with $a_\ell > 0$ and

giving it to party p . Moreover, $\lceil q(v_\ell) \rceil = 1$ for all parties $\ell \in [p]$. By [AUQ](#), we must then have $a_p = 1$. Therefore,

$$\frac{n}{2} \cdot (\omega_1(1) + \omega_1(2)) \leq \left(\frac{n}{2} + 1\right) \cdot \omega_1(1) \implies n/2 \cdot \omega_1(2) \leq \omega_1(1).$$

Since ω_1 does not depend on n (which can be arbitrarily large), we obtain $\omega_1(2) = 0$.

Step 2. Next, assume that there are three voters divided in two parties with sizes $A = (1, 2)$. Let $k = 3$. If $\omega = (\omega_1)$, that is, ω has only one entry, then some $\mathbf{a}$ with $\mathbf{a} = (2, 1)$ is among the winning allocations, which violates [AUQ](#). Thus, ω has at least two entries. Note that we can assume $\omega_1 \neq \omega_2$ without loss of generality, which implies $\omega_2(2) > 0$.

Step 3. Now, assume that $\omega_2(\ell) = 0$ for some $\ell \in \mathbb{N}$. Pick ℓ^* to be the smallest such ℓ , that is, $\ell^* = \min \{\ell \in \mathbb{N} : \omega_2(\ell) = 0\}$. Observe that $\ell^* > 2$. Let $k = 2\ell^* - 2$. Construct a two-party profile $A = (\ell^*, \ell^* - 2)$ with $n = 2\ell^* - 2$ voters. Note that $\lceil q_1 \rceil = \ell^*$ and $\lceil q_2 \rceil = \ell^* - 2$. However, since $\omega_1(2) = \omega_2(\ell^*) = 0 < \omega_2(\ell^* - 1)$, any winning allocation must have $a_1 = a_2 = \ell^* - 1$, violating [AUQ](#). Therefore, $\omega_2(\ell) > 0$ for every $\ell \in \mathbb{N}$.

Step 4. Fix $\ell \in \mathbb{N} \setminus \{1, 2\}$. Pick an arbitrarily small $\epsilon \in \mathbb{Q}$ where $0 < \epsilon < 1$ and some $h \in \mathbb{N}$ such that $h \cdot \epsilon \in \mathbb{N}$. Let $g = \lceil 1/\epsilon \rceil$ and observe that $1 \leq g\epsilon$. Let $k = g(\ell - 1) + 2$ and $n = h + gh(\ell - 1 + \epsilon)$. Consider profile

$$A = (h, \underbrace{h(\ell - 1 + \epsilon), \dots, h(\ell - 1 + \epsilon)}_{g \text{ many}}).$$

Observe that there are $p = g + 1$ parties. Then, we have

$$q(v_1) = \frac{h}{h + gh(\ell - 1 + \epsilon)} [g(\ell - 1) + 2] \leq 1 \iff 1 \leq g\epsilon$$

which holds by definition of g . Similarly, for party $i \in [g + 1] \setminus \{1\}$,

$$q(v_i) = \frac{h(\ell - 1 + \epsilon)}{h + gh(\ell - 1 + \epsilon)} [g(\ell - 1) + 2] > \ell - 1 \iff \ell - 1 + 2\epsilon > 0,$$

which holds because $\ell > 2$ and $\epsilon < 1$. Since any winning $\mathbf{a}$ must be ω_1 -optimal, given $k \geq g + 1$, we must have $a_i > 0$ for all $i \in [g + 1]$. Therefore, $a_1 = 1$, as otherwise we would violate [AUQ](#). Therefore, allocation

$$\mathbf{a} = (1, \ell, \ell - 1, \dots, \ell - 1)$$

must have optimal ω_2 -score. This holds because every two parties $i, j \in [g + 1] \setminus \{1\}$ have $v_i = v_j$, and hence we can achieve maximal ω_2 -score by spreading the seats over

the parties in $[g+1] \setminus \{1\}$ as evenly as possible. Consequently, $\mathbf{a}$ must have strictly lower ω_2 -score than allocation

$$\mathbf{a}' = (2, \ell-1, \dots, \ell-1),$$

which violates AUQ. This yields

$$\begin{aligned} h(\omega_2(1) + \omega_2(2)) + gh(\ell-1+\epsilon) \sum_{s=1}^{\ell-1} \omega_2(s) < \\ h \cdot \omega_2(1) + h(\ell-1+\epsilon) \cdot \omega_2(\ell) + gh(\ell-1+\epsilon) \sum_{s=1}^{\ell-1} \omega_2(s), \end{aligned}$$

which in turn implies

$$\frac{\omega_2(2)}{\ell-1+\epsilon} < \omega_2(\ell).$$

Since $\omega_2(2) = 1$ and this holds for any $\ell > 2$ and arbitrarily small ϵ we get that $\omega_2(\ell) \geq 1/(\ell-1)$ for all $\ell > 1$.

Step 5. Finally, we will show that $\omega_2(\ell) = 1/(\ell-1)$ for all $\ell \in \mathbb{N} \setminus \{1\}$ by induction on ℓ . The base case ($\ell = 2$) is trivial, as we already know $\omega_2(2) = 1$. Hence, fix $\ell > 2$ and consider the induction hypothesis

$$“\omega_2(s) = 1/(s-1) \text{ for } s \in [\ell-1] \setminus \{1\}.”$$

Pick an arbitrarily small $\epsilon \in \mathbb{Q}$ where $0 < \epsilon < 1$ and some $h \in \mathbb{N}$ such that $h \cdot \epsilon \in \mathbb{N}$. Let

$$g = \left\lceil \frac{\ell-1}{\epsilon} \right\rceil.$$

Fix $k = g + \ell$ and consider the profile with $n = h(g + \ell - 1 - \epsilon)$ voters

$$A = (\underbrace{h, \dots, h}_{g \text{ many}}, \ell-1-\epsilon).$$

Observe that there are $p = g+1$ parties and that, for any $i \in [g]$, we have

$$q(v_i) = \frac{h}{h(g + \ell - 1 - \epsilon)}(g + \ell) > 1 \iff 1 + \epsilon > 0,$$

which holds by definition of ϵ . Next,

$$q(v_{g+1}) = \frac{h(\ell-1-\epsilon)}{h(g + \ell - 1 - \epsilon)}(g + \ell) \leq \ell-1 \iff g+1 \geq \frac{\ell-1}{\epsilon},$$

which again by definition of g . Moreover, observe that since any winning allocation must be ω_1 -optimal and $k > g+1$, every party gets assigned at least one seat. Consequently,

we have $k - (g + 1) = \ell - 1$ seats left to assign. We now prove that $a_{g+1} \geq \ell - 1$. Indeed, consider any allocation where $a_{g+1} = \ell - 1 - z$ for some positive $z \in \mathbb{N}$. Among these allocations, we claim that

$$\mathbf{a} = (\underbrace{2, \dots, 2}_{(z+1) \text{ many}}, \underbrace{1, \dots, 1}_{(g-z-1) \text{ many}}, \ell - 1 - z)$$

is ω_2 -optimal. This holds because, after accounting for the first seat given to the parties in $[g]$ as well as the seats given to party $g + 1$, we are left with

$$k - g - (\ell - 1 - z) = z + 1 \leq \ell - 1 \leq \frac{\ell - 1}{\epsilon} \leq g$$

seats. Hence, we cannot give at least 2 seats to all parties in $[g]$, and it is therefore optimal to spread the seats among the parties in $[g]$ as evenly as possible, which is what $\mathbf{a}$ does. We now show that $\mathbf{a}$ has lower ω_2 -score than

$$\mathbf{a}' = (2, 1, \dots, 1, \ell - 1),$$

showing that a_{g+1} must be at least $\ell - 1$. We compare the scores of $\mathbf{a}$ and $\mathbf{a}'$:

$$\begin{aligned} gh \cdot \omega_2(1) + h(z + 1)\omega_2(2) + h(\ell - 1 - \epsilon) \sum_{s=1}^{\ell-1-z} \omega_2(s) &< \\ gh \cdot \omega_2(1) + h \cdot \omega_2(2) + h(\ell - 1 - \epsilon) \sum_{s=1}^{\ell-1} \omega_2(s), \end{aligned}$$

which for positive z holds if

$$z \cdot \omega_2(1) = z < \sum_{s=\ell-z}^{\ell-1} (\ell - 1 - \epsilon) \cdot \omega_2(s). \quad (3.1)$$

Now, by the induction hypothesis, the smallest addend involved in the summation on the right-hand side is

$$\frac{\ell - 1 - \epsilon}{\ell - 2} > 1,$$

because $\epsilon < 1$. Therefore

$$z = \sum_{s=\ell-z}^{\ell-1} 1 < \sum_{s=\ell-z}^{\ell-1} (\ell - 1 - \epsilon) \cdot \omega_2(s),$$

which shows that Equation (3.1) holds, and hence, that $a_{g+1} \geq \ell - 1$. Finally, since $a_{g+1} = \ell$ violates [AUQ](#), we must have

$$gh \cdot \omega_2(1) + h \cdot \omega_2(2) + h(\ell - 1 - \epsilon) \sum_{s=1}^{\ell-1} \omega_2(s) > gh \cdot \omega_2(1) + h(\ell - 1 - \epsilon) \sum_{s=1}^{\ell} \omega_2(s),$$

which in turn implies

$$\frac{\omega_2(2)}{\ell - 1 - \epsilon} > \omega_2(\ell).$$

Since $\omega_2(2) = 1$ and this holds for arbitrarily small ϵ , we have that $\omega_2(\ell) \leq 1/(\ell-1)$, and thus, by the previous arguments in this proof, $\omega_2(\ell) = 1/(\ell-1)$. This closes the induction step, and we are done. $\square$

In light of Theorem 3.6 one might wonder whether Thiele rules, which maximise satisfaction, are just ill-suited to avoiding overrepresentation or whether JUQ is just hard to satisfy. The latter seems to be the case, as, to the best of our knowledge, no other known ABC voting rule satisfies JUQ. Indeed, in Appendix A (Section A.2.1), we show counterexamples for a wide set of familiar rules. This includes, for example, Phragmén’s rules [Brill et al., 2024], the method of equal shares [Peters and Skowron, 2020], as well as all rules in the ABCVOTING Python package [Lackner et al., 2023b] at the time of writing. Interestingly, even the sequential formulation of Adams-AV fails JUQ.¹⁵

This concludes our initial exploration of JUQ, which so far has focused on satisfiability and on which rules are guaranteed to satisfy the axiom. Next, we turn to a structural aspect of JUQ: its swap dynamics. Indeed, the axiom is naturally stated in terms of pairs of candidates, one “worse” and one “better”. One can therefore think of JUQ as prescribing a swap whenever such a swap is available. In the next section, we investigate the effect of these swaps, and the evolution of committees under repeated swaps. This will help us gain a deeper understanding of the axiom and, in the process, identify some of its potential limitations and our responses to them.

3.4.4 Swap Dynamics of Justified Upper Quota

As mentioned above, we can think of JUQ as being defined in terms of *swaps* of candidates. Formally:

Definition 3.10. A JUQ swap for a committee W is a distinct pair of candidates $c \in W$ and $d \notin W$ such that c and d witness the JUQ violation of W together with some $N' \subseteq N(d)$. A sequence of JUQ swaps starting from W is a sequence of form

$$(c_1, d_1), (c_2, d_2), \dots$$

where (c_1, d_1) is a JUQ swap for W , (c_2, d_2) is a JUQ swap for $(W \setminus \{c_1\}) \cup \{d_1\}$, and so on.

A natural question is whether swap sequences always terminate, i.e., whether all sequences must eventually reach a committee that satisfies JUQ. If this were not the case, we could have infinite sequences of swaps. This would challenge the intuition that JUQ

¹⁵We have not defined sequential composite Thiele rules. However, one can naturally extend Definition 2.4 to composite rules.

swaps are “improvements”: in our finite setting, any infinite sequence must eventually revisit a committee and thus contain a cycle. If the improvement relation is transitive, this would imply that some committee is an improvement of itself.

One first approach to show that such sequences converge would be to prove that, for any starting committee W , there is no sequence of JUQ swaps of form

$$(c_1, d_1), (c_2, d_2), \dots, (c_{\ell-1}, d_{\ell-1}), (c_{\ell}, c_1),$$

i.e., that it is never possible to add a candidate c_1 back after removing it through a JUQ swap. Notice that this would also immediately give us a polynomial-time algorithm for computing an exhaustive JUQ committee: we could start from an arbitrary exhaustive committee W , perform at most $m - k$ swaps, and find a committee that satisfies JUQ. Unfortunately, the following shows that this approach does not work in general.

Example 3.8. Consider profile

$$A = (\{a, b, c\}, \{a, b, d\}, \{e, f\}, \{e, f\})$$

and let $k = 4$. If we start with the committee $W = \{a, b, c, d\}$, the sequence

$$(a, e), (c, f), (d, a)$$

shows that we can add candidate a back in.

The above problem persists even assuming some clever schemata on how to pick which JUQ swap to perform next.

Example 3.9. Consider the following schemata: “always choose a swap (c, d) such that c has the least supporters”, “always choose a swap (c, d) such that d has the least supporters”, and “always choose a swap (c, d) such that d has the most supporters”. The following counterexample works for any of these schemata, and also for their combination when one criterion is applied first and the other is used to break ties. Consider the following election where $k = n = 15$:

Candidate/Voter	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
c_1, c_2, c_3	✓	✓	✓												
c_4, c_5				✓	✓	✓									
c_6	✓			✓											
c_7		✓			✓										
c_8			✓			✓									
c_9, c_{10}, c_{11}							✓	✓	✓						
$c_{12}, \dots, c_{19}$										✓	✓	✓	✓	✓	✓

In the above table, ticks indicate approvals. For instance, $A(1) = \{c_1, c_2, c_3, c_6\}$ and

$$A(10) = \dots = A(15) = \{c_{12}, \dots, c_{19}\}.$$

Suppose we start with the following size k committee:

$$\{c_1, c_2, c_3, c_4, c_6, c_7, c_8, c_{12}, \dots, c_{19}\}.$$

Notably, c_5 is not in this committee, while candidates in $\{c_1, c_2, c_3, c_{12}, \dots, c_{19}\}$ have all their (respective) supporters above upper quota. According to our schemata, we can replace c_3 (a minimally-supported, upper-quota-violating candidate) with, e.g., c_5 . After this change, candidates in $\{c_6, c_7, c_8\}$ have their respective supporters over their upper quotas. We can replace each of these with a candidate from $\{c_9, c_{10}, c_{11}\}$. Note that now we can add c_3 back (e.g., by swapping it with c_{19}).

Let us then consider a different strategy. Another approach would be to use a potential function

$$P: \mathcal{W}_k \rightarrow \mathbb{R}_{\geq 0}$$

over the committees and show that each swap strictly decreases (or increases) this function. Since there are only finitely many committees, this would show that there are no infinite sequences of JUQ swaps. Were this not the case, the same committee would eventually have to appear twice along the sequence, contradicting the fact that the potential changes strictly with every swap. This is a familiar approach from game theory; see, e.g., the seminal work of [Monderer and Shapley \[1996\]](#).

Given the scope of this chapter, a natural starting point for such a potential would be to consider some measure of “overrepresentation” of a committee. To name a few natural candidates, one might consider the number of candidates c for which $N(c)$ is overrepresented, the number of unique voters that are overrepresented, or the magnitude of total overrepresentation (i.e., sum of the satisfaction above upper quota).

Unfortunately, this approach also does not work—there are JUQ swaps that *increase* overrepresentation according to these measures.

Example 3.10. Consider profile

$$A = (\{a, b, c\} \times 4, \{d, e, f\} \times 3, \{d, g, h\} \times 3, \{i\} \times 6)$$

and let $k = 8$. Observe that $n = 16$ and $n/k = 2$. Consider committee

$$W = \{a, b, c, e, f, g, h, i\}.$$

This committee violates JUQ, and (c, d) is a JUQ swap. However, this change increases overrepresentation with respect to the natural measures discussed above. The number of candidates whose supporters are overrepresented increases: from $\{a, b, c\}$ to $\{e, f, g, h\}$. Similarly, the number of unique overrepresented voters grows from 4 (namely, voters approving $\{a, b, c\}$) to 6 (namely, all voters approving

at least d). The sum of the magnitudes of the violations also increases (from a total of 4 points over upper quota to 6).

Aside from the fact that this illustrates that another approach is required for showing that sequences of swaps eventually terminate, this might seem like a problem with JUQ. Indeed, an axiom that aims to prevent overrepresentation should not require swaps that increase overrepresentation. However, we would argue that this is not as problematic as it might first seem. Indeed, from the example above, W is the unique committee that minimises the three measures we mentioned. On the other hand, this committee violates a rather weak notion of efficiency: it includes e but not d , even though $N(e) \subset N(d)$. In other words, including d instead of e would make everyone equally satisfied or better off. So it appears that minimising overrepresentation (at least in the sense above) is not a reasonable objective, as it clashes with an extremely weak and desirable notion of efficiency. Any axiom that aims to capture overrepresentation in ABC voting has to cope with this fact: we choose to prioritise efficiency.

Still, we need a different potential function. Upon closer inspection, however, the proof of Theorem 3.5 already gives us an option: the Adams-AV score of a committee. Technically, the potential function here is the tuple-valued function

$$P(W) = (\text{score}_{\omega_1}(A, W), \text{score}_{\omega_2}(A, W)),$$

where ω_1 and ω_2 are the weight functions defining Adams-AV and tuples are compared lexicographically.¹⁶ Since there are only finitely many committees, no infinite sequences can exist.

Proposition 3.7. There are no infinite sequences of JUQ swaps.

Proof. By inspecting the proof of Theorem 3.5, we see that any JUQ swap either increases the number of covered voters or keeps the covered voters constant while increasing the ω_2 -score (where $\omega_2 = (1, 1, 1/2, 1/3, \dots)$). Since there are only finitely many committees, it follows that sequences of JUQ swaps must always terminate. $\square$

We will next study a related question: the complexity of computing and verifying JUQ committees.

3.4.5 Complexity of Justified Upper Quota: The Upper Quota Elimination Rule

We have shown above that sequences of JUQ swaps eventually terminate. This is normatively appealing; however, we have no bound on the size of such sequences. Clearly, every sequence of JUQ swaps has length at most $\binom{m}{k}$, since visiting the same

¹⁶If one insists in having a single-valued function, then we can use score_{ω} with $\omega = (1, 1/nk, 1/2nk, 1/3nk, \dots)$ instead (see Lemma 3.2). This makes no difference.

committee twice would lead to the possibility of an infinite sequence. On the other hand, in principle, one could need super-polynomially many swaps to reach a locally optimal committee (in terms of [Adams-AV](#) score) where no [JUQ](#) swaps are possible. To give an analogous example, for the [PAV](#) score, one might need super-polynomially many swaps to converge to a local optimum [[Kraiczy and Elkind, 2024](#)].

Therefore, [Proposition 3.7](#) gives us no information on whether one can compute an exhaustive committee that satisfies [JUQ](#) in polynomial time.¹⁷ Moreover, we have already mentioned that all known polynomial-time computable rules fail the axiom ([Section A.2.1](#) of the appendix), and computing a winning committee under [Adams-AV](#) cannot be done in polynomial time. This holds because [Adams-AV](#) refines the [CC](#) rule, which is NP-hard to compute [[Procaccia et al., 2008](#)].

Observation 3.8. For any refinement F of [Adams-AV](#), unless $P = NP$, there is no polynomial-time algorithm that computes F .

Proof. NP-hardness of [CC](#) is formulated by [Procaccia et al. \[2008\]](#) as: “Is there a committee W of size k with a [CC](#) score of at least ℓ ?”. Since [Adams-AV](#) refines [CC](#), a polynomial-time algorithm for any refinement of [Adams-AV](#) would solve this question. $\square$

Hence, the complexity of computing exhaustive [JUQ](#) committees is still open; we will tackle it in the remainder of this section. As a first observation, we note that checking whether a committee satisfies [JUQ](#) is decidable in polynomial time. This is similar to [EJR+](#) [[Brill and Peters, 2023](#)], whereas the same does not hold for the original [EJR](#) axiom [[Aziz et al., 2017](#)].

Proposition 3.9. Whether a committee satisfies [JUQ](#) is decidable in polynomial time.

Proof. Fix a committee W . Clearly, there are only polynomially many pairs of candidates c and d to check. Furthermore, there are polynomially many possible sizes of N' to check, since $n_d \in [n]$ for all $d \in C$. We show how to check, for each $(c, d) \in C^2$ and $s \in [n_d]$, whether a counterexample for [JUQ](#) exists with candidate pair (c, d) and some N' with $|N'| = s$.

First, we check whether the condition on $N(c)$ holds (i.e., all voters in this set are above their upper quota). If not, then we are done. If yes, then we compute the quota associated to s as $\lceil ks/n \rceil$ and check, for every $i \in N(d)$, whether

$$|A(i) \cap ((W \cup \{d\}) \setminus \{c\})| \leq \lceil ks/n \rceil$$

holds. If there exist at least s voters in $N(d)$ such that the previous conditions holds, then we can use such voters to form N' and [JUQ](#) is failed. Otherwise, no counterexample with N' of size s exists. $\square$

¹⁷Clearly, inexhaustive committees can be computed in constant time since the empty committee satisfies [JUQ](#).

Having made this observation, we will now exhibit an exhaustive and polynomial-time computable rule that satisfies **JUQ**. Our rule is similar to the **Greedy Justified Candidate Rule (GJCR)** [Brill and Peters, 2023], a polynomial-time computable rule that satisfies **EJR+** (but fails **JUQ**, see Section A.2.1 of the appendix). We recall the latter rule for convenience.

Definition 3.11. **GJCR** begins with the empty committee $W = \emptyset$ and iteratively adds candidates witnessing an **EJR+** violation to it. Among all **EJR+** violations witnessed by a candidate d and a nonempty set of voters $N' \subseteq N(d)$, **GJCR** picks one (according to some tiebreaking) maximising $|N'|$ and adds it to W . The rule terminates once no **EJR+** violations remain, returning a committee W with $|W| \leq k$.¹⁸

Observe that the above is a resolute rule, i.e., it always returns a single committee. The same will hold for the rule we introduce. This does not impact our analysis, and simplifies the exposition.

We take a similar sequential approach to devise the **Upper Quota Elimination Rule (UQER)**, showing that exhaustive committees satisfying **JUQ** can be found in polynomial time in the process. We take the “dual approach” to **GJCR**: instead of adding **EJR+** witnesses, our rule starts off with the entire set of candidates C and then iteratively deletes candidates witnessing an upper quota violation. This is repeated until either no candidate is overrepresented (and hence the remaining set of candidates satisfies **UQ**) or until the rule reaches a committee of size k . We will show that this committee and any set considered in between indeed satisfies **JUQ**.

Definition 3.12. **UQER** works as follows. Start by initialising $W_0 = C$, the set of all candidates. In each step $\ell \in \{1, 2, \dots\}$, we choose a candidate c_ℓ to remove from the set. To do so, we find all candidates $c \in W_{\ell-1}$ that cause a **UQ** violation for $W_{\ell-1}$, that is, candidates $c \in W_{\ell-1}$ such that

$$|A(i) \cap W_{\ell-1}| > \lceil q_c \rceil \quad \text{for all } i \in N(c).$$

Let c_ℓ be a candidate among these with minimal $|N(c)|$, with ties broken arbitrarily. Set $W_\ell = W_{\ell-1} \setminus \{c_\ell\}$, and repeat.

We continue until the first step ℓ^* in which either $|W_{\ell^*}| = k$ or W_{ℓ^*} satisfies **UQ**. In the first case, our rule outputs W_{ℓ^*} , and in the second case it outputs some subset $W \subseteq W_{\ell^*}$ with $|W| = k$ (again, ties broken arbitrarily).

We note for the interested reader that **UQER** is similar to the dual-quota apportionment method of Mayberry [1978]. However, even for apportionment profiles, **UQER** differs as we explicitly use the committee size k , and restrict our deletions to candidates who are strictly (not weakly) above the upper quota. Further, Mayberry [1978] uses a dif-

¹⁸This is technically a refinement of **GJCR** as defined by Brill and Peters [2023], which selects the candidate maximising $\lfloor q(N') \rfloor$ instead.

ferent method to select which candidates to delete. The general idea of starting from the full candidate set and iteratively removing candidates also appears in Thiele's elimination methods [Janson, 2018], also known as reverse sequential Thiele methods [Lackner and Skowron, 2023]. These methods provide an alternative to the sequential methods discussed in Section 2.2.7 for computing Thiele rules in polynomial time, and were already proposed by Thiele [1895] himself.

We show the following.

Theorem 3.10. **UQER** returns an exhaustive committee satisfying **JUQ** in polynomial time.

Proof. Firstly, observe that **UQER** runs in polynomial time, as we only need to check which candidates are overrepresented for at most $m - k$ iterations. To prove that the output of **UQER** satisfies **JUQ**, we first show that **JUQ** is always maintained during the execution of **UQER**. In this proof, let $\{W_0, W_1, \dots, W_{\ell^*}\}$ be the intermediate sets of candidates produced during the execution of **UQER**, as described in Definition 3.12.

It is easy to see that the initial set $W_0 = C$ satisfies **JUQ**, as there are no candidates outside the set that can be part of a **JUQ** violation. We can also observe that the set of **UQ**-violating candidates never increases during the execution of **UQER**.

Suppose at some point we encounter a **JUQ** violation in the set W_ℓ , witnessed by some candidate pair (c, d) and nonempty voter set $N' \subseteq N(d)$. Note that c need not be the same as $c_{\ell+1}$, i.e., the candidate that will be removed from W_ℓ to obtain $W_{\ell+1}$. This means that

$$|A(i) \cap (W_\ell \setminus \{c\} \cup \{d\})| \leq \lceil q(N') \rceil \leq \lceil q_d \rceil$$

for all $i \in N'$, and thus $|A(i) \cap (W_\ell \setminus \{c\})| < \lceil q_d \rceil$.

Since $d \notin W_\ell$, candidate d must have been removed in some earlier step s , i.e., $d = c_s$ for some $s < \ell$. As d was removed in this step, we know that $|A(i) \cap W_{s-1}| > \lceil q_d \rceil$ for all $i \in N(d)$ as d was a **UQ**-violating candidate for W_{s-1} . Thus, after d was removed, $|A(i) \cap W_s| \geq \lceil q_d \rceil$ for all $i \in N(d)$. Furthermore, by our selection of d we know that d was among the overrepresented candidates with the fewest approvals. Hence, $q_d \leq q_e$ for any other candidate e whose supporters are overrepresented at step s . Now let $i \in N' \setminus N(c)$. Note that such an i must exist, as otherwise we would have $N' \subseteq N(c)$ and thus $q(N') \leq q_c$, which contradicts $\lceil q_c \rceil < |A(j) \cap W_\ell| \leq \lceil q(N') \rceil$ for $j \in N' \cap N(c)$.

As $|A(i) \cap W_s| \geq \lceil q_d \rceil$ and $|A(i) \cap W_\ell| = |A(i) \cap (W_\ell \setminus \{c\})| < \lceil q_d \rceil$, there must exist a candidate in $A(i) \cap (W_s \setminus W_\ell)$. Let $e \in A(i) \cap (W_s \setminus W_\ell)$ be the last candidate from this set to be removed from the outcome. In particular, assume that e was removed at step r between step s and step ℓ . As we assumed e to be the last candidate in $A(i) \cap (W_s \setminus W_\ell)$ to be removed, we know that $|A(i) \cap (W_{r-1} \setminus \{e\})| < \lceil q_d \rceil$.

Observe that when d is removed, both e and d witness an **UQ** violation, and that d is removed first. Hence, by our assumption that d was a candidate with the fewest approvals when removed at step s we know that $n_e \geq n_d$ and hence $\lceil q_e \rceil \geq \lceil q_d \rceil$. Therefore, since i approves e , it must hold that $|A(i) \cap W_{r-1}| > \lceil q_e \rceil \geq \lceil q_d \rceil$. Since all

these quantities are integers, this is a contradiction to $|A(i) \cap (W_{r-1} \setminus \{e\})| < \lceil q_d \rceil$, which we derived earlier. Thus (c, d) could not have been a **JUQ** violation. Hence, every candidate set in $\{W_0, W_1, \dots, W_{\ell^*}\}$ satisfies **JUQ**.

With this, we can finally prove our theorem. Indeed, we know that W_{ℓ^*} satisfies **JUQ**. If $|W_{\ell^*}| = k$ then we are done. Meanwhile, if $|W_{\ell^*}| > k$ we know that W_{ℓ^*} additionally satisfies **UQ**. Thus, any subset $W \subseteq W_{\ell^*}$ will satisfy **UQ** and thus **JUQ**. $\square$

Similarly to Theorem 3.10, we can show **UQER** maintains **EJR+** during its execution up to, but not including, the final step. We will discuss this result later, in Section 3.6. We stress that the result does not necessarily hold for the arbitrary subset $W \subseteq W_{\ell^*}$ that is returned if $|W_{\ell^*}| > k$, i.e., **UQER** fails **EJR+** (Proposition 3.22).

Proposition 3.11. Every intermediate set of candidates $\{W_0, W_1, \dots, W_{\ell^*}\}$ produced during the execution of **UQER** satisfies **EJR+**.

Proof. It is easy to see that the initial committee $W_0 = C$ satisfies **EJR+**, as there are no candidates outside the committee that can witness an **EJR+** violation. We show by induction that **UQER** maintains **EJR+**. Again, let $\{W_0, W_1, \dots, W_{\ell^*}\}$ be the intermediate sets of candidates produced during the execution of **UQER**, as described in Definition 3.12.

Suppose removing a candidate c_ℓ is the first time we encounter an **EJR+** violation. As we started from the set of all candidates C , this violation must involve some previously removed candidate c_s with $s < \ell$ and voter group $N' \subseteq N(c_s)$. This means that $|N'| \leq |N(c_s)| \leq |N(c_\ell)|$ from the definition of our rule, and thus $q(N') \leq q_{c_s} \leq q_{c_\ell}$.

We know that $|A(i) \cap W_{\ell-1}| > \lceil q_{c_\ell} \rceil$ for all $i \in N(c_\ell)$, as c_ℓ violated **UQ** in $W_{\ell-1}$. We also know that $|A(i) \cap (W_{\ell-1} \setminus \{c_\ell\})| < \lfloor q(N') \rfloor$ for all $i \in N'$, because of our **EJR+** violation assumption.

Suppose $N(c_\ell) \cap N' \neq \emptyset$ and let $i^* \in N(c_\ell) \cap N'$. From the above,

$$\lfloor q(N') \rfloor \geq |A(i^*) \cap W_{\ell-1}| > \lceil q_{c_\ell} \rceil \geq \lceil q(N') \rceil,$$

which is a contradiction. Thus, $N(c_\ell) \cap N' = \emptyset$. This means that

$$|A(i) \cap W_{\ell-1}| = |A(i) \cap (W_{\ell-1} \setminus \{c_\ell\})| < \lfloor q(N') \rfloor$$

for all $i \in N'$, and thus c_s with voter group N' would have violated **EJR+** in $W_{\ell-1}$, which contradicts our assumption that W_ℓ is the first time we encounter an **EJR+** violation.

Thus every candidate set from $\{W_0, W_1, \dots, W_{\ell^*}\}$ satisfies **EJR+**. $\square$

This concludes our analysis of **JUQ**.

3.5 Balancing Underrepresentation and Overrepresentation

We now turn to our second view on overrepresentation, based on **Apportionment Near Quota (ANQ)** [Balinski and Young, 2001]. Under this view, there should not exist a party

i which is overrepresented and a party j which is underrepresented such that by moving a seat from party i to party j we can bring both parties closer to their respective quotas. Our goal is to generalise this to the **ABC** voting setting, by again interpreting single candidates as parties that represent their voters.

In this spirit, there should not exist an in-committee candidate whose approvers are overrepresented and an out-of-committee candidate with underrepresented approvers such that swapping them would make both groups better represented. To better interpret this, we again return to the apportionment setting. We first give the following equivalent formulation of **ANQ** (Definition 3.3). Two parties i and j witnessing an **ANQ** violation is logically equivalent to the first party being more than $1/2$ above its quota, and the second party being more than $1/2$ below its quota.

Observation 3.12. An allocation a satisfies **ANQ** if and only if there are no two parties $i, j \in [p]$ such that

$$a_i - q(v_i) > \frac{1}{2} \text{ and } q(v_j) - a_j > \frac{1}{2}.$$

Proof. The claim follows immediately via algebraic manipulation. $\square$

To generalise this to committee voting, we have to account for potentially overlapping approval sets. Assume that $c \in W$ is the overrepresented candidate and $d \in C \setminus W$ is the underrepresented one and let $N' \subseteq N(d)$ be the set of underrepresented voters approving d . Then we say that c and d witness a **Justified Near Quota (JNQ)** violation if swapping c with d moves every voter in $N(c)$ closer to their quota q_c *whenever their satisfaction changes* (i.e., when they do not also approve of d) and every voter in N' moves closer to their quota $q(N')$ whenever their satisfaction changes. For voters approving of both of c and d the satisfaction does not change. For these voters we merely require that they are above the quota q_c and below the quota $q(N')$ to ensure these voters are not on the “wrong side” of their respective quotas.

Definition 3.13. A set W satisfies **JNQ** if there does not exist a pair of candidates $c \in W$ and $d \in C \setminus W$ and nonempty set of voters $N' \subseteq N(d)$ such that

$$\begin{aligned} |A(i) \cap W| - q_c &> \frac{1}{2} \cdot \mathbb{1}[d \notin A(i)] && \text{for all } i \in N(c), \text{ and} \\ q(N') - |A(i) \cap W| &> \frac{1}{2} \cdot \mathbb{1}[c \notin A(i)] && \text{for all } i \in N'. \end{aligned}$$

The following example shows that, as one would expect, **JNQ** and **JUQ** are two distinct notions.

Example 3.11. Consider the profile $A = (\{a\}, \{a\}, \{a\}, \{a, c\}, \{b\})$ and let $k = 2$, so the quota of a single voter is $2/5$. The committee $\{a, c\}$ satisfies **JNQ**. Indeed, replacing either a or c with b would change the satisfaction of the voter submitting $\{b\}$ from 0 to 1, increasing the distance from their quota of $2/5$. On the other hand, this committee fails **JUQ** as the voter approving $\{a, c\}$ is overrepresented and we could swap c with b . Conversely, the committee $\{b, c\}$ satisfies **JUQ**, but fails **JNQ**. Indeed, here we should swap b with a , and the set $N' \subseteq N(a)$ can consist for example of voters 1 and 2. Finally, the committee $\{a, b\}$ satisfies both axioms.

Following Observation 3.12 we can see that **JNQ** generalises **ANQ**.

Proposition 3.13. For apportionment profiles, **JNQ** is equivalent to **ANQ**.

Proof. If a committee satisfies **JNQ**, then it clearly satisfies **ANQ**. Let us prove the inverse direction.

Suppose that a satisfies **ANQ**. Assume towards a contradiction that c, d and N' are witnesses to a **JNQ** violation. Let i be the party of voters approving c and j the party of voters approving j .

First, assume $i = j$. This implies $N' \subseteq N(c)$ and hence $q(N') \leq q_c$. However, the inequalities that define **JNQ** reduce to $a_i > q_c$ and $q(N') > a_i$. This is a contradiction.

Hence, assume $i \neq j$. In this case, the first condition of **JNQ** reduces to $a_i - q_c > 1/2$ and the second to $q(N') - a_j > 1/2$. Since $N' \subseteq N(d)$, we have $q(N') \leq q_d$ and hence $q_d - a_j < 1/2$. This is a violation of **ANQ**, completing the proof. $\square$

Next, similarly to **Adams-AV** and **JUQ**, we get that every refinement of **SL** satisfies **JNQ**. The proof is analogous to the proof of Theorem 3.5. That is, we show that if a committee W does not satisfy **JNQ** as witnessed by two candidates c and d then swapping c with d increases the **SL** score.

Theorem 3.14. Any refinement of **SL** satisfies **JNQ**.

Proof. To begin, observe that a **JNQ** violation for W implies $2 \cdot q_c < 2|A(i) \cap W| - 1$ for all voters $i \in N(c) \setminus N(d)$ and $2q(N') > 1 + 2|A(i) \cap W|$ for all voters $i \in N' \setminus N(c)$.

Assume towards a contradiction that we have a **JNQ** violation witnessed by $c \in W$, $d \in C \setminus W$ and $N' \subseteq N(d)$ that changes the **SL** score by:

$$\begin{aligned} \text{score}_\omega(A, (W \cup \{d\}) \setminus \{c\}) - \text{score}_\omega(A, W) &\geq \\ \sum_{i \in N(c) \setminus N'} -\frac{1}{2|A(i) \cap W| - 1} + \sum_{i \in N' \setminus N(c)} \frac{1}{2|A(i) \cap W| + 1} &> \\ -\frac{|N(c) \setminus N'|}{2q_c} + \frac{|N' \setminus N(c)|}{2q(N')} &. \end{aligned}$$

We consider two cases. If $N' \cap N(c) = \emptyset$, then:

$$\begin{aligned} \text{score}_\omega(A, (W \cup \{d\}) \setminus \{c\}) - \text{score}_\omega(A, W) &> -\frac{|N(c) \setminus N'|}{2q_c} + \frac{|N' \setminus N(c)|}{2q(N')} = \\ &= -\frac{|N(c)|}{2q_c} + \frac{|N'|}{2q(N')} = -\frac{n}{2k} + \frac{n}{2k} = 0. \end{aligned}$$

Suppose instead that there exists some voter $i \in N' \cap N(c)$. By definition of **JNQ**, this implies $q_c < q(N')$. Then:

$$\begin{aligned} \text{score}_\omega(A, (W \cup \{d\}) \setminus \{c\}) - \text{score}_\omega(A, W) &> \\ -\frac{|N(c) \setminus N'|}{2q_c} + \frac{|N' \setminus N(c)|}{2q(N')} &= -\frac{|N(c)|}{2q_c} + \frac{|N(c) \cap N'|}{2q_c} + \frac{|N'|}{2q(N')} - \frac{|N(c) \cap N'|}{2q(N')} = \\ |N(c) \cap N'| \left(\frac{1}{2q_c} - \frac{1}{2q(N')} \right) &> 0. \end{aligned}$$

Hence, the **SL** score would strictly increase by swapping candidate c with d . Contradiction: **SL** satisfies **JNQ**. $\square$

Next, we again obtain a characterisation, this time of refinements of **SL**.

Theorem 3.15. A composite Thiele rule satisfies **ANQ** if and only if it is a refinement of **SL**.

Proof. The first direction of the theorem follows from Theorem 3.14 and the fact that **JNQ** and apportionment near quota are equivalent on apportionment profiles (Proposition 3.13).

We show the second direction by proving that for any simple Thiele rule except for **SL** there exists an apportionment profile on which it selects only a single committee which violates **ANQ**. As it only selects a single committee, every refinement of this rule must also select this single committee and therefore violate **ANQ**. Recall that we indicate apportionment profiles as $A = (v_1, \dots, v_p)$.

Consider any other Thiele vector $\omega = (1, \omega(2), \omega(3), \dots)$ and let $\omega(\ell)$ be the first coordinate in which ω differs from **SL**. We consider two cases, namely $\omega(\ell) > 1/(2\ell-1)$ and $\omega(\ell) < 1/(2\ell-1)$.

Case 1: $\omega(\ell) > 1/(2\ell-1)$ or, equivalently $1/\omega(\ell) < 2\ell-1$. We fix $s/r \in \mathbb{Q}$ such that

$$\frac{1}{\omega(\ell)} < \frac{s}{r} < 2\ell-1$$

and consider a two-party apportionment profile with $A = (s, r)$ and $k = \ell$. We note that $s \cdot \omega(\ell) > r$. Hence, the Thiele rule would give all ℓ seats to the first party. However, the

quota of the first party is

$$q(v_1) = \ell \cdot \frac{s}{s+r} = \ell \cdot \frac{1}{1+\frac{r}{s}} < \ell \cdot \frac{1}{1+\frac{1}{(2\ell-1)}} = \ell - \frac{1}{2},$$

while the quota of the second party is

$$q(v_2) = \ell \cdot \frac{r}{s+r} = \ell \cdot \frac{1}{\frac{s}{r}+1} > \ell \cdot \frac{1}{2\ell-1+1} = \frac{1}{2}.$$

This is a violation of [ANQ](#).

Case 2: $\omega(\ell) < 1/(2\ell-1)$ or, equivalently, $1/\omega(\ell) > 2\ell-1$. We fix $s/r \in \mathbb{Q}$ such that $1/\omega(\ell) > s/r > 2\ell-1$ and again consider an apportionment profile with $A = (s, r)$ and $k = \ell$. We claim that the Thiele rule would select only $\mathbf{a} = (\ell-1, 1)$. Indeed, it holds that $r > \omega(\ell) \cdot s$, which implies that the second party receives at least one seat. Moreover, since ℓ is the first entry for which ω differs from the weights of [SL](#), and since one can verify that $r/3 < s/(2\ell-3)$, the second party receives at most one seat. For the quotas we compute

$$\begin{aligned} q(v_1) &= \ell \cdot \frac{s}{s+r} = \ell \cdot \frac{1}{1+\frac{r}{s}} > \ell \cdot \frac{1}{1+\frac{1}{(2\ell-1)}} = \ell - \frac{1}{2} \quad \text{and} \\ q(v_2) &= \ell \cdot \frac{r}{s+r} = \ell \cdot \frac{1}{\frac{s}{r}+1} < \ell \cdot \frac{1}{2\ell-1+1} = \frac{1}{2}. \end{aligned}$$

This is again a violation of [ANQ](#), which concludes the proof. $\square$

What about rules other than composite Thiele rules? Just as we did for [JUQ](#), we exhibit in [Appendix A](#) (Section [A.2.2](#)) examples of essentially all familiar voting rules failing [JNQ](#). We note that the variance-Phragmén voting rule [[Brill et al., 2024](#)], which is equivalent to Sainte-Lagué’s method on apportionment instances, also fails [JNQ](#).

As an immediate consequence of [Theorem 3.14](#), analogously to the [JUQ](#) case, we get that there are no infinite sequences of [JNQ](#) swaps, as any swap increases the [SL](#) score.

Proposition 3.16. There are no infinite sequences of [JNQ](#) swaps.

Proof. The proof is analogous to the proof of [Proposition 3.7](#) by referencing the proof of [Theorem 3.14](#) instead of [Theorem 3.5](#). $\square$

Finally, similarly to [JUQ](#) and [EJR+](#), we show that it can be verified in polynomial time whether a given committee satisfies [JNQ](#).

Proposition 3.17. Whether a committee satisfies **JNQ** is decidable in polynomial time.

Proof. This proof is essentially analogous to the proof of Proposition 3.9. The only difference is that, for every voter in $i \in N(d)$, we check whether $k^\ell/n - |A(i) \cap W| > 1/2 \cdot \mathbb{1}[c \notin A(i)]$ holds instead. $\square$

We leave it as an open question whether exhaustive committees satisfying **JNQ** can be computed in polynomial time. Observe that **SL** is hard to compute [Skowron et al., 2016].

3.6 Combining Lower, Near, and Upper Quota Axioms

In this section, we discuss the connections between standard proportionality axioms for **ABC** elections and the axioms we propose. In particular, we investigate whether upper, near, and lower quota axioms can be simultaneously satisfied.

3.6.1 Some Remarks on Thiele Rules

First, we observe that, as a refinement of **CC**, **Adams-AV** satisfies a rather weak under-representation axiom, **JR**, in addition to satisfying **JUQ** and being exhaustive. Can we do better? In particular, is it always possible to find an (exhaustive) rule that satisfies **EJR+** and **JUQ**? What about **EJR+**, **JUQ** and **JNQ**?

Unfortunately, in the class of Thiele rules, this is not the case. We have seen a “tripartition” with regard to proportionality axioms. **PAV**, and its refinements, are the only composite Thiele rules satisfying **EJR+**, **SL** refinements are the only ones satisfying **JNQ**, and **Adams-AV** refinements are the only ones satisfying **JUQ**, with these (im)possibilities mirroring the apportionment setting [Balinski and Young, 2001].

For the apportionment setting, however, Balinski and Young [2001, Proposition 6.5] prove a curious fact about divisor methods: any allocation selected by a divisor method either satisfies **ALQ** or **AUQ**. Interestingly enough, we can generalise this result to the **ABC** voting setting and show that for a large class of Thiele rules (corresponding to the divisor methods as defined by Balinski and Young), every selected committee satisfies either **UQ** or **EJR+**. This class includes, among others, **PAV** and **SL**, and the result extends to **Adams-AV** (see Corollary 3.19).

Theorem 3.18. Let ω be a weight function with strictly positive weights for which there exists an $\alpha \in \mathbb{R}_{>0}$ satisfying either

1. $\ell - 1 < \alpha/\omega(\ell) \leq \ell$ for all $\ell \in \mathbb{N}$ or
2. $\ell - 1 \leq \alpha/\omega(\ell) < \ell$ for all $\ell \in \mathbb{N}$.

Then any committee selected by the simple Thiele rule defined by ω satisfies **UQ** or **EJR+** (or both).

Proof. Assume that W is selected by THL_ω and that it fails both **UQ** and **EJR+**. Let c and d be the respective witnesses of this. Thus, it holds that $|A(i) \cap W| > \lceil q_c \rceil$ for every $i \in N(c)$ as well as $|A(i) \cap W| < \lfloor q(N') \rfloor$ for some nonempty $N' \subseteq N(d)$ and every $i \in N'$.

Now consider the score change by exchanging c with d . The score of every voter $i \in N' \setminus N(c)$ increases by $\omega(|A(i) \cap W| + 1)$, while the score of every $i \in N(c) \setminus N'$ decreases by $\omega(|A(i) \cap W|)$. However,

$$\begin{aligned} \sum_{i \in N' \setminus N(c)} \omega(|A(i) \cap W| + 1) &- \sum_{i \in N(c) \setminus N'} \omega(|A(i) \cap W|) \geq \\ \sum_{i \in N'} \omega(|A(i) \cap W| + 1) &- \sum_{i \in N(c)} \omega(|A(i) \cap W|) \geq \\ \sum_{i \in N'} \omega(\lfloor q(N') \rfloor) &- \sum_{i \in N(c)} \omega(\lceil q_c \rceil + 1) > \alpha \left(\frac{|N'|}{q(N')} - \frac{n_c}{q_c} \right) = 0. \end{aligned}$$

Hence, by swapping c with d the score would increase. This gives us a contradiction and the proof is concluded. $\square$

As a corollary, we immediately obtain that this is true for **PAV**, **SL**, and **Adams-AV**.

Corollary 3.19. Any committee selected by **PAV**, **SL**, or **Adams-AV** satisfies **UQ** or **EJR+** (or both).

Proof. **PAV** and **SL** satisfy the definition with $\alpha = 1$ and $\alpha = 1/2$, respectively. For **Adams-AV**, we leverage Lemma 3.2 and observe that, for every $n, k \in \mathbb{N}$ such that $nk > 1$, the vector $\omega = (1, 1/nk, 1/2nk, 1/3nk, \dots)$ satisfies the condition with $\alpha = 1/nk$. The remaining case with $nk = 1$ is trivial. $\square$

We stress that Theorem 3.18 and Corollary 3.19 do *not* imply that a rule in this class satisfies either **EJR+** or **UQ** globally. Rather, they guarantee that every committee selected by the rule satisfies at least one of the two conditions. Which condition is satisfied may vary across instances: on some instances, a selected committee may satisfy **EJR+** but fail **UQ**, whereas on others the converse may hold. An immediate consequence is that these rules necessarily satisfy **EJR+** on instances where no **UQ** committee exists.

Moreover, the classical result of Balinski and Young [2001] was partially strengthened by Cembrano et al. [2025] for a restricted class of divisor methods: they showed that, for every fixed apportionment profile, there exists a divisor method satisfying both **AUQ** and **ALQ** in that instance. An interesting open question is whether a similar result holds for **ABC** voting—namely, whether for every instance there exists a (composite) Thiele method that satisfies both **EJR+** and **JUQ** in that very instance.

3.6.2 Compatibility for Inexhaustive Committees: The Proportional Elimination Rule

If we consider arbitrary **ABC** voting rules, we can show that **UQ** (and thus **JUQ**), **JNQ** and **EJR+** are jointly satisfiable, as long as we do not require exhaustiveness. To this end, we define a voting rule in a similar way to **UQER**. The differences lie in the initial committee and in the method for selecting candidates to remove.

Definition 3.14. The **Proportional Elimination Rule (PER)** works as follows. Start from any committee W_0 output by **GJCR** (Definition 3.11). This committee has size $|W_0| \leq k$. In each step $\ell \in \{1, 2, \dots\}$, we find the set C^* of candidates $c \in W_{\ell-1}$ whose supporters are all strictly above lower quota, i.e.,

$$|A(i) \cap W_{\ell-1}| > \lfloor q_c \rfloor \quad \text{for all } i \in N(c).^{19}$$

If there are no such candidates then the rule outputs $W_{\ell-1}$. Otherwise, let c_ℓ be the candidate in C^* with fewest approvals (with ties broken arbitrarily). Set $W_\ell = W_{\ell-1} \setminus \{c_\ell\}$, and repeat.

Observe, similarly to **UQER**, that this rule is resolute—it returns a singleton, that is, a single committee.

Proposition 3.20. The (possibly inexhaustive) committee W produced by **PER** satisfies **UQ**, **JNQ** and **EJR+**.

Proof. Let W be the output of **PER**. This rule ensures that for each $c \in W$ there exists $i \in N(c)$ with $|A(i) \cap W| \leq \lfloor q_c \rfloor \leq \lceil q_c \rceil$, and thus satisfies **UQ** by construction.

Suppose that W violates **JNQ**, witnessed by $c \in W$ and $d \in C \setminus W$. Thus, for all $i \in N(c)$, we have $|A(i) \cap W| > q_c$. However, the rule ensures $|A(i) \cap W| \leq \lfloor q_c \rfloor$ for some $i \in N(c)$ and thus $|A(i) \cap W| \leq q_c$, a contradiction.

Suppose now that W violates **EJR+**, witnessed by $d \in C \setminus W$ and nonempty $N' \subseteq N(d)$. Thus, we know that $|A(i) \cap W| < \lfloor q(N') \rfloor$ for every voter $i \in N'$. Note that the committee W_0 output by **GJCR** satisfies **EJR+**. We let $c^* \in W_0$ be the first candidate after whose deletion the **EJR+** violation with d occurs. Call the committee after c^* was deleted W' . As W' is the first committee witnessing the **EJR+** violation, there must exist a voter $i \in N'$ such that $|A(i) \cap W'| < \lfloor q(N') \rfloor$ and $|A(i) \cap (W' \cup \{c^*\})| \geq \lfloor q(N') \rfloor$ as, otherwise, the violation would have already occurred before. This, in particular, implies that $i \in N(c^*)$. Hence, by definition of **PER**, for this same voter we get that $\lfloor q_{c^*} \rfloor \leq |A(i) \cap W'| < \lfloor q(N') \rfloor \leq \lfloor q_d \rfloor$. Note that this implies that $|N(c^*)| < |N'| \leq |N(d)|$.

First, suppose that $d \in W_0$. As $d \notin W'$, we know that d must have been removed before c^* . However, at the step in which d was removed, c^* would also have been eligible

¹⁹Note that this is different to **UQER**, which removed candidates violating **UQ**, i.e., those with all supporters above their upper quota.

for removal. As $|N(c^*)| < |N(d)|$ our rule would have removed c^* over d , leading to a contradiction.

Hence, we know that $d \notin W_0$. In other words, d was not selected by **GJCR**. During the execution of the latter rule, consider the first moment in which a candidate e with $|N(e)| \leq |N(c^*)|$ is selected by **GJCR**, and let W'' be the set of candidates—as per the definition of **GJCR**—before the inclusion of e . To obtain W' from W_0 , **PER** only removed candidates with weakly fewer than $|N(c^*)|$ supporters, which means $W'' \subseteq W'$. Therefore, $|A(i) \cap W''| \leq |A(i) \cap W'| < \lfloor q(N') \rfloor$ for all voters $i \in N'$, and hence, d and N' would have been an **EJR+** violation for W'' . Since

$$|N(e)| \leq |N(c^*)| < |N'| \leq |N(d)|,$$

candidate d would have been chosen to be added to W'' instead of e . Contradiction. $\square$

With exhaustiveness, the picture gets more complicated. Indeed, it remains an open problem whether any two of the axioms **JUQ**, **JNQ**, and **EJR+** can be satisfied simultaneously for exhaustive committees. In the next section, we explore the interplay of **JUQ**, **EJR+** and exhaustiveness in more detail. In particular, we show that some natural avenues for proving compatibility do not work.

3.6.3 Some Difficulties in Combining Exhaustiveness, Lower and Upper Quota

A voting rule satisfying **JUQ**, **EJR+** and exhaustiveness is unlikely to be as straightforward as some of the voting rules we proposed earlier (i.e., **UQER** and **PER**). One might be tempted to modify the outcome of a rule that satisfies two of **EJR+**, **JUQ** and exhaustiveness to also satisfy the third. Below we provide negative results for three of the most natural modifications.

First, we show that **PER** cannot be made exhaustive while satisfying **JUQ** by adding additional candidates.

Proposition 3.21. There exists an instance where every exhaustive superset of the outcome of **PER** fails **JUQ**.

Proof. Consider the following election with $n = 63$ and $k = 7$:

Candidate/Approvals	3	3	3	14	7	7	14	12
c_1, c_2				✓	✓			✓
c_3, c_4							✓	✓
c_5	✓	✓	✓					
c_6	✓	✓		✓				
c_7	✓		✓		✓	✓		
c_8		✓	✓					✓

The column headers, here, refer to the number of voters submitting a certain ballot. For instance, there are three voters approving exactly $\{c_5, c_6, c_7\}$. Here, independently of

the tiebreaking, **GJCR** selects the inexhaustive committee $W = \{c_1, c_2, c_3, c_4, c_5\}$. Notice that, for every $c \in W$ and $i \in N(c)$, we have that $|A(i) \cap W| \leq \lfloor q_c \rfloor$. Thus **PER** also selects W .

Now consider any possible completion of this committee to 7 candidates. Without loss of generality, consider $W' = \{c_1, \dots, c_7\}$. Then c_5 and c_8 witness a JUQ violation with $N' = N(c_8)$. Indeed, $q_{c_5} = 1$ and every voter $i \in N(c_5)$ has $|A(i) \cap W'| = 2$. Moreover, $\lceil q_{c_8} \rceil = 3$ and every voter $i \in N(c_8)$ has $|A(i) \cap ((W' \setminus \{c_5\}) \cup \{c_8\})| \leq 3$. $\square$

Secondly, recall that the intermediate sets of candidates $\{W_0, \dots, W_{\ell^*}\}$ produced during the execution of **UQER** all satisfy **EJR+** (Proposition 3.11). One might tempted to modify the final step of **UQER**, where we pick an arbitrary subset of W_{ℓ^*} of size k , to ensure that the subset which we pick satisfies **EJR+**. Unfortunately, this cannot work in general.

Proposition 3.22. Consider the intermediate set of candidates W_{ℓ^*} produced by **UQER** as defined in Definition 3.12. This set need not contain a subset of size k that satisfies **EJR+**.

Proof. Consider the following election with $n = 15$ and $k = 10$:

Candidate/Voter	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
$\hat{c}_1, \hat{c}_2$	✓	✓	✓												
c_1, c_2	✓									✓	✓	✓			
c_3, c_4		✓								✓	✓	✓			
c_5, c_6			✓							✓	✓	✓			
$\hat{c}_3, \hat{c}_4$				✓	✓	✓									
c_7, c_8				✓						✓	✓	✓			
c_9, c_{10}					✓					✓	✓	✓			
c_{11}, c_{12}						✓				✓	✓	✓			
$\hat{c}_5, \hat{c}_6$							✓	✓	✓						
c_{13}, c_{14}							✓			✓	✓	✓			
c_{15}, c_{16}								✓		✓	✓	✓			
c_{17}, c_{18}									✓	✓	✓	✓			
c_1^*, c_2^*													✓	✓	✓

(Dashed lines and colours are intended as visual helpers and carry no additional meaning.) In order to obtain W_{ℓ^*} , **UQER** removes $\{\hat{c}_1, \dots, \hat{c}_6\}$, as these are the **UQ**-violating candidates with fewest supporters. Then, the rule needs to select some subset $W \subseteq W_{\ell^*}$ of size k . In order to satisfy **EJR+**, W must contain:

- At least 3 candidates from each of $\{c_1, \dots, c_6\}$, $\{c_7, \dots, c_{12}\}$ and $\{c_{13}, \dots, c_{18}\}$. This holds because each group of three voters among the first nine must have one voter who approves of at least two candidates in W and another who approves of at least one.
- Both c_1^* and c_2^* , as the last three voters together deserve at least 2 candidates according to **EJR+**.

This is $11 > k$ candidates. Thus W_{ℓ^*} need not contain a subset of size k that satisfies **EJR+**. $\square$

Finally, one might wonder whether we can start from an exhaustive committee satisfying **EJR+**, and modify it by *cleverly* choosing **JUQ** swaps, such that **EJR+** is maintained. Unfortunately, this does not work in general.

Proposition 3.23. There exist an exhaustive committee W such that W satisfies **EJR+** and, for every **JUQ** swap (c, d) , committee $(W \setminus \{c\}) \cup \{d\}$ fails **EJR+** but satisfies **JUQ** (i.e., no further swaps are possible).

Proof. Consider an **ABC** election with $k = 5$ and

$$A = (\{a, b\} \times 3, \{a, b, c\}, \{c, d, e\}, \{d, e\} \times 3, \{f\}, \{g\}).$$

Committee $W = \{a, c, d, f, g\}$ satisfies **EJR+**. Indeed, consider $b \notin W$. If we consider a group $N' \subseteq N(b)$ with $|N'| \leq 3$, this group will have lower quota at most $\lfloor 3 \cdot 5/10 \rfloor = 1$, but all supporters of b have satisfaction 1 with W . On the other hand, group $N(b)$ has lower quota 2 but includes a voter approving $\{a, b, c\}$ with satisfaction 2. A similar argument shows that e is involved in no **EJR+** violations.

However, c witnesses a **JUQ** violation with either b and voter set $N(b)$ or e and voter set $N(e)$. These are the only **JUQ** violations, and either swap would lead to a committee that fails **EJR+**. For example, suppose we swap c with b . Then $N(e)$ has lower quota 2, but all its members would have a satisfaction of 1. Similarly, one can check that both swaps would result in a committee satisfying **UQ** and thus **JUQ**. $\square$

We find that this claim holds even if we select a particularly well chosen **EJR+** committee, such as the output of the method of equal shares [Peters and Skowron, 2020]. Since this statement, as well as its proof, are rather involved and require an additional definition, we defer this result to Appendix A (Proposition A.1).

Despite these difficulties, for other popular notions of proportionality, we can give a clear answer. We do so in the next section, which concludes the technical content of the chapter.

3.6.4 Justified Upper Quota and Other Proportionality Notions

In this section, we will present some other proportionality notions and show that they are incompatible with **JUQ**. In particular, we will consider *priceability* [Peters and Skowron, 2020], *fully justified representation* [Peters et al., 2021], and *perfect representation* [Sánchez-Fernández et al., 2026]. The impossibilities regarding the latter two hold even without requiring exhaustiveness (though perfect representation requires exhaustiveness by itself). We will keep this discussion brief, that is, we will not provide examples and in-depth explanations, as this would require too much of a detour. The interested reader can consult the book of Lackner and Skowron [2023, Chapter 4].

Let us begin with priceability. Intuitively, priceability tries to distribute power to voters equally. Each voter is (virtually) endowed with the same budget, and voters should be able to pool their money together and buy candidates they jointly approve of. A committee that *could* be obtained with this process is said to be priceable.

Definition 3.15. A committee W is *priceable* if there exists a per-voter budget $\alpha \in \mathbb{R}_{>0}$ and a payment function $\delta_i: C \rightarrow [0, 1]$ for each $i \in N$ such that:

1. $\sum_{c \in C} \delta_i(c) \leq \alpha$ for each $i \in N$;
2. $\delta_i(c) = 0$ for each $i \in N$ and $c \notin A(i)$;
3. $\sum_{i \in N} \delta_i(c) = \mathbb{1}[c \in W]$;
4. $\sum_{i \in N(c)} (\alpha - \sum_{d \in W} \delta_i(d)) \leq 1$ for each $c \notin W$.

Fully justified representation is a strengthening of [EJR](#). This axiom guarantees representation to groups that are only *partially* cohesive.

Definition 3.16. Given an integer $\ell \in \mathbb{N}$ and a set of candidates $C' \subseteq C$, a group of voters $N' \subseteq N$ is *weakly (ℓ, C') -cohesive* if $|N'| \geq |C'| \cdot n/k$ and if for all $i \in N'$ it holds that $|A(i) \cap C'| \geq \ell$.

A committee W satisfies *fully justified representation* if for every $\ell \in \mathbb{N}$, set of candidates $C' \subseteq C$, and weakly (ℓ, C') -cohesive group of voters N' there is some $i \in N'$ with $|A(i) \cap W| \geq \ell$.

Finally, perfect representation requires that, whenever possible, we should represent all voters perfectly. It applies only to profiles with a very specific structure. As discussed by [Sánchez-Fernández et al. \[2026\]](#), the idea behind this axiom is similar to that of the rule proposed by [Monroe \[1995\]](#), which we mentioned in [Section 2.1](#).

Definition 3.17. A committee W satisfies *perfect representation* if N can be partitioned in k disjoint groups $N_1, \dots, N_k$ of size n/k each and if we can assign to each group N_ℓ a distinct candidate $c_\ell \in W$ such that $c_\ell \in \bigcap_{i \in N_\ell} A(i)$.

While priceability does not imply (nor is implied by) [EJR](#), priceable rules must extend D'Hondt's method [[Peters and Skowron, 2020](#)]. Moreover, fully justified representation implies [EJR](#) [[Peters et al., 2021](#)]. Hence, both are lower quota axioms. On the other hand, perfect representation reduces in apportionment profiles to a weaker, more fundamental proportionality requirement: when we have exactly k equal-sized parties, we must select one candidate for each party. Hence, this is neither a lower nor upper quota axiom, and is in fact incompatible with [EJR](#) [[Sánchez-Fernández et al., 2026](#)].

We now state and prove the incompatibilities.

Proposition 3.24. **JUQ** is incompatible with (i) fully justified representation, (ii) perfect representation, and, if we additionally require exhaustiveness, (iii) priceability.

Proof. (i) Consider the following election with $k = 10$:

Candidate/Voter	{1, 2, 3, 4}	{5, 6, 7, 8}	{9, 10, 11, 12}	13 ... 20
c_1	✓	✓		
c_2		✓	✓	
c_3	✓		✓	
c_4	✓			
c_5		✓		
c_6			✓	
c_7				✓
$\vdots$				...
c_{14}				✓

Voters in $\{1, \dots, 12\}$ are 3-cohesive over $\{c_1, \dots, c_6\}$. Thus, without loss of generality, a committee satisfying fully justified representation must contain c_1, c_3 and c_4 . Then c_4 , with $q_{c_4} = 2$, witnesses a **JUQ** violation. Furthermore, for any committee of size $k \leq 10$ there will always be an unselected candidate $c \in \{c_7, \dots, c_{14}\}$, such that (c_4, c) is a **JUQ** swap.

(ii) Consider the following election with $k = 4$:

Candidate/Voter	1	2	3	4	5	6	7	8	9	10	11	12
c_1	✓	✓	✓	✓	✓	✓						
c_2	✓	✓	✓									
c_3							✓	✓	✓	✓		
c_4							✓	✓			✓	✓
c_5									✓	✓	✓	✓

A committee satisfying perfect representation exists and must contain c_1, c_2 and two of $\{c_3, c_4, c_5\}$. Without loss of generality, consider $\{c_1, c_2, c_3, c_4\}$. Then, (c_2, c_5) is a **JUQ** swap, as $q_{c_2} = 1$ and $\lceil q_{c_5} \rceil = 2$.

(iii) Consider the following election with $k = 3$:

Candidate/Voter	1	2	3	4	5	6
c_1		✓				
c_2			✓			
c_3				✓	✓	
c_4					✓	✓
c_5				✓	✓	✓

Consider any committee W of size $|W| = 3$. We will show that W either violates priceability or **JUQ**.

First, we consider the case that W contains c_5 and at least one of c_3 and c_4 . Without loss of generality, assume $c_4 \in W$. Then, we observe that $\lceil q_{c_4} \rceil = 1$ and thus c_4 together with the unselected candidate among c_1 and c_2 witnesses a **JUQ** violation.

Next, assume that either c_1 or c_2 is contained in W . Then, if W were priceable voters 1 and 2 must have a budget of at least 1. If W additionally contains c_5 , we know by averaging that one of the groups $\{3, 4\}$ or $\{5, 6\}$ must have a leftover budget of $3/2$ therefore violating the fourth condition of priceability together with the candidate they commonly approve. If on the other hand, the committee contains either c_3 or c_4 the other of the two candidates could also be afforded by its supporters.

Hence, W either violates priceability or **JUQ**. □

This concludes our analysis.

3.7 Conclusion

In this chapter, we have studied overrepresentation in **ABC** voting.

Summary. Let us summarise what we learned from this chapter.

- First, we observed that most work on proportionality in **ABC** voting focuses on preventing underrepresentation. The notion of preventing *overrepresentation* has received far less attention, despite also being an appealing notion.
- In Section 3.4, we defined an axiom aimed at preventing overrepresentation, **Justified Upper Quota (JUQ)**. We have shown that, among the class of composite Thiele rules (introduced in Section 3.3), only **Adams-AV** satisfies it. To the best of our knowledge, we are the first to study this rule. We have also studied the swap dynamics of **JUQ**, and showed that exhaustive committees satisfying our axiom can be constructed in polynomial time via a rule we introduced, the **Upper Quota Elimination Rule (UQER)**.
- In Section 3.5, we changed our perspective and focused on the problem of balancing underrepresentation and overrepresentation. Inspired by the near quota notion from the apportionment literature, we introduced an axiom called **Justified Near Quota (JNQ)**. We have shown that among the class of composite Thiele rules, only **SL** satisfies this axiom.
- Finally, in Section 3.6, we discussed the relationship between our notions and established proportionality axioms, mainly **EJR+**. We have shown that **JUQ** is incompatible with some of these axioms (priceability, fully justified representation and perfect representation). Moreover, we have shown that certain natural avenues for proving that **JUQ** and **EJR+** are (in)compatible for exhaustive committees fail.

We note that some of our results, such as the characterisation among composite Thiele rules of **Adams-AV** and **SL** through **JUQ** and **JNQ**, mirror and extend results obtained in the classical apportionment literature [Balinski and Young, 2001].

Future work. We identify several interesting open questions for future work.

- A first task is to close some gaps regarding [JNQ](#). For example, can we find exhaustive committees that satisfy this axiom in polynomial time?
- Next, it would be interesting to characterise the relationship between [JUQ](#) and [JNQ](#) for exhaustive committees and, similarly, the relationship between our axioms and [EJR+](#). In other words, does there exist an exhaustive rule that avoids under- and overrepresentation at the same time? We know from our results that this is not possible within the composite Thiele class and that such a rule must fail the commonly studied axiom of priceability.
- Furthermore, we note that all other established voting rules we examined fail both our axioms (see Section [A.2](#) in appendix). This suggests that our axioms embody a rather different approach to voting rule design compared to the existing literature. Therefore, we believe that finding natural and normatively-appealing rules that satisfy our axioms could lead to fundamentally new insights in [ABC](#) voting.
- Moreover, we point out that approval voting is only one possible ballot type in multi-winner voting. It would be interesting to explore overrepresentation for ranked or cardinal preferences or even for subset selection problems beyond voting, such as recently done for clustering by [Jerrett and Anshelevich \[2025\]](#).
- Finally, composite Thiele rules might be interesting in themselves beyond the scope of overrepresentation. One natural question is the following. In the context of single-winner voting with ranked preferences, in the celebrated axiomatic characterisation by [Young \[1975\]](#), simple and composite scoring functions are separated by an axiom called *continuity*. Continuity states that large enough majorities should dictate the outcome. Interestingly, [Lackner and Skowron \[2021\]](#) use the same axiom—translated in the context of [ABC](#) voting—to characterise simple Thiele rules. Non-simple Thiele rules, for instance [Adams-AV](#), can fail this axiom.²⁰ It would be interesting to see if, by dropping continuity, the characterisation by [Lackner and Skowron \[2021\]](#) extends to composite Thiele rules. This would give a complete characterisation of [Adams-AV](#) (through [AUQ](#)).

²⁰We give a counterexample without entering in the details of the definition of continuity. Fix $k = 2$. In the single-ballot profile $A = (\{a, b\})$, [Adams-AV](#) selects only committee $\{a, b\}$. Then, continuity prescribes that, for some large enough $\ell \in \mathbb{N}$, [Adams-AV](#) should select the same committee in $A' = (\{a, b\} \times \ell, \{c\})$. However, this rule always selects some committee containing c .

CHAPTER 4

Manipulation

But in fact, hundreds, thousands, or even millions of other citizens are also eligible to vote; so each man's ballot is only one drop in a vast sea.

— Anthony Downs [1957]

4.1 Introduction

In the previous chapter, we studied the issue of overrepresentation in parliamentary and committee elections. However, many other settings can be conceptualised as multi-winner elections. One example is repeated single-winner voting, a setting originally proposed by Gottlob Frege (see the historical overview in Chapter 2). These can be viewed as a form of multi-winner elections, where the “multiple winners” are the members of the temporal sequence of selected candidates. This perspective—which is conceptually and, as we will see in Section 4.3.2, formally related to ABC voting—shifts the focus from representation *within a body* to representation *over time*. In particular, this chapter investigates the interplay between temporal fairness guarantees and the strategic incentives induced by such mechanisms.

Frege's idea concerned single-member constituencies (a kind of single-winner election) [Harrenstein et al., 2020]. He argued that a candidate's votes should be carried over from election to election, and when a candidate is elected, some of their cumulative votes should be subtracted. In principle, this should allow minorities to be represented eventually, since their cumulative votes can, at some point, outweigh the majority. Indeed, Harrenstein et al. [2020] analysed this system with the tools of modern social choice theory and found that, under certain conditions, it *does* provide some timid guarantees of “temporal proportionality”. More importantly, it can be slightly modified (e.g., by

adjusting the “price” of winning an election) to achieve a temporal notion of upper quota proportionality and avoid overrepresenting groups. This is similar in spirit to the axioms discussed in the previous chapter.

We might take Frege’s concept further and reformulate it from a voter-centred perspective rather than a candidate-centred one without losing the underlying idea. For instance, a concrete implementation would be a system that assigns weights to all voters and enforces a “cost” on the weights of the citizens whose supported candidate wins.¹ Then, unsatisfied voters will have a greater influence in the next election. The crucial advantage here is that a voter-centred perspective more readily accommodates situations in which the set of candidates changes drastically from election to election. This setting, where the set of candidates may change across elections, has recently attracted some attention in social choice theory, particularly when combined with approval voting. This model is called *perpetual voting* [Lackner, 2020, Lackner and Maly, 2023]. In this chapter, we will adopt this voter-centred perspective.

Beyond fairness, according to Harrenstein et al. [2020], Frege “made his proposal with the intention of reducing strategic voting”. This term refers to the situation where a voter misrepresents her true preferences to gain an advantage in the outcome. A familiar example is deciding whether to vote for one’s favourite party or for a party with an actual chance of winning (the “wasted vote” dilemma). The latter vote would be a strategic one: for our purposes, it would be a form of manipulation. Frege’s claim, at an intuitive level, checks out: this method should ameliorate the wasted vote problem. This is because, in principle, every wasted ballot increases a voter’s chances of being satisfied in later elections. Hence, voters might feel less pressure to misrepresent their true preferences. Indeed, we can think of repeated elections as a system to compensate citizens for their votes being momentarily wasted.

Frege’s intentions regarding strategic voting notwithstanding, these fairness considerations open the door to a straightforward way to manipulate the election. Suppose that we are using the weighted voter-centred method outlined above. An election comes along, and there are two options: one that almost everyone, including you, will support, and a rather unpopular one. For instance, this could be the re-election of an incumbent representative who did an outstanding job against a fringe opponent. In this situation, why should you pay for this election with your voting weight? After all, it is clear that the incumbent candidate will win; strategically, you should abstain, or, if that is not possible, even vote for the opponent. Then, the elections still go in your favour, without you paying the cost: you are a *free-rider*.

We define free-riding in sequential elections as untruthfully opposing a popular candidate, who is likely to win, with the objective of obtaining increased consideration in subsequent decisions. Note that free-riding is a notorious problem in economics, studied well beyond our area of interest. It occurs when it is strategically convenient for

¹Note that one family of rules we consider in this chapter, sequential Thiele rules (Section 4.3.4), can be expressed in these terms, with each voter having assigned a weight. We present them slightly differently, but the two formulations are equivalent.

an individual to abstain from contributing to the production of a public good, since other individuals will likely pay the price [Tuck, 2008, Cullity, 2025]. In particular, free-riding has been studied already in the context of non-temporal proportional representation. For example, Schulze [2004] analyses free-riding under STV, in which voters might untruthfully oppose popular candidates they support in order not to “waste” their votes on such clear winners. This type of manipulation, known both as Hylland free-riding (from the work of Hylland [1992]) and as subset-manipulation, has been studied in the context of ABC voting [Peters, 2018, Botan, 2021]. In particular, Peters [2018] has shown that any ABC voting rule that is proportional—in, arguably, the weakest sense of this word—must be susceptible to Hylland free-riding.

We argue that free-riding is especially natural in sequential decision-making because temporal elections have an explicit cost structure: voters are repeatedly asked to “pay” for the representation they receive. If some voters can strategically avoid this cost while still benefiting from the outcome, then the fairness of the rule may be undermined by its strategic incentives. Still, to our knowledge, there has been no systematic study of free-riding in this context. In this chapter, we set out to fill this gap. We establish two main claims. First, free-riding is, in some sense, endemic to sequential elections: we prove that minimal temporal fairness already implies susceptibility to free-riding. Second, as a positive counterpoint to the previous observation, we argue that free-riding is often strategically less attractive than it initially appears. In particular, we show that free-riding might have negative consequences for the manipulator and that, in general, it is nontrivial to predict when this will be the case.

Although our focus is on sequential elections, our formal model naturally extends to scenarios in which several decisions are made in parallel, for example, multiple referenda [Amanatidis et al., 2015]. The concerns about free-riding, as well as many of our results, carry over, and we discuss this extension in Section 4.7.

4.2 Related Work

We begin by surveying the related work.

Other models. Our approach falls within the broad domain of voting in combinatorial domains [Lang and Xia, 2016], which involves voting over complex objects formed by combining simpler elements (in our case, multiple single-winner elections). In contrast to many works in this field (see, e.g., the work of Brams et al. [1997, 1998], Lang and Xia [2009], Ahn and Oliveros [2012], and Conitzer and Xia [2012]), we assume that voters’ preferences are separable (that is, independent) between different decisions. Separability is a very natural assumption in the sequential setting, as voters may be only asked to state their preferences after preceding elections have already been settled.

From a technical standpoint, our work is most closely related to papers on multiple referenda, where several yes-or-no questions have to be settled in parallel. In this setting, Amanatidis et al. [2015] study strategic manipulation and the computational complexity

of OWA voting rules.² In a similar model, Barrot et al. [2017] consider questions of manipulability, such as the relation between the OWA vector and the susceptibility to strategic manipulation. In contrast to this chapter, these two papers do not consider free-riding specifically. We discuss more technical connections between these papers and our work later in the text.

Another related formalism is perpetual voting [Lackner, 2020], which we mentioned in the introduction. It essentially corresponds to our model when voting is done sequentially, but, as mentioned above, we also consider the case where voting is done in parallel. The work of Lackner [2020] and the follow-up paper by Lackner and Maly [2023] do not consider strategic issues. Kozachinskiy et al. [2025] further study this setting, providing bounds on the dissatisfaction of voters. By contrast, Bulteau et al. [2021] study proportional representation in an offline version of perpetual voting, in which all preferences are submitted upfront. In the same model, Elkind et al. [2024a] study proportionality and strategyproofness (the property of never incentivising strategic behaviour) for the utilitarian and egalitarian rules (which we will define later). Elkind et al. [2025] and Chandak et al. [2026] study proportionality in both the sequential and offline settings. Perpetual voting and its offline variant have also been referred to as temporal voting (see the survey by Elkind et al. [2024b]).

A third related formalism is that of public decision-making, which has been studied both in the offline [Conitzer et al., 2017] and sequential [Freeman et al., 2017] formulations. As in our model, public decision-making considers several decisions and, for each one, an alternative has to be chosen. This model is more general than ours in that it allows arbitrary cardinal additive utilities (whereas we consider only binary utilities, i.e., approval ballots).³ Our work differs in that Conitzer et al. [2017] and Freeman et al. [2017] focus on fairness properties, whereas our focus is on strategic aspects. Fairness considerations in public decision-making have further been explored by Skowron and Górecki [2022] and Kahana and Hazon [2023].

Next, we discuss two other approaches aimed at empowering minorities in repeated (binary) decision-making, though technically distinct from ours. In storable voting [Casella, 2005], for each issue, voters receive one “voting credit”; at each step, they can choose to spend as many credits as they wish (including zero), while unspent credits are accumulated for future decisions. In principle, with storable voting, minorities can focus their influence on issues that affect them most. However, as in our model, the fairness of storable voting may be challenged by its strategic incentives: for example, Lalley and Weyl [2018] argue that storable voting leads to the voter with the strongest preference determining each decision. This vulnerability is addressed by quadratic voting [Lalley and Weyl, 2018]. This system is similar to storable voting, with the main difference being that the price of casting ℓ votes is ℓ^2 credits. Lalley and Weyl [2018] show that the

²OWA rules in multiple referenda are analogous to OWA rules for ABC voting (Definition 2.3). We will define them in detail for our setting as well (Section 4.3).

³A voter having “cardinal utility” refers to the situation where we can measure with a number the satisfaction this voter enjoys from every possible outcome. “Additive” here means that the utility of multiple outcomes is simply the sum of the utilities of each outcome. Approval ballots can be thought of as a special case of this, where each voter can only report two numbers, 0 (no approval) and 1 (approval).

quadratic cost function is the only cost function to incentivise voters to vote according to their true preference intensities. This allows minorities to exert greater influence on issues they value most without enabling a single highly intense voter to dominate collective decisions.

We finally mention two models that, respectively, generalise and specialise our setting. First, [Masářík et al. \[2024\]](#) recently proposed a generalisation of [ABC](#) voting that includes our framework as a special case; they study proportionality in that setting. Second, [Delemazure et al. \[2023\]](#) study strategyproofness in party-approval elections, a model related to [ABC](#) voting which is a special case of ours (where alternatives and preferences are constant across elections).

Manipulation and free-riding. Focusing on [ABC](#) voting, there is substantial literature on the relationship between fairness (often proportionality) and strategyproofness. For example, see the work of [Peters \[2018\]](#), [Lackner and Skowron \[2018\]](#), [Lee \[2018\]](#), and [Kluiving et al. \[2020\]](#). More specifically, as already mentioned in the introduction, free-riding has been studied in the context of [ABC](#) voting through a type of manipulation called Hylland free-riding, also known as subset-manipulation [[Peters, 2018](#), [Botan, 2021](#)]. Here, a voter manipulates by only reporting a subset of the candidates she truly approves of. We note that, while this notion is related to ours in spirit, it is technically distinct—even after accounting for the fact that these two notions are defined for different models. First, in subset-manipulation, there is no requirement for any of the candidates that the manipulator untruthfully opposes to be winners of the original election, in contrast to our definition of free-riding. Second, we do not restrict our manipulators to report a subset of their truthful approvals: in a given election, the manipulator can submit any ballot that does not contain the original winner.

Note that subset-manipulation has also been studied in the setting of fair mixing [[Brandl et al., 2021](#)], a generalisation of multi-winner voting where the output is a probability distribution over candidates. A related notion of “excludable strategyproofness” has been studied in fair mixing by [Aziz et al. \[2019\]](#) and, for a more general model, [Bei et al. \[2025\]](#). This corresponds to a scenario where a voting rule is manipulable *only* by subset-manipulation, which can be desirable in certain settings.⁴

To conclude, as mentioned before, free-riding is a very general phenomenon and has been widely studied in the economic literature on public goods [[Samuelson, 1954](#), [Olson, 1965](#), [Groves and Ledyard, 1977](#)] and philosophy [[Tuck, 2008](#)]. It has also been considered in more technical domains, such as free-riding in memory sharing [[Friedman et al., 2019](#)] and in peer-to-peer networks [[Adar and Huberman, 2000](#)].

⁴The intuition here is that, sometimes, a voter might benefit from the election of an alternative only if this voter publicly supported (i.e., approved) this alternative. For example, think of voting on timeslots for a group activity. If someone pretends to be unavailable in a certain timeslot (that is, she attempts subset-manipulation), then it might be socially dangerous for her to participate in the activity in that timeslot—her manipulative behaviour might be revealed. In this context, excludable strategyproofness would suffice.

4.3 Preliminaries

In this section, we introduce our model, the voting rules, and our formal definition of free-riding. We will also explore the formal connection between our model and ABC voting.

4.3.1 Multi-Issue Decisions

We study a scenario where several single-winner elections, which we call “issues”, have to be settled. We assume that, for each issue, voters cast approval ballots. We will define our model in such a way to accommodate both the case of sequential and parallel decisions, though we focus on the former case. Whether a multi-issue election is sequential or parallel is a feature of the voting rule; this will be clarified later.

Formally, k denotes the number of issues and $C = (C_1, \dots, C_k)$ the respective nonempty sets of candidates. We will typically index issues with variable $\ell \in [k]$. Note that the finiteness of k implies that we do not capture sequential elections of infinite length. Let $N = [n]$ denote the set of voters, which is fixed across all issues. The profile of preferences is a vector of form

$$A = (A_1, \dots, A_k)$$

where each A_ℓ is an approval profile over C_ℓ , that is,

$$A_\ell: N \rightarrow 2^{C_\ell} \setminus \{\emptyset\}$$

for all $\ell \in [k]$. The nonempty set $A_\ell(i) \subseteq C_\ell$ denotes the approval ballot of voter i concerning issue ℓ . Combining all the above, we call the triple $\mathcal{E} = (C, N, A)$ a *multi-issue election*. For the sake of succinctness, we might just call such an object an election, dropping the “multi-issue” qualifier.

To capture sequential elections, we need a formal object representing the result of an election up to a certain issue. A *partial outcome* of length $\ell \in [k]$ is a vector

$$\mathbf{o} = (o_1, \dots, o_\ell) \in C_1 \times \dots \times C_\ell.$$

If $\ell = k$, then we say that $\mathbf{o}$ is a *complete outcome*. The generic term “outcome” might refer to either a partial or a complete outcome. Context should clarify this.

Given an election $\mathcal{E}$ and an outcome $\mathbf{o}$ of length $\ell \in [k]$, the satisfaction of voter $i \in N$ with $\mathbf{o}$ indicates the number of issues decided in i ’s favour:

$$\text{sat}_{\mathcal{E}}(i, \mathbf{o}) = |\{s \in [\ell]: o_s \in A_s(i)\}|.$$

If the election $\mathcal{E}$ is clear from the context, we omit it in the notation and just write $\text{sat}(i, \mathbf{o})$.

A *multi-issue voting rule* F (from now on, called simply a *voting rule*) maps elections to complete outcomes, that is,

$$F(\mathcal{E}) = (o_1, \dots, o_k)$$

for every election $\mathcal{E}$. Note that, in contrast to the previous chapters, we are restricting ourselves to *resolute* voting rules, that is, voting rules that return a single outcome. In practice, such rules typically need to be paired with a tiebreaking mechanism. We make this restriction for two reasons: first, for mathematical convenience. Secondly, it is much harder to model the behaviour of manipulating voters when rules might return multiple tied outcomes. Indeed, given a voter's approvals, we can compare two outcomes (by the number of issues decided in her favour), but comparing two sets of outcomes is ambiguous.⁵

As mentioned earlier, whether an election is decided sequentially or not is captured, in our model, as a feature of the voting rule. A voting rule F is *sequential* if, given any election $\mathcal{E}$, the ℓ -th of the outcome $F(\mathcal{E}) = \mathbf{o}$ can be computed as a function of the current profile A_ℓ , the previous preferences $(A_1, \dots, A_{\ell-1})$, and the previous decisions $(o_1, \dots, o_{\ell-1})$. That is, there exists a function f such that, for every election $\mathcal{E} = (C, N, A)$ with $F(\mathcal{E}) = (o_1, \dots, o_k)$, satisfies

$$f((A_1, \dots, A_\ell), (o_1, \dots, o_{\ell-1})) = o_\ell.$$

We note that in this chapter we focus on the sequential case, unless explicitly noted. We will treat the non-sequential case in Section 4.7.

A straightforward example of a sequential rule is the one which returns, in each issue, the candidate with most approvals—this corresponds to repeatedly applying single-winner approvals voting. This ignores the outcomes of the previous issues: it can be thought of as a “limiting case” of sequential rules, since it is “sequential” only in a purely formal sense.

4.3.2 Connection to Approval-Based Committee Voting

Before exploring the model further, let us clarify the relationship between this setting and ABC voting. As outlined in the introduction, the main connection between multi-issue and ABC elections is conceptual: they are two distinct ways of electing multiple winners. In one case, we are electing a representative body; in the other, we are making several independent decisions. Nonetheless, the two models are strongly related at a technical level. More precisely, assuming resolute voting rules,⁶ multi-issue voting can be seen as a constrained form of ABC voting—and *vice versa*.

Consider the multi-issue election $\mathcal{E} = (C, N, A)$ with k issues. Assume without loss of generality that all candidate sets in C are pairwise disjoint.⁷ In other words, no issue shares any candidates with any other issue. For concreteness, we might assume that, for

⁵Deducting a preference over 2^S given a preference over S is a problem known as preference lifting [Barberà et al., 2004].

⁶Similar constructions to what we show here can be constructed if we assume irresolute voting rules in both models.

⁷While we can safely make this assumption in this chapter, one could exhibit a similar construction as the one we make here without it. It would just be more tedious to write down.

every issue $\ell \in [k]$, the set of candidates is of the form

$$C_\ell = \{c_{(\ell,1)}, \dots, c_{(\ell,|C_\ell|)}\}.$$

We can encode this as an instance of [ABC](#) voting with target committee size k . The voter set of the latter is still N . The set of candidates becomes

$$C = \bigcup_{\ell \in [k]} C_\ell$$

while the approval ballot for every $i \in N$ becomes

$$A(i) = \bigcup_{\ell \in [k]} A_\ell(i).$$

To model the fact that one candidate must be selected for every issue, we additionally need the constraint that any output committee $W \subseteq C$ has $|W \cap C_\ell| = 1$ for all $\ell \in [k]$.

Let us now look at the opposite direction. Indeed, consider an instance of [ABC](#) voting with voters N , candidates C , and target committee size k . Again, we can build a multi-issue election with k issues and voters N . For every $\ell \in [k]$ we have $C_\ell = C$ and for every voter $i \in N$ we have $A_\ell(i) = A(i)$. Again, we need an additional constraint: for any outcome $\mathbf{o} = (o_1, \dots, o_k)$, we must have $o_\ell \neq o_{\ell'}$ for all distinct $\ell, \ell' \in [k]$. This constraint models the fact that we cannot have repeated elements in the output committee.

Due to this connection (in particular, the first direction), we can (and will) immediately use all the voting rules defined in Chapter 2. The only change required is to apply the constraint outlined above to the output, which is straightforward. In particular, one can check that the sequential rules defined in Section 2.2.7 map to rules that are sequential in the sense outlined in the previous section. For clarity, in the next two sections we shall explicitly rewrite sequential Thiele and [OWA](#) rules in the language of this chapter.

As a final note, observe that, while the [ABC](#) voting axioms defined in Chapter 2 can be, in principle, applied to multi-issue decision-making due to the above connection, they are not meaningful, as they do not account for the output constraint. A more fruitful direction would be to consider a more general model that natively accounts for such constraints and to define meaningful axioms for it. Such direction is explored in the work of [Masařík et al. \[2024\]](#), which we already mentioned. They consider a model of constrained [ABC](#) voting in which voters have approval preferences over a set of candidates, and constraints are specified as an explicit list of subsets of candidates that are allowed as outcomes. This model captures both [ABC](#) and multi-issue elections. [Masařík et al.](#) introduce proportionality axioms for their model which meaningfully translate to these two sub-models. We shall not explore this idea further, however, as our focus is on free-riding.

We will now present the rules. We cover first [OWA](#) rules, as it is more convenient for our presentation.

4.3.3 Sequential OWA Rules

Given the connection between multi-issue and [ABC](#) voting outlined above, Definitions 2.3 and 2.4 work immediately in this setting, save for the additional constraint that in issue ℓ we might only pick a candidate from set C_ℓ . We restate the definition for convenience.

Definition 4.1. Analogously to Definition 2.3, we define the sorted score vector $s(\mathcal{E}, \mathbf{o})$ as the vector $(sat_{\mathcal{E}}(1, \mathbf{o}), \dots, sat_{\mathcal{E}}(n, \mathbf{o}))$ sorted in nondecreasing order. In issue $\ell \in [k]$, given a partial outcome $\mathbf{o}$ for the first $\ell - 1$ issues, the *sequential OWA* rule defined by the family of weight functions $\omega^{(n)}$ selects a candidate in

$$\arg \max_{\mathbf{o} \in C_\ell} \left(s(\mathbf{A}, (\mathbf{o}, o_\ell)) \cdot (\omega^{(n)}(1), \dots, \omega^{(n)}(n)) \right)$$

where $(\mathbf{o}, o_\ell)$ is the ℓ -tuple formed by appending o_ℓ to $\mathbf{o}$. If more than one candidate optimises the above objective, we use a fixed ordering over the candidates to break the tie before proceeding to the next issue.

[OWA](#) voting rules for multi-issue domains were proposed by [Amanatidis et al. \[2015\]](#) in the context of multiple referenda. At the end of this section, we will show an example.

We now name some specific [OWA](#) rules. The “limiting case” rule discussed above, which picks the most-approved candidate in each issue, is defined by $\omega^{(n)} = (1, \dots, 1)$. Note that this is the weight vector of [AV](#) in [ABC](#) voting. We refer to this rule as the *utilitarian* rule. Another notable rule is the *egalitarian* rule, which corresponds to $\omega^{(n)} = (1, 0, \dots, 0)$. This rule selects, at each step, the candidate that maximises the satisfaction of the least-satisfied voter. Note that there are typically many such candidates, and thus this rule is rather indecisive—it requires some serious tiebreaking. One can thus consider a refinement of the egalitarian rule, the *leximin* rule, which—among outcomes that are optimal for the least satisfied voter—maximises the satisfaction of the second-least satisfied voter, then the third-least, etc.⁸

We now introduce the leximin rule formally, and then prove that it belongs to the [OWA](#) class.

Definition 4.2. The *leximin* rule is based on the strict leximin ordering $\mathcal{L}(\mathcal{E})$, defined

⁸Note that this terminology—utilitarian, egalitarian, and leximin—is standard in social choice theory. For examples of the terms “utilitarian” and “egalitarian” in the context of [OWA](#) rules, see the work of [Amanatidis et al. \[2015\]](#) and [Elkind and Ismaili \[2015\]](#). The term “leximin” is based on the welfare-aggregation operator based on lexicographic comparisons introduced by John Rawls [[Rawls, 1971](#), [Sen, 2017](#)], already mentioned in Chapter 3.

as follows. Given two outcomes $\mathbf{o}$ and $\mathbf{o}'$ with

$$\mathbf{s}(\mathcal{E}, \mathbf{o}) = (s_1, \dots, s_n) \quad \text{and} \quad \mathbf{s}(\mathcal{E}, \mathbf{o}') = (s'_1, \dots, s'_n),$$

$\mathbf{o}$ is preferred by $\mathcal{L}(\mathcal{E})$ over $\mathbf{o}'$ if and only if there exists an index $i \in [n]$ such that

$$s_1 = s'_1, \quad \dots, \quad s_{i-1} = s'_{i-1} \quad \text{and} \quad s_i > s'_i.$$

Let $\mathbf{o} = (o_1, \dots, o_{\ell-1})$ be the partial outcome selected thus far. In issue $\ell \in [k]$, the sequential leximin rule selects a candidate that together with $\mathbf{o}$ is maximal with respect to $\mathcal{L}(\mathcal{E})$. Precisely, it selects a candidate belonging to the set

$$\{o_\ell \in C_\ell : \text{there is no } o \in C_\ell \text{ such that } (\mathbf{o}, o) \text{ is preferred by } \mathcal{L}(\mathcal{E}) \text{ over } (\mathbf{o}, o_\ell)\}.$$

If more than one candidate belongs to the above set, we use a fixed ordering over the candidates to break the tie before proceeding to the next issue.

We show that this rule corresponds to the **OWA** rule defined by the vector $\omega^{(n)} = (1, 1/kn, 1/k^2n^2, \dots)$.

Proposition 4.1. The **OWA** rule defined by

$$\omega^{(n)} = \left(1, \frac{1}{kn}, \frac{1}{k^2n^2}, \dots\right)$$

is equivalent to the leximin rule.

Proof. Assume that $(\mathbf{o}, \mathbf{o}') \in \mathcal{L}(\mathcal{E})$, i.e., for $\mathbf{s}(\mathcal{E}, \mathbf{o}) = (s_1, \dots, s_n)$ and $\mathbf{s}(\mathcal{E}, \mathbf{o}') = (s'_1, \dots, s'_n)$, there exists an index $i \in [n]$ such that $s_1 = s'_1, \dots, s_{i-1} = s'_{i-1}$ and $s_i > s'_i$. Then

$$\begin{aligned} & \left(\omega^{(n)}(1), \dots, \omega^{(n)}(n) \right) \cdot \mathbf{s}(\mathcal{E}, \mathbf{o}) - \left(\omega^{(n)}(1), \dots, \omega^{(n)}(n) \right) \cdot \mathbf{s}(\mathcal{E}, \mathbf{o}') = \\ & \underbrace{(s_i - s'_i)}_{\geq 1} \cdot \frac{1}{(kn)^{i-1}} + \sum_{j=i+1}^n \underbrace{(s_j - s'_j)}_{\geq -k} \cdot \frac{1}{(kn)^{j-1}} \geq \frac{1}{(kn)^{i-1}} - k \sum_{j=i+1}^n \frac{1}{(kn)^{j-1}} \geq \\ & \frac{1}{(kn)^{i-1}} - k(n-1) \frac{1}{(kn)^i} > 0. \end{aligned}$$

This argument is symmetric in $\mathbf{o}$ and $\mathbf{o}'$. Hence, we have shown that $(\mathbf{o}, \mathbf{o}') \in \mathcal{L}(\mathcal{E})$ if and only if $\mathbf{o}$ has a higher **OWA** score than $\mathbf{o}'$. Thus, a maximal element with respect to $\mathcal{L}(\mathcal{E})$ achieves a maximum score and *vice versa*. $\square$

Let us conclude with an example.

Example 4.1. Consider an election with 100 voters and 4 issues with the same three candidates, $\{a, b, c\}$. There are 66 voters that approve $\{a\}$ in all issues, 33 voters that approve $\{b\}$ in all issues, and one voter that approves always $\{c\}$. The utilitarian rule selects the outcome $\mathbf{o}^{(1)} = (a, a, a, a)$ as a is the most approved candidate in each issue, achieving a total satisfaction of $\sum_{i \in N} \text{sat}(i, \mathbf{o}^{(1)}) = 4 \cdot 66$. The leximin rule selects $\mathbf{o}^{(2)} = (a, b, c, a)$. Indeed, in the first issue, all alternatives give a satisfaction of 0 to (at least) 34 voters; alternative a is the alternative that maximises the satisfaction of the 35th least-satisfied voter. In the second issue, all alternatives give satisfaction 0 to the least-satisfied voter; alternative b is the only alternative where the second least-satisfied voter has satisfaction of 1, and is hence selected. In the third issue, c is the only alternative giving a satisfaction of at least 1 to everyone. In the last issue it selects a , similarly as the first issue. On the other hand, the egalitarian rule is highly indecisive in this election, as many different outcomes are optimal; depending on the tiebreaking rule, it can return, e.g., $\mathbf{o}^{(1)} = (a, a, a, a)$, $\mathbf{o}^{(3)} = (b, b, b, b)$ or even the questionable $\mathbf{o}^{(4)} = (c, c, c, c)$.

4.3.4 Sequential Thiele Rules

Similarly to [OWA](#) rules, sequential Thiele (Definition 2.2) rules can be adapted from the [ABC](#) voting model to our setting straightforwardly.

Definition 4.3. In issue $\ell \in [k]$, given a partial outcome $\mathbf{o}$ for the first $\ell - 1$ issues, the *sequential Thiele rule* defined by ω selects a candidate in

$$\arg \max_{o_\ell \in C_\ell} \left(\sum_{i \in N} \sum_{\ell=1}^{\text{sat}(i, (\mathbf{o}, o_\ell))} \omega(\ell) \right).$$

If more than one candidate optimises the above objective, we use a fixed ordering over the candidates to break the tie before proceeding to the next issue.

In issue ℓ , a voter who has accumulated satisfaction s contributes a marginal score of $\omega(s + 1)$ to each candidate she approves. Thus, under the sequential ω -Thiele rule, each candidate's score is obtained by summing these marginal contributions over all its supporters, and the candidate with highest score is selected.⁹ We will typically use this marginal-score phrasing in proofs, as it is easier to work with.

The utilitarian rule is, again, defined by $\omega = (1, \dots, 1)$. The egalitarian rule does not appear in this class,¹⁰ whereas, for fixed n and k , leximin is equivalent to the Thiele

⁹If a voter does not approve a candidate, her contribution to the objective remains $\sum_{r \in [s]} \omega(r)$. If she does approve the candidate, this increases by the marginal term $\omega(s + 1)$. Since the common baseline is the same for all candidates, only marginal contributions matter.

¹⁰One might first think that the Thiele rule with $\omega = (1, 0, \dots, 0)$ is the egalitarian rule. It is not. The key difference is that once every voter has satisfaction at least 1, this Thiele rule becomes indifferent among the outcomes of the remaining issues. By contrast, the egalitarian rule always tries to maximise the satisfaction

rule with $\omega(\ell) = 1/(k \cdot n)^{\ell-1}$ (among other quickly decreasing functions). We will also consider the harmonic weight function $\omega(\ell) = 1/\ell$, which resembles the PAV rule in the ABC voting model.

Example 4.2. Continuing Example 4.1, we see that for PAV in the first issue a has maximum score of 66. In the second issue, both partial outcomes (a, a) and (a, b) achieve the same score of $66 + 33 = 99$. The final outcome is either $\mathbf{o}^{(5)} = (a, b, a, a)$ or $\mathbf{o}^{(6)} = (a, a, b, a)$, depending on tiebreaking. Note that PAV is more majoritarian than leximin as it essentially ignores the single voter approving $\{c\}$.

4.3.5 Free-Riding

In this section we provide the formal definition of free-riding we will work with. To set the stage, recall that by free-riding we intuitively mean the situation in which a voter untruthfully opposes a popular candidate to gain an advantage in later issues.

Example 4.3. Consider an election with three voters and two issues. The first issue is uncontroversial: all voters approve candidate a . The second issue is highly controversial and all voters approve different candidates:

$$A_2 = (\{b\}, \{c\}, \{d\}).$$

If we settle the election via sequential PAV (with alphabetical tiebreaking), the outcome is $\mathbf{o} = (a, b)$. This leaves voters 2 and 3 less satisfied than voter 1. Namely, the latter has a final satisfaction of 2, while the other two voters have a final satisfaction of only 1. Both voter 2 and voter 3 can free-ride to improve their satisfaction. Let us suppose voter 2 *does* free-ride. If she changes her ballot on the first issue to, say, $\{b\}$, then the outcome in this issue is still a , because a majority of voters here still supports this candidate. By contrast, in the second issue the winner is now c , as this alternative yields—according to the declared ballots—a total PAV score of $1 + 1 + 1$, while the other alternatives both yield $1 + 1 + 1/2$. As the true preferences of voter 2 are positive towards a in the first issue, this manipulation was successful.

We now state the formal definition. Note that it applies to all voting rules, not only the sequential ones. Moreover, observe that the definition makes no assumption about the final satisfaction of the free-riding voter: free-riding might lead to the same, higher, or lower satisfaction for the manipulator. This allows us to study the potential risks associated with free-riding.

of the worst-off voter. A simple example shows the difference. Consider an election with two issues. In the first issue, the only candidate is a , and all voters approve it. In the second issue, there are two candidates, a and b , and all voters approve only b . With alphabetical tiebreaking (i.e., a is preferred over b), the ω -Thiele rule selects the dubious outcome (a, a) , because all voters have already satisfaction 1 in the first issue and a is preferred over b by the tiebreaking. The egalitarian rule instead selects (a, b) , since in the second issue b increases the satisfaction of all voters, while a does not.

Definition 4.4. Fix a voting rule F and consider two elections $\mathcal{E} = (C, N, A)$ and $\mathcal{E}^* = (C, N, A^*)$, a voter $i \in N$, and a set of issues $L \subseteq [k]$. Suppose that $\mathcal{E}$ and $\mathcal{E}^*$ only differ in the ballots submitted by i in issues L , i.e.,

$$\begin{aligned} A_\ell(j) &= A_\ell^*(j) \text{ for all } j \in N \setminus \{i\} \text{ and } \ell \in [k], \quad \text{and} \\ A_\ell(i) &= A_\ell^*(i) \text{ for all } \ell \in [k] \setminus L. \end{aligned}$$

Consider outcomes

$$F(\mathcal{E}) = (o_1, \dots, o_k) \quad \text{and} \quad F(\mathcal{E}^*) = (o_1^*, \dots, o_k^*).$$

If, for every $\ell \in L$, we have

$$\begin{aligned} o_\ell &= o_\ell^*, \\ o_\ell &\in A_\ell(i), \text{ and} \\ o_\ell &\notin A_\ell^*(i), \end{aligned}$$

then we say that i can *free-ride* in $\mathcal{E}$ via $\mathcal{E}^*$ on L .

In words: a voter can free-ride in $\mathcal{E}$ via $\mathcal{E}^*$ on some issues L if she was approving the (true) winners of these issues, and by changing her ballots on those issues (in a way that does not include the winners) the winners remain the same. Since the previous definition does not require the free-riding voter to increase her satisfaction, we additionally define the following:

Definition 4.5. We say that a voting rule F can be *manipulated by free-riding* if there exists an election $\mathcal{E}$ and a voter i such that i can free-ride in $\mathcal{E}$ via some election $\mathcal{E}^*$ (over some issues) and, additionally, $\text{sat}_\mathcal{E}(i, F(\mathcal{E})) < \text{sat}_\mathcal{E}(i, F(\mathcal{E}^*))$.

To sum up, we say that a voter can free-ride if she can get the outcome of an issue “for free”. She can manipulate by free-riding if this move results in a higher total satisfaction. This distinction will be useful in our analysis.

Sometimes (Section 4.5.2), we will consider a more general notion of manipulation, called *generalised* free-riding, where we lift the constraint that the outcome on the issues where free-riding occurs remains exactly the same. Most of our results hold for both models,¹¹ but in some cases we are able to get stronger results for the generalised case. Since this distinction plays a role only when we study computational complexity, we defer the formal definition to Section 4.5.2.

We have defined our model and main notion of manipulation. We are now ready to start the analysis. Note that, in our study, as usual in the context of strategic voting,¹²

¹¹The only exceptions are Propositions 4.17 and 4.18.

¹²For classic examples, see the work of Gibbard [1973], Satterthwaite [1975], Bartholdi III et al. [1989a]

we will assume that there is a single free-rider with full information—the manipulator knows all the preferences from the onset (even of future issues). We leave the study of coalitional manipulation and imperfect information for future work.

4.4 Possibility and Risk of Free-Riding

In this section, we ask two main questions. The first is: for which voting rules manipulation by free-riding is actually possible? We approach this question from two angles. First, we identify a small set of mild representation axioms and show that any rule satisfying them is susceptible to this type of manipulation. Secondly, we focus on sequential Thiele and OWA rules and show that manipulation by free-riding is possible for essentially all these rules. These two approaches complement each other, as they do not cover the same sets of voting rules.

The second question is: under which conditions free-riding may lead to a decrease in the final satisfaction of the manipulator? We analyse for which rules free-riding entails this possibility, i.e., rules for which free-riding is potentially harmful.

4.4.1 Possibility of Free-Riding

We first identify a multi-issue voting rule where manipulating by free-riding is not possible: the utilitarian rule. Observe that with the utilitarian rule the outcomes for different issues do not influence each other. Thus, artificially decreasing one's satisfaction in some issues brings no advantage with this rule. Therefore, the utilitarian rule cannot be manipulated by free-riding.

Observation 4.2. The utilitarian rule cannot be manipulated by free-riding.

Next, we show that for essentially all other reasonable rules, free-riding is possible.

4.4.1.1 Axiomatic Results

In order to formulate axioms that are as mild as possible, we will focus on a class of elections with a particularly simple structure. We claim that in such elections it is easy to argue about how a rule should behave.

Definition 4.6. For every $k \in \mathbb{N}$, $r \in [k]$ and $n \in \mathbb{N} \setminus [2]$, an r -simple (k, n) -election satisfies the following:

- there are k issues and n voters;

and Moulin [1988]. Coincidentally, Moulin studied manipulation by abstention, which the author compares to free-riding. For a more recent example, see the work of Peters [2018]. The chapter by Conitzer and Walsh [2016] contains additional references, as well as examples of papers dealing with coalitional manipulation and manipulation with imperfect information.

- in each issue there are two candidates, namely, $C_1 = \dots = C_k = \{a, b\}$;
- every voter $i \in N \setminus \{n\}$ approves $\{a\}$ in every issue;
- voter n approves $\{b\}$ in r issues, and $\{a\}$ in the remaining $k - r$.

We illustrate this with an example.

Example 4.4. An example 2-simple $(6, 4)$ -election is the following, depicted in tabular form:

Voter/Issue	1	2	3	4	5	6
1	$\{a\}$	$\{a\}$	$\{a\}$	$\{a\}$	$\{a\}$	$\{a\}$
2	$\{a\}$	$\{a\}$	$\{a\}$	$\{a\}$	$\{a\}$	$\{a\}$
3	$\{a\}$	$\{a\}$	$\{a\}$	$\{a\}$	$\{a\}$	$\{a\}$
4	$\{b\}$	$\{a\}$	$\{a\}$	$\{a\}$	$\{b\}$	$\{a\}$

Observe that there are

$$\frac{6!}{4!2!} = 15$$

possible 2-simple $(6, 4)$ -elections, corresponding to the permutations of the last row.

For every number of issues $k \in \mathbb{N}$ and of voters $n \geq 3$, there is exactly one k -simple (k, n) -election (namely, the election where voter n always approves only b) and exactly k different 1-simple (k, n) -elections.

Building on this definition, we now introduce two axioms for multi-issue voting rules. Both apply to all possible rules, not only the sequential ones.

Note that, given the different model and the restricted scope, we define these axioms with a slightly different language with respect to elsewhere in the thesis: we explicitly list the allowed outcomes for a given class of elections. A rule satisfies the axiom if it selects one of these outcomes.

- **Near-unanimity:** The complete outcome of every 1-simple (k, n) -election must be $(a, \dots, a)$.
- **Incentive for minorities:** For every $n \in \mathbb{N} \setminus [2]$, there must be a number of issues k_n such that the complete outcome of the k_n -simple (k_n, n) -election is a permutation of $(a, \dots, a, b)$.

Importantly, our axioms only concern binary elections, i.e., elections in which every issue has exactly two alternatives. Hence, all our axiomatic results apply in particular to multiple referenda.

Normatively, the first axiom states that, if every issue but one is unanimous, and in this exceptional issue everyone except one voter agrees on the same candidate, then we should elect the majority winner in all issues. The second axiom states that, if preferences are constant across issues and everyone but voter n always votes for a , while voter n always votes for b , then if the number of issues is large enough b should win once.

The two axioms are logically independent; the utilitarian rule satisfies near-unanimity but not incentive for minorities, and *vice versa* for the constant rule always selecting $(b, a, a, \dots)$. Both rules are also immune to manipulation by free-riding.

First, we show that any rule that satisfies these two axioms is susceptible to free-riding.

Theorem 4.3. Any rule that satisfies near-unanimity and incentive for minorities can be manipulated by free-riding.

Proof. Consider a rule F that satisfies the conditions of the theorem. Consider any number of voters $n \in \mathbb{N} \setminus [2]$ and let k_n be the number of issues whose existence is prescribed by incentive for minorities. Focus on the k_n -simple (k_n, n) -election $\mathcal{E}^*$ and let ℓ be the unique issue where b is selected by F on $\mathcal{E}^*$. Here is a visualisation of this scenario:

Voter/Issue	1	2	$\dots$	ℓ	$\dots$	k_n
1	$\{a\}$	$\{a\}$	$\dots$	$\{a\}$	$\dots$	$\{a\}$
2	$\{a\}$	$\{a\}$	$\dots$	$\{a\}$	$\dots$	$\{a\}$
$\vdots$	$\vdots$	$\vdots$		$\vdots$		$\vdots$
$n-1$	$\{a\}$	$\{a\}$	$\dots$	$\{a\}$	$\dots$	$\{a\}$
n	$\{b\}$	$\{b\}$	$\dots$	$\{b\}$	$\dots$	$\{b\}$
outcome	a	a	$\dots$	b	$\dots$	a

Consider now the 1-simple (k_n, n) -election $\mathcal{E}$ where voter n approves of b in issue ℓ . By near-unanimity, the outcome must be a in all issues. Visually:

Voter/Issue	1	2	$\dots$	ℓ	$\dots$	k_n
1	$\{a\}$	$\{a\}$	$\dots$	$\{a\}$	$\dots$	$\{a\}$
2	$\{a\}$	$\{a\}$	$\dots$	$\{a\}$	$\dots$	$\{a\}$
$\vdots$	$\vdots$	$\vdots$		$\vdots$		$\vdots$
$n-1$	$\{a\}$	$\{a\}$	$\dots$	$\{a\}$	$\dots$	$\{a\}$
n	$\{a\}$	$\{a\}$	$\dots$	$\{b\}$	$\dots$	$\{a\}$
outcome	a	a	$\dots$	a	$\dots$	a

But then voter n can manipulate by free-riding in $\mathcal{E}$ via $\mathcal{E}^*$ on issues $[k_n] \setminus \{\ell\}$, completing the proof. $\square$

This shows that, as soon as we want to give *some* (but not all) representation to minorities, manipulation by free-riding becomes unavoidable.

Observe that the proof of Theorem 4.3 requires the manipulator to free-ride in multiple issues. If we additionally enforce the following axioms, then free-riding can occur even if we restrict the manipulator to only free-ride in one issue.

- **Issue-wise unanimity:** For every r -simple (k, n) -election, if in some issue all voters approve of a , then a must be selected in that issue.
- **Monotonicity:** If the complete outcome of the k -simple (k, n) -election contains b exactly once, then for all $r \in [k]$, the complete outcome of any r -simple (k, n) -election contains b at most once.

Note that adding issue-wise unanimity is only a mild requirement, here, as its core idea is already embodied by near-unanimity (although the two are logically independent). The intuition behind the monotonicity is that, for a fixed number of issues, if alternative b wins only once when voter n *always* approves only b , then b should not win multiple times when voter n does *not* always approve b .

Theorem 4.4. Any rule that satisfies near-unanimity, incentive for minorities, issue-wise unanimity, and monotonicity can be manipulated by free-riding on only one issue.

Proof. Consider a rule that satisfies the conditions of the theorem. Consider any $n \in \mathbb{N} \setminus [2]$ and let k_n be the number of issues whose existence is prescribed by incentive for minorities. Let r be the minimal r for which there is an r -simple (k_n, n) -election where b wins exactly once, and let $\mathcal{E}^*$ be such an election. Observe that r must exist (by incentive for minorities), that we must have $r > 1$ (by near-unanimity), and that for all $(r - 1)$ -simple (k_n, n) -elections b is not in the outcome (by monotonicity).

Let ℓ be the issue where b wins in $\mathcal{E}^*$. Note that in this issue voter n must approve of b (by issue-wise unanimity). Next, let ℓ' be some issue distinct from ℓ where voter n approves of b (such a ℓ' exists as $r > 1$). Let $\mathcal{E}$ be the $(r - 1)$ -simple (k_n, n) -election obtained by letting voter n approve of a instead of b in issue ℓ' (all else stays the same).

Observe that voter n can manipulate by free-riding in $\mathcal{E}$ via $\mathcal{E}^*$ on issue ℓ' , completing the proof. $\square$

Note that a weaker formulation of incentive for minorities could require for b to win *at least* once instead of *exactly* once (though, clearly, not in all issues):

- **Weak incentive for minorities:** For every $n \in \mathbb{N} \setminus [2]$, there must be a number of issues k_n such that the complete outcome of the k_n -simple (k_n, n) -election contains both a and b at least once.

We have two arguments in favour of our stronger formulation. First, every sequential rule that satisfies the weaker variant of this axiom must also satisfy our stronger notion.

Observation 4.5. If F is sequential and satisfies weak incentive for minorities, then it satisfies incentive for minorities.

Proof. Fix $n \in \mathbb{N} \setminus [2]$ and let k_n be the number of issues whose existence is prescribed by weak incentive for minorities. Let $\mathcal{E}$ be the k_n -simple (k_n, n) -election and let $\ell \in [k_n]$ be the first issue in which F selects b on $\mathcal{E}$. If we consider the election $\mathcal{E}'$ that consists of first ℓ issues of $\mathcal{E}$, then by sequentiality of F we must have $F(\mathcal{E}') = (a, \dots, a, b)$.

Hence, for an arbitrary n , we have shown that there is a number of issues ℓ such that the outcome of the ℓ -simple (ℓ, n) -election (namely, $\mathcal{E}'$) is $(a, \dots, a, b)$. Thus F satisfies incentive for minorities. $\square$

Secondly, one can interpret the additional requirement of the stronger notion (where b must win exactly once, instead of at least once) as a continuity condition: if given $k - 1$ issues voter n does not deserve any representation (satisfaction 0), then she should not suddenly deserve to decide more than one issue when there is exactly one additional issue. Nevertheless, we can consider this weaker notion and get results analogous to the results obtained above by slightly strengthening the other axioms.

- **Strong near-unanimity:** The complete outcome of any r -simple (k, n) -election can contain b at most $r - 1$ times.
- **Strong monotonicity:** For any $r \in [k]$, the complete outcome of any r -simple (k, n) -election cannot contain b more times than the complete outcome of the k -simple (k, n) -election.

The first axiom states that the special voter n should not dictate the outcome. The second axiom is a mild strengthening of monotonicity.

Theorem 4.6. Any rule that satisfies strong near-unanimity, weak incentive for minorities and issue-wise unanimity can be manipulated by free-riding. If it additionally satisfies strong monotonicity, then it can be manipulated by free-riding on only one issue.

Proof. The proof is similar to the proofs of Theorems 4.3 and 4.4.

Consider a rule that satisfies strong near-unanimity, weak incentive for minorities and issue-wise unanimity. Consider any $n \in \mathbb{N} \setminus [2]$ and let k_n be the number of issues whose existence is prescribed by weak incentive for minorities. Focus on the k -simple (k_n, n) -election $\mathcal{E}^*$ and let L be the set of issues where b is the outcome for $\mathcal{E}^*$. Consider now a $|L|$ -simple (k, n) -election $\mathcal{E}$ where voter n approves of b in every issue in L . By strong near-unanimity, the outcome here must contain b at most $|L| - 1$ times. By issue-wise unanimity, b can win only in the issues in L . But then voter n can manipulate by free-riding in $\mathcal{E}$ via $\mathcal{E}^*$ on issues $[k_n] \setminus L$.

Next, suppose the rule additionally satisfies strong monotonicity. Again, fix $n \in \mathbb{N} \setminus [2]$ and let k_n be the number of issues whose existence is prescribed by weak incentive

for minorities. Let s be the number of times where b wins in the k_n -simple (k_n, n) -election. We know by weak incentive for minorities that $1 \leq s \leq k_n - 1$. Let r be the minimal r for which there is an r -simple (k_n, n) -election where b wins exactly s times, and let $\mathcal{E}^*$ be such an election. Observe that $r > s$. Indeed, we get $r \neq s$ via strong near-unanimity and $r \geq s$ via issue-wise unanimity. Furthermore, for all $(r - 1)$ -simple (k_n, n) -elections, b wins at most $s - 1$ times (by strong monotonicity). Let ℓ be some issue where b wins in $\mathcal{E}^*$. Note that in this issue voter n must approve of b (by issue-wise unanimity). Next, let ℓ' be some issue distinct from ℓ where voter n approves of b and where a is in the outcome (such a ℓ' exists as $r > s$). Let $\mathcal{E}$ be the $(r - 1)$ -simple (k_n, n) -election obtained by letting voter n approve of a instead of b in issue ℓ' (all else stays the same). Observe that voter n can manipulate by free-riding in $\mathcal{E}$ via $\mathcal{E}^*$ on issue ℓ' , completing the proof. $\square$

4.4.1.2 Results about Thiele and OWA Rules

Observe that the previous results do not fully cover the families of rules we have introduced. For example, the sequential egalitarian rule does not satisfy near-unanimity (the outcome might depend on the tiebreaking).¹³ Clearly, one might be concerned by what failing such mild axioms means for a rule; nonetheless, we now present a complementary result that shows that both Thiele and OWA rules can be manipulated by free-riding (excluding the utilitarian rule, as per Observation 4.2).

Theorem 4.7. Every sequential Thiele and sequential OWA rule except the utilitarian rule can be manipulated by free-riding on only one issue.

Note that the following proof does not rely on simple elections as used thus far.

Proof of Theorem 4.7. First, consider a sequential ω -Thiele rule different from the utilitarian rule. Then, there exists a k such that $\omega(k) > \omega(k + 1) \geq 0$. Consider an election with $k + 1$ issues, four voters and with two candidates a and b for all issues. Suppose that for the first k issues all voters only approve candidate a . Moreover, on issue $k + 1$ voters 1 and 2 approve b while voter 3 and 4 approve a . Assume further that a is prioritised over b in the tiebreaking order. The following table summarises the construction:

Voter	Issues $[k]$	Issue $k + 1$
1	$\{a\}$	$\{b\}$
2	$\{a\}$	$\{b\}$
3	$\{a\}$	$\{a\}$
4	$\{a\}$	$\{a\}$

¹³For example, consider the 1-simple $(1, 3)$ -election, which has a single issue. If the tiebreaking prioritises b over a , the outcome is (b) , since for both this outcome and outcome (a) the worst-off voter has satisfaction 0. This violates near-unanimity.

Clearly, a wins in the first k issues. Hence in issue $k + 1$ all voters contribute with a score of $\omega(k + 1)$ to the election of their approved candidates. This means that both candidates have a score of $2\omega(k + 1)$. By tiebreaking, a wins.

We claim that voter 1 can manipulate by changing her vote in one of the first k issues to $\{b\}$. Let ℓ be the issue on which 1 manipulates. Then, in issue ℓ , candidate a has a score of $3\omega(\ell)$ while b has a score of $\omega(\ell)$. Now, $\omega(\ell) \geq \omega(k)$ implies that $\omega(\ell) > 0$. Therefore $3\omega(\ell) > \omega(\ell)$, which means a still wins in issue ℓ . It is clear that a also wins in the remaining issues up to issue k . In issue $k + 1$, a has a score of $2\omega(k + 1)$ while b has a score of $\omega(k + 1) + \omega(k)$. By assumption, this means that b wins on issue $k + 1$. Therefore, voter 1 did manipulate by free-riding.

Next, consider a sequential $\omega^{(n)}$ -OWA rule that is not the utilitarian rule. Then there exists a n for which $\omega^{(n)}$ satisfies $\omega^{(n)}(1) > \omega^{(n)}(n)$. Consider an election with 2 issues and n voters. In both issues there are n candidates, namely $\{a_1, \dots, a_n\}$. In the first issue, voters 1 and 2 approve only a_1 . Every other voter $i \in \{3, \dots, n\}$ approves only a_i . In the second issue voter 1 approves a_1 , voter 2 approves a_2 and all other voters approve both a_1 and a_2 . We assume that candidates with a lower index are prioritised by the tiebreaking, which is applied lexicographically (i.e., component-wise, and the first issue in which they differ determines the preferred outcome). Again, we summarise the construction in tabular form:

Voter	Issue 1	Issue 2
1	$\{a_1\}$	$\{a_1\}$
2	$\{a_1\}$	$\{a_2\}$
$i \in \{3, \dots, n\}$	$\{a_i\}$	$\{a_1, a_2\}$

In the first issue a_1 has to be selected, as no other candidate can have a higher OWA score. In the second issue, outcomes (a_1, a_1) and (a_1, a_2) both lead to the highest possible OWA score. By tiebreaking, (a_1, a_1) wins.

Now, we claim that voter 2 can free-ride by approving a_2 instead of a_1 in the first issue. If she does so, then all candidates are tied in the first issue, hence a_1 still wins by tiebreaking. However, in the second issue outcomes (a_1, a_1) and (a_1, a_2) yield the sorted satisfaction vectors

$$(1, 1, \dots, 1, 1) \quad \text{and} \quad (0, 1, \dots, 1, 2),$$

respectively. Here, $\omega^{(n)}(1) > \omega^{(n)}(n)$ implies that a_2 is selected in the second issue. It follows that voter 2 did manipulate by free-riding. $\square$

Hence, free-riding is essentially unavoidable if we want to guarantee fairer outcomes using Thiele or OWA rules.

Observe that we will discuss non-sequential rules in more detail in Section 4.7.

4.4.2 The Risk of Free-Riding

As we argued in the introduction, free-riding seems to offer a simple manipulative strategy. In a sense, it might seem even safe: if I can increase my voting weight by untruthfully opposing a popular candidate, why should I not do that? In this section we challenge this intuition—this will pave the way for the rest of our analysis.

In particular, we observe that for most sequential voting rules, free-riding may lead to a decrease in overall satisfaction. First, we can show that this holds for all sequential Thiele rules, except the utilitarian rule.

Proposition 4.8. Let ω be a weight function for which there is an $\ell \in \mathbb{N}$ such that $\omega(\ell) > \omega(\ell + 1)$. Then, under the sequential ω -Thiele rule, free-riding can reduce the final satisfaction of the free-riding voter.

Proof. Consider a sequential ω -Thiele rule such that $\omega(\ell) > \omega(\ell + 1)$ and consider the following election with 9 voters, $\ell + 3$ issues, and candidates $\{a, \dots, g\}$ for every issue. We assume alphabetical tiebreaking, i.e., a is preferred over b , and so on. The approvals for each issue are:

Voter/Issue	1, ..., $\ell - 1$	ℓ	$\ell + 1$	$\ell + 2$	$\ell + 3$
1	$\{a\}$	$\{a\}$	$\{b\}$	$\{b\}$	$\{b\}$
2	$\{a\}$	$\{a\}$	$\{a\}$	$\{a\}$	$\{c\}$
3	$\{a\}$	$\{a\}$	$\{c\}$	$\{c\}$	$\{d\}$
4	$\{a\}$	$\{b\}$	$\{b\}$	$\{b\}$	$\{e\}$
5	$\{a\}$	$\{b\}$	$\{b\}$	$\{e\}$	$\{b\}$
6	$\{a\}$	$\{c\}$	$\{a\}$	$\{a\}$	$\{f\}$
7	$\{a\}$	$\{d\}$	$\{a\}$	$\{f\}$	$\{a\}$
8	$\{a\}$	$\{e\}$	$\{a\}$	$\{a, b\}$	$\{a\}$
9	$\{a\}$	$\{f\}$	$\{b\}$	$\{g\}$	$\{g\}$

Then, a is clearly the winner for the first $\ell - 1$ issues. Thus, all voters have a satisfaction of $\ell - 1$ before issue ℓ and a wins on issue ℓ as it has the most supporters. In issue $\ell + 1$, a and b increase the Thiele score by $\omega(\ell + 1) + 3\omega(\ell)$, while c increases the score only by $\omega(\ell + 1)$. By tiebreaking, a wins again. Then, in issue $\ell + 2$, a increases the score by $2\omega(\ell + 1) + \omega(\ell + 2)$, b by $\omega(\ell) + 2\omega(\ell + 1)$, while all other candidates by at most $\omega(\ell)$. As $\omega(\ell) > \omega(\ell + 1) \geq \omega(\ell + 2)$ (together with the tiebreaking rule if $\omega(\ell + 1) = 0$), it follows that b wins in issue $\ell + 2$. Finally, in issue $\ell + 3$, a increases the score by $\omega(\ell + 1) + \omega(\ell + 2)$, b by $\omega(\ell) + \omega(\ell + 2)$ and all other candidates increase the score by at most $\omega(\ell)$. Hence b wins.

Now assume voter 1 changes her preferences and free-rides in issue ℓ by approving, e.g., $\{c\}$. It is straightforward to check that a remains the winner for issue ℓ , but winner in issue $\ell + 1$ changes to b while a now wins for issue $\ell + 2$ and $\ell + 3$. Therefore, 1 now additionally approves of the winner on issue $\ell + 1$ but does not approve the winners of

issues $\ell + 2$ and $\ell + 3$ any more. Hence, free-riding led to a lower satisfaction for the free-riding voter. $\square$

Let us describe informally what “goes wrong” for the manipulator in the above construction. The manipulator free-rides in issue ℓ , causing the next issue to be decided in her favour. However, changing the winner on issue $\ell + 1$ from a to b also changes the satisfaction, and hence the subsequent weights, of the other voters. In particular, several voters who benefit from the selection of b also support b in later issues, so that their increased satisfaction reduces the subsequent (weighted) support for b . Conversely, the voters who would have benefited from the selection of a also support a in later issues, and thus retain a higher weight after a loses in issue $\ell + 1$. This causes a to be selected instead of b in the final two issues. Thus, the immediate benefit from free-riding comes at the cost of losing two approved winners later on.

An analogous result holds for the following large class of sequential OWA rules. It should cover most practically appealing rules.

Proposition 4.9. Consider a sequential $\omega^{(n)}$ -OWA rule such that there exists an $n \geq 8$ for which $\omega^{(n)}$ satisfies $\omega^{(n)}(3) > \omega^{(n)}(n - 2)$. Then, free-riding can reduce the final satisfaction of the free-riding voter.

Proof. Fix $n \geq 8$ and $\omega^{(n)}$ that satisfy the conditions of the proposition. Consider an election with 4 issues and n voters. In each issue the candidate set is a subset of $\{a_1, \dots, a_n\}$. The specific set of candidates for each issue is defined as the candidates that receive at least one approval in that issue. The approvals are given by the following description:

- In issue 1, voters 3, $\dots$, $n - 3$, and n approve $\{a_n\}$. Every other voter i approves $\{a_i\}$.
- In issue 2, voter 1, 2, 3 approve $\{a_n\}$, voters $n - 2$, $n - 1$, n approve $\{a_1\}$. Every other voter i approves a_i .
- In issue 3, voter 1 and voter 4 approve $\{a_4\}$, voter $n - 1$ and n approve $\{a_n\}$, Every other voter i approves $\{a_i\}$.
- In issue 4, voter 2 and 3 approve $\{a_2\}$, voter $n - 2$ and n approve $\{a_n\}$ and every other voter i approves $\{a_i\}$.

We assume that candidates with a higher index are prioritised in the tiebreaking.

Let us determine the outcome of this election. In the first issue, a_n has to be selected, as no other candidate can have a higher OWA score. This leads to sorted satisfaction vector $(0, 0, 0, 0, 1, \dots, 1)$. In the second issue selecting a_1 or a_n both lead to a satisfaction vector of $(0, 0, 1, \dots, 1, 2)$. Selecting any other candidate leads to a satisfaction vector of $(0, 0, 0, 0, 1, \dots, 1, 2)$. Clearly, it is again the case that no candidate can have a higher OWA score than a_n . Hence a_n wins again.

In the third issue, selecting any candidate other than a_4 , a_3 , a_{n-2} or a_n leads to a satisfaction vector of $(0, 0, 1, \dots, 1, 2, 2)$. Selecting a_3 leads to a satisfaction vector of $(0, 0, 1, \dots, 1, 3)$, while selecting a_4 leads to a satisfaction vector of $(0, 0, 1, \dots, 1, 2, 2, 2)$ and a_{n-2} to $(0, 1, \dots, 1, 2)$. Finally, selecting a_n leads to a vector of $(0, 1, \dots, 1, 2, 2)$. Since

$$\omega^{(n)}(2) \geq \omega^{(n)}(3) > \omega^{(n)}(n-2) \geq \omega^{(n)}(n),$$

selecting a_n leads to a higher **OWA** score than selecting a candidate a_i with $i < n-2$. Moreover, the **OWA** score of a_{n-2} cannot be higher than the score of a_n . Hence, a_n wins in issue three.

In the fourth issue, selecting any candidate other than a_2 or a_n leads to a satisfaction vector of $(0, 1, \dots, 1, 2, 2, 2)$. Selecting a_2 leads to a satisfaction vector of $(0, 1, \dots, 1, 2, 2, 3)$, while selecting a_n leads to a satisfaction vector of $(1, \dots, 1, 2, 3)$. As $\omega^{(n)}(1) \geq \omega^{(n)}(3) > \omega^{(n)}(n-2)$ selecting a_n leads to the highest **OWA** score.

Now, we claim that voter n can free-ride by approving any other candidate in the first issue. Indeed, if voter n approves any other candidate it is still the case that no candidate can have a higher **OWA** score than a_n . Let us consider how the other issues change.

In the second issue selecting a_1 leads to a satisfaction vector of $(0, 0, 1, \dots, 1)$. Selecting a_n leads to a satisfaction vector of $(0, 0, 0, 1, \dots, 1, 2)$, while selecting any other candidate leads to a satisfaction vector of $(0, 0, 0, 0, 0, 1, \dots, 1, 2)$. As $\omega^{(n)}(3) > \omega^{(n)}(n)$ we know that the **OWA** score of a_1 is higher than that of a_n which is at least as high as the **OWA** score of every other candidate.

In the third issue, selecting any candidate other than a_4 , a_2 or a_n leads to a satisfaction vector of $(0, 0, 1, \dots, 1, 2)$. Selecting a_2 leads to a satisfaction vector of $(0, 1, \dots, 1)$, while selecting a_4 leads to a satisfaction vector of $(0, 1, \dots, 1, 2)$. Finally, selecting a_n leads to a vector of $(0, 0, 1, \dots, 1, 2, 2)$. As $\omega^{(n)}(2) \geq \omega^{(n)}(3) > \omega^{(n)}(n-1)$ selecting a_4 leads to a higher **OWA** score than selecting a candidate a_i with $i \notin \{2, 4\}$. Moreover, the **OWA** score of a_2 cannot be higher than the score of a_4 . Hence, a_4 wins in issue three.

In the fourth issue, selecting any candidate other than a_2 , a_4 , or a_n leads to a satisfaction vector of $(0, 1, \dots, 1, 2, 2)$. Selecting a_2 leads to a satisfaction vector of $(1, \dots, 1, 2, 2)$, while selecting a_4 leads to a satisfaction vector of $(0, 1, \dots, 1, 3)$ and a_n to a satisfaction vector of $(0, 1, \dots, 1, 2, 2, 2)$. As $\omega^{(n)}(1) > \omega^{(n)}(n-2)$ selecting a_2 leads to the highest **OWA** score.

Therefore, free-riding led to a decrease in the final satisfaction of voter n , completing the proof. $\square$

Note that the conditions of Proposition 4.9 do not include the sequential egalitarian rule, but an equivalent statement holds nonetheless:

Proposition 4.10. Free-riding can decrease the final satisfaction of the free-riding voter under the sequential egalitarian rule.

Proof. Consider an election with 5 issues and 5 voters. In each issue there are 3 candidates a , b and c . We assume that tiebreaking always prefers a . The approval sets are given as follows:

Voter/Issue	1	2	3	4	5
1	$\{b\}$	$\{a\}$	$\{b\}$	$\{b\}$	$\{b\}$
2	$\{b\}$	$\{a\}$	$\{a\}$	$\{b\}$	$\{a\}$
3	$\{b\}$	$\{a\}$	$\{a\}$	$\{c\}$	$\{b\}$
4	$\{a\}$	$\{a\}$	$\{b\}$	$\{a\}$	$\{b\}$
5	$\{a\}$	$\{a\}$	$\{b\}$	$\{a\}$	$\{b\}$

Let us determine the winners under the sequential egalitarian rule. In the first issue, every option leads to a minimal satisfaction of 0. Therefore, a is winning by tiebreaking. Then, in the second issue, a must be winning as it raises the minimal satisfaction to 1. In the third issue, again, no alternative can increase the minimal satisfaction and therefore a wins by tiebreaking. This leads to a situation where every voter except 1 has a satisfaction of 2 while 1 has a satisfaction of 1. Hence, in issue four, b must be the winner, as it increases the minimal satisfaction to 2. Finally, in the fifth issue, electing b leads to a minimal satisfaction of 3, which is better than electing a . We observe that voter 1 has a satisfaction of 3 in the end.

Now, we claim that voter 1 can free-ride on issue two. If voter 1 approves b instead, then all candidates lead to the same minimal satisfaction of 0. Hence, a wins by tiebreaking. This manipulation attempt changes the winners of the subsequent issues as follows: in issue three, b must now be the winner, as it increases the minimal satisfaction to 1. Then, in issue four and five, no candidate increases the minimal satisfaction and hence a wins both issues by tiebreaking. However, this leads to a final satisfaction of 2 for voter 1. This is lower than what she would have received if she did not free-ride. $\square$

So far, we have argued that free-riding becomes possible as soon as we try to guarantee minimal representation in repeated elections. However, we have also shown that free-riding is not always risk-free, at least not when using sequential Thiele or OWA rules. This is encouraging.

In the remainder of the chapter, we will try to better understand the risky nature of free-riding from two perspectives: by studying the complexity of assessing its benefits and by evaluating its associated risk through computer simulations.

4.5 Computational Complexity

In this section, we will study the computational complexity of successful free-riding: that is, increasing one's total satisfaction by free-riding. We will show that it is generally hard to do so, even with full information.

This approach—showing hardness of manipulation—is a standard one in computational social choice. Indeed, as outlined in Chapter 2, it was one of the first contributions of computer scientists to voting theory [Bartholdi III et al., 1989a]. This is usually called a “computational barrier to manipulation” [Brandt et al., 2016]. Still, since the late 1980s, things have progressed: nowadays, efficient SAT or ILP solvers allow to solve very large instances of NP-complete problems in a matter of seconds.¹⁴ Additionally, hardness is a worst-case notion and does not provide information about the average-case complexity of a problem. Hence, the status of hardness results (in particular, NP-hardness) as “barriers” against manipulation might be disputed on these grounds; for example, Conitzer and Walsh [2016, Section 6.5] discuss these issues.

One first counterpoint is that, in many scenarios, we cannot expect voters to have access to, or be willing to use, such powerful solvers. For instance, think of a committee of managers voting each week on a topic—it is not particularly realistic to imagine them using such solvers. A second and more essential point is that the criticism above concerns manipulation in general, not free-riding in particular. Free-riding is attractive precisely because it is a restricted and seemingly simple form of manipulation, and this is why we focus on it here. By showing NP-hardness of free-riding, we show that the usual worst-case complexity barrier already applies in this restricted setting. Thus, even though hardness does not by itself resolve average-case behaviour, it shows that the intuitive simplicity of the manipulation move does not eliminate computational difficulty. In particular, a manipulator who is not able to solve computationally hard problems does not gain any computational advantage by restricting attention to free-riding.

Having addressed these preliminary points, let us start with the technical content of the section.

4.5.1 Complexity of Free-Riding

First, observe that the outcome of a sequential rule is always polynomial-time computable: for every issue, we can just iterate over all the candidates involved in that issue and pick the one maximising the score. Therefore, it is computationally easy to decide if free-riding is possible.

Observation 4.11. For any sequential Thiele and OWA rule, election $\mathcal{E}$, and voter $i \in N$, whether i can free-ride in $\mathcal{E}$ on any number of issues is decidable in polynomial time.

Proof. Observe that in sequential rules, if i can free-ride on issues $L \subseteq [k]$, then i can free-ride on the first issue in L , i.e., $\min(L)$. Hence, we can restrict our attention to single-issue free-riding. First, compute the outcome under the truthful preferences. Then, for each issue $\ell \in [k]$ where i approves of the winner o , iterate over all singleton ballots $A_\ell(i) = \{o'\}$ with $o' \in C_\ell \setminus \{o\}$ and check if under any such ballot o still wins.

¹⁴See, for reference, the “SAT competition”: <https://satcompetition.github.io/> (accessed: 2026-05-08).

Note that restricting oneself to singleton ballots is safe under OWA and Thiele rules—if there exists a ballot $A_\ell^*(i)$ via which voter i can free-ride in issue ℓ , then she can free-ride in the same issue also via any singleton ballot $\{o'\}$ where $o' \in A_\ell^*(i)$. If this is possible for some issue, return “yes”. Otherwise, return “no”. $\square$

However, although voters can easily verify if free-riding is possible, it might be hard to judge its long-term consequences. If this is unfeasible, voters might be discouraged from free-riding because, as argued in Section 4.4.2, it can have negative consequences. Hence, we study the following problem.

Definition 4.7. SUCCESSFUL- F -FREE-RIDING is the following decision problem: “Given an election $\mathcal{E} = (\mathbf{C}, N, \mathbf{A})$ and a voter $i \in N$, is there an election $\mathcal{E}^*$ such that i can free-ride in $\mathcal{E}$ via $\mathcal{E}^*$ and $\text{sat}_{\mathcal{E}}(i, F(\mathcal{E})) < \text{sat}_{\mathcal{E}}(i, F(\mathcal{E}^*))$?”

First, we show that free-riding is NP-complete for a large class of sequential Thiele rules and the sequential egalitarian rule.

Theorem 4.12. SUCCESSFUL- F -FREE-RIDING is NP-complete for every sequential ω -Thiele rule for which there exists $\ell \in \mathbb{N}$ such that ω is constant on $[\ell]$ and strictly decreasing on $\mathbb{N} \setminus [\ell]$.

In other words, ω has a constant prefix and is strictly decreasing thereafter. Notice that being strictly decreasing implies that this function never reaches zero. This condition applies to several appealing rules, including the sequential PAV rule.

Proof of Theorem 4.12. Fix the ω -Thiele rule F satisfying the conditions of the theorem. First, notice that SUCCESSFUL- F -FREE-RIDING is in NP: given a sequence of insincere ballots for the manipulator, we can check whether they improve her final satisfaction and they constitute a case of free-riding in polynomial time.

We show hardness by a reduction from 3-SAT [Garey and Johnson, 1979], the problem of deciding whether a 3-CNF formula is satisfiable. A 3-CNF formula is a CNF formula where each clause has exactly three literals. Let φ be such a formula with n variables and m clauses. In particular, the sets of variables and clauses are

$$\{x_1, \dots, x_n\} \quad \text{and} \quad \{\gamma_1, \dots, \gamma_m\},$$

respectively. Crucially, we assume that φ is not satisfied by setting all variables to false. This is without loss of generality, since we can check in polynomial time whether the all-false assignment satisfies φ and, if so, return a trivial yes-instance of SUCCESSFUL- F -FREE-RIDING.

Furthermore, let $h \in \mathbb{N}$ be the smallest h such that $\omega(h+1) < \omega(h)$. As F is not the utilitarian rule, such a h exists, and is constant in the size of φ (as ω does not depend on the number of voters and issues). Moreover, let $g \in \mathbb{N}$ be the smallest g such that $\omega(h+1)(g+1) < \omega(h)g$. Again, note that g can be large, but does not depend

on the instance. We will construct an instance of SUCCESSFUL- F -FREE-RIDING with $3n(g+1)+5$ voters and $h+3n$ issues.

Let us describe the reduction at a high level. The construction consists of three parts. First, we use an initial block of issues to equalise the satisfaction levels of all voters to $h-1$. Second, for each variable, we create a three-issue gadget in which the manipulator can choose whether or not to free-ride, thereby encoding a truth assignment. Finally, a last issue checks whether the assignment encoded through these free-riding choices satisfies the formula. We will show that the manipulator can improve her satisfaction by free-riding if and only if φ is satisfiable.

For each propositional variable x_i , there are three voters s_i , v_i and $\bar{v}_i$ plus $3g$ voters

$$r_1^x, t_1^x, w_1^x, \dots, r_g^x, t_g^x, \text{ and } w_g^x.$$

Furthermore, we have four additional voters q , u_1 , u_2 , and u_3 . Finally, we have a distinct voter i^* , the free-rider.

In the first $h-1$ issues, there are two candidates, c and $\bar{c}$, and every voter approves both. Until here, i^* cannot free-ride, as whatever she votes for will win. Moreover, all voters win in each issue. Thus, the satisfaction of every voter after the first $h-1$ issues will be exactly $h-1$.

Now, focus on issues from h to $h+3n-1$, that is, all remaining issues excluding the final one. We partition this set of issues into triples. That is, the first triple is $(h, h+1, h+2)$, the second $(h+3, h+4, h+5)$, and so on. We refer to the ℓ' -th entry of ℓ -th triple as issue (ℓ, ℓ') . For example, issue $(2, 3)$ corresponds to issue $h+5$. In the first issue of every triple, there is one candidate c , plus one candidate c_i for every voter $i \in N$. In the second and third issues, we additionally have a voter $\bar{c}$. The tiebreaking order is of form $\bar{c} \succ c \succ \{c_i : i \in N\}$.

For every triple ℓ , in issue $(\ell, 1)$, voters i^* , q and s_ℓ approve of candidate c . Every other voter i approves only of c_i . In issue $(\ell, 2)$, voters in $\{i^*, v_\ell, r_1^\ell, \dots, r_g^\ell\}$ vote for c , while voters in $\{q, \bar{v}_\ell, t_1^\ell, \dots, t_g^\ell\}$ vote $\bar{c}$. Every other voter i approves only of c_i . In issue $(\ell, 3)$, voters in set $\{i^*, r_1^\ell, \dots, r_g^\ell\}$ vote for c , while $\{w_1^\ell, \dots, w_g^\ell\}$ for $\bar{c}$. Again, every other voter i approves only of c_i . We provide a visualisation of this construction:

Voter/Issue	$(\ell, 1)$	$(\ell, 2)$	$(\ell, 3)$
i^*	$\{c\}$	$\{c\}$	$\{c\}$
q	$\{c\}$	$\{\bar{c}\}$	$\{c_q\}$
s_ℓ	$\{c\}$	$\{c_{s_\ell}\}$	$\{c_{s_\ell}\}$
v_ℓ	$\{c_{v_\ell}\}$	$\{c\}$	$\{c_{v_\ell}\}$
$\bar{v}_\ell$	$\{c_{\bar{v}_\ell}\}$	$\{\bar{c}\}$	$\{c_{\bar{v}_\ell}\}$
$i \in \{r_1^\ell, \dots, r_g^\ell\}$	$\{c_i\}$	$\{c\}$	$\{c\}$
$i \in \{t_1^\ell, \dots, t_g^\ell\}$	$\{c_i\}$	$\{\bar{c}\}$	$\{c_i\}$
$i \in \{w_1^\ell, \dots, w_g^\ell\}$	$\{c_i\}$	$\{c_i\}$	$\{\bar{c}\}$
Any other voter i	$\{c_i\}$	$\{c_i\}$	$\{c_i\}$

First, note that no “voter-candidate” c_i is ever selected. We show this by induction. In issue h every voter has satisfaction $h - 1$, and hence c , the most approved candidate, wins. Now, suppose that no voter-candidates wins until issue ℓ . In issue $\ell + 1$, we know that there is at least one voter approving c or $\bar{c}$ that has never been satisfied since issue $h - 1$, as she always voted for her voter-candidate up to issue ℓ . For example, if $\ell + 1$ is the first entry of a triple, then this would be voter s_ℓ . Such voters contribute to the score of c and $\bar{c}$ with a value of $\omega(h)$. Since no voter-candidate can have a score greater than $\omega(h)$, our claim follows.

Consequently, in every issue $(\ell, 1)$, voter i^* can free-ride by voting for her voter-candidate. We claim that if i^* free-rides in $(\ell, 1)$ then she wins in issue $(\ell, 2)$ and loses in issue $(\ell, 3)$; if she does not, the opposite happens. Furthermore, we claim that at the beginning of every triple, i^* and q have the same (declared) satisfaction. We show both claims by induction.

As base case, consider issue $(1, 1)$. If i^* does not free-ride, c wins, and the satisfaction of q and i^* will be h . Now consider issue $(1, 2)$. The score of both c and $\bar{c}$ is $\omega(h + 1) + \omega(h) + \omega(h)g$, and hence $\bar{c}$ wins by tiebreaking. Thus, in the next issue, the score of c will be $\omega(h)g + \omega(h + 1)$, and the score of $\bar{c}$ will be $\omega(h)g$. Hence, c wins. Now, suppose that i^* free-rides. Then, her satisfaction after issue $(1, 1)$ will be $h - 1$. Therefore, the score of c would increase to $2\omega(h) + \omega(h)g$, making it the winner in $(1, 2)$. Hence, in issue $(1, 3)$ the scores of c and $\bar{c}$ would be $\omega(h + 1)(g + 1)$ and $\omega(h)g$, respectively. As we assume $\omega(h + 1)(g + 1) < \omega(h)g$, here we have that $\bar{c}$ wins, as desired. Observe that in both cases i^* and q won the same number of issues (according to the ballot reported by i^*).

Now suppose the claim holds up to triple ℓ . Then, let α be the (declared) satisfaction of i^* and q at the beginning of triple $\ell + 1$. Observe that if i^* does not free-ride in $(\ell + 1, 1)$, then in issue $(\ell + 1, 2)$ the score of c and $\bar{c}$ is $\omega(\alpha + 2) + \omega(h) + \omega(h)g$, and hence $\bar{c}$ wins by tiebreaking. Thus, in issue $(\ell + 1, 3)$, the score of c will be $\omega(h)g + \omega(\alpha + 2)$, and of the score of $\bar{c}$ will be $\omega(h)g$. Therefore, c wins. If i^* does free-ride, then in issue $(\ell + 1, 2)$ the score of c raises to $\omega(\alpha + 1) + \omega(h) + \omega(h)g$, making it the winner. Furthermore, in issue $(\ell + 1, 3)$, the scores of c and $\bar{c}$ would be $\omega(h + 1)g + \omega(\alpha + 2)$ and $\omega(h)g$, respectively. Clearly, $\bar{c}$ wins here. Observe again that i^* and q won the same number of issues (according to the ballot reported by i^*).

This concludes the inductive argument. We have shown, in particular, that before the last issue, i^* and q have the same (declared) satisfaction.

Another crucial aspect is that i^* can free-ride in the first issue of every triple, but not elsewhere. Fix a triple ℓ . In issue $(\ell, 2)$, free-riding is never an option as i^* does not approve the truthful winner of the election. Free-riding only in issue $(\ell, 3)$ is not possible, as in this case the score contributed by i^* breaks the tie, and her not approving of c would change the true outcome of the election. Finally, since free-riding on issue $(\ell, 1)$ causes $\bar{c}$ to win in issue $(\ell, 3)$, free-riding on both issues $(\ell, 1)$ and $(\ell, 3)$ is not an option, as i^* truly approves of c in the latter.

Let us describe the final issue. Here, there are $m + 1$ candidates, namely

$\{c, c_1, \dots, c_m\}$. Here, each voter v_i votes for c_j if $x_i \in \gamma_j$ (and similarly for $\bar{v}_i$ and $\neg x_i$). Furthermore, i^* , u_1 , u_2 , and u_3 all approve c . Voter q approves all candidates except for c . Everyone else votes for all candidates. If there is a tie, now c loses against every other candidate.

We can interpret c (respectively, $\bar{c}$) winning in issue $(\ell, 2)$ as setting x_ℓ to true (respectively, false). Observe that the inductive argument above shows that c wins in issue $(\ell, 2)$ if and only if voter i^* free-rides in issue $(\ell, 1)$. Hence, we can in turn interpret the fact that i^* free-rides in issue $(\ell, 1)$ as setting x_ℓ to true, and *vice versa*. Observe that we have already shown that i^* cannot free-ride anywhere else. We claim that this assignment satisfies φ if and only if c wins in the final issue.

Let α be the score contributed in the last issue by the voters who vote for all candidates. Furthermore, let β be the score contributed by i^* or by q , which are the same (as shown in the inductive argument above). Observe that u_1 , u_2 , and u_3 won exactly $h - 1$ issues. Furthermore, v_i won exactly h issues if c won in issue $(i, 2)$, and $h - 1$ otherwise (and conversely for $\bar{v}_i$ and $\bar{c}$). Thus, the score of c in the final issue will be $\alpha + \beta + 3\omega(h)$. Furthermore, given a clause γ_j , if all of its literals are unsatisfied, the score of c_j will also be $\alpha + \beta + 3\omega(h)$. By our rule on tiebreaking, such a c_j blocks the victory of c . If, on the other hand, some literals in γ_j are satisfied, the score of c_j will be at most $\alpha + \beta + 2\omega(h) + \omega(h + 1)$. Hence, if all clauses are satisfied, c wins, as desired.

We finally claim that the only way for i^* to raise her (true) satisfaction is by making c win in the last issue by forcing a satisfying assignment for φ by free-riding. Indeed, if i^* votes truthfully, then the induced assignment is the all-false assignment, which by assumption does not satisfy φ , making c lose in the last issue. Moreover, as shown earlier, i^* will be truthfully satisfied with the $h - 1$ first issues, plus two issues per triple (irrespective of whether she free-rides or not).

It follows that i^* can increase her satisfaction by free-riding if and only if φ is satisfiable, and we are done. $\square$

We further consider the egalitarian rule.

Theorem 4.13. SUCCESSFUL- F -FREE-RIDING is NP-complete for the sequential egalitarian rule.

Proof. First, note that, for the egalitarian rule, SUCCESSFUL- F -FREE-RIDING is in NP. Given an insincere ballot sequence for the manipulator, we can check whether they improve her final satisfaction and they constitute a case of free-riding in polynomial time.

We again show hardness by a reduction from 3-SAT. Let φ be a 3-CNF with n variables and m clauses. We use the same notation as in the proof of Theorem 4.12. In particular, we refer to the j -th clause as γ_j . We again assume that φ is not satisfied by setting all variables to false. We construct an instance of SUCCESSFUL- F -FREE-RIDING with $2(n + 1)$ voters and $5n + 1$ issues. In particular, we will have two voters v_i and $\bar{v}_i$

for each variable x_i , a voter u , and a distinguished voter i^* , the manipulator.

Let us start with the first $4n$ issues. We subdivide this set of issues into quadruples; that is, the first quadruple consists of issues 1, 2, 3 and 4, the second are the issues 5, 6, 7 and 8, and so on. We refer to the ℓ' -th entry of quadruple ℓ as issue (ℓ, ℓ') . For example, issue $(2, 3)$ corresponds to issue 7. In each issue of every quadruple, there are two candidates c and $\bar{c}$. This notation is analogous to that of the proof of Theorem 4.12. Here (and in all subsequent issues), we assume that if there is a tie then c loses.

Consider a generic quadruple ℓ . In issue $(\ell, 1)$, all voters vote for $\bar{c}$. In issue $(\ell, 2)$, voters i^* and $\bar{v}_\ell$ vote for c , while voters u and v_ℓ vote for $\bar{c}$. Everyone else approves of both. Furthermore, in issue $(\ell, 3)$, voters i^* and u approve of c and $\bar{c}$, respectively; everyone else approves of both. Finally, in issue $(\ell, 4)$, voter i^* approves of both c and $\bar{c}$, while everyone else approves only of c . We summarise the construction visually:

Voter/Issue	$(\ell, 1)$	$(\ell, 2)$	$(\ell, 3)$	$(\ell, 4)$
i^*	$\{\bar{c}\}$	$\{c\}$	$\{c\}$	$\{c, \bar{c}\}$
v_ℓ	$\{\bar{c}\}$	$\{\bar{c}\}$	$\{c, \bar{c}\}$	$\{c\}$
$\bar{v}_\ell$	$\{\bar{c}\}$	$\{c\}$	$\{c, \bar{c}\}$	$\{c\}$
u	$\{\bar{c}\}$	$\{\bar{c}\}$	$\{\bar{c}\}$	$\{c\}$
Any other voter i	$\{\bar{c}\}$	$\{c, \bar{c}\}$	$\{c, \bar{c}\}$	$\{c\}$

For each quadruple ℓ , we claim that (1) i^* can free-ride (only) in issue $(\ell, 1)$, that (2) if i^* does not free-ride, the winners in this quadruple are $(\bar{c}, \bar{c}, c, c)$, and that (3) if i^* free-rides, the winners in this quadruple are $(\bar{c}, c, \bar{c}, \bar{c})$. We show so by induction over the quadruples.

Consider quadruple 1. Observe that, by tiebreaking, $\bar{c}$ wins in issue $(1, 1)$ (irrespective of what i^* votes for). Hence, here i^* can free-ride. Suppose that she votes truthfully. Thus, in the next issue, the minimal satisfaction if c or $\bar{c}$ win is the same (namely, 1). By tiebreaking, $\bar{c}$ wins. Now, up to here, every voter has satisfaction 2, save for i^* and $\bar{v}_1$, who have satisfaction 1. In the next issue, then, the minimal satisfaction if c or $\bar{c}$ win is 2 and 1, respectively; hence, c wins. Finally, in issue $(1, 4)$, the minimal satisfaction of c winning is 3, whereas the minimal satisfaction in case $\bar{c}$ wins is 2 (namely, of voter u); hence, c wins. With a similar line of reasoning, one can show that $(\bar{c}, c, \bar{c}, \bar{c})$ is the result if i^* does free-ride. Observe that i^* can indeed free-ride only in issue $(1, 1)$: in every other issue, either she is losing, or her vote would change the outcome.

Now, suppose this property holds up to quadruple ℓ , and consider quadruple $\ell + 1$. If this holds, observe that no voter v_j or $\bar{v}_j$ can have won fewer issues than i^* (according to her ballot) or u , and i^* and u won the same number of issues (again, according to what i^* declares). Therefore, at the beginning of each quadruple, i^* and u are among the voters with the lowest (declared) satisfaction. Let this minimal satisfaction be s .

Again, observe that i^* can free-ride in issue $(\ell + 1, 1)$. Suppose she votes truthfully. Then, she and u will have the same satisfaction of $s + 1$ in issue $(\ell + 1, 2)$, and $\bar{c}$ will win by tiebreaking. Next, in issue $(\ell + 1, 3)$, i^* will have the minimal satisfaction of $s + 1$,

and hence c will win. Finally, in issue $(\ell + 1, 4)$, if c wins the minimal satisfaction will be $s + 3$, whereas if $\bar{c}$ wins it will be $s + 2$ (namely, of u); hence c wins. With similar arguments, we could show that $(\bar{c}, c, \bar{c}, \bar{c})$ is the result if i^* does free-ride. Observe that i^* can indeed free-ride only in issue $(\ell + 1, 1)$: in every other issue, either she is losing, or her vote would change the outcome. This concludes the inductive argument.

From the previous discussion we know that, in each quadruple ℓ , all voters v_j and $\bar{v}_j$ (with $j \neq \ell$) win the same amount of issues (either 3 or 4, depending on whether i^* free-rides or not). Let this number be $r(\ell)$. Furthermore, one voter in v_ℓ and $\bar{v}_\ell$ wins $r(\ell)$ issues, while the other wins $r(\ell) - 1$ issues. Finally, both i^* and u win exactly $r(\ell) - 1$ issues. Thus, we can partition the set $\{v_\ell: \ell \in [n]\} \cup \{\bar{v}_\ell: \ell \in [n]\}$ into two groups according to their satisfaction, with satisfaction differing of exactly 1 point between the two. Let s be the satisfaction of the least satisfied voters in $\{v_\ell: \ell \in [n]\} \cup \{\bar{v}_\ell: \ell \in [n]\}$. Observe that the (declared) satisfaction of i^* and u will be exactly $s + 1 - n$, because in each quadruple ℓ , both i^* and u lose exactly one issue compared to the voters v_j and $\bar{v}_j$ with $\ell \neq j$.

So then, let us consider issue $4n + 1$ to issue $5n - 1$. There are two candidates: c and $\bar{c}$. Here, i^* approves of $\bar{c}$, u of both candidates, and everyone else only of c . Observe that i^* and u win every such issue, as i^* always has a strictly lower satisfaction than the rest of the voters, and u always approves of all candidates. Furthermore, i^* cannot free-ride here: as she always has the lowest satisfaction, she is always pivotal, and hence her vote decides the outcome. Thus, after these issues, both i^* and u will have satisfaction s .

In issue $5n$, there are two candidates: $\bar{c}$ and c . Here, i^* votes for $\bar{c}$, whereas every other voter votes for c . If either candidate wins, the minimal satisfaction will be s , and thus $\bar{c}$ wins by tiebreaking. Furthermore, i^* cannot free-ride: if she were to vote for c , then c would win.

Finally, in issue $5n + 1$, we know that i^* has satisfaction of $s + 1$ and u of s . Furthermore, v_ℓ has satisfaction $s + 1$ if c won in issue $(\ell, 2)$ and s otherwise (and conversely for $\bar{v}_\ell$ and $\bar{c}$). In this issue, there are $m + 1$ candidates, namely $\{c, c_1, \dots, c_m\}$. Here, u approves of all candidates, voter v_ℓ (respectively, $\bar{v}_\ell$) approves of c and of all candidates c_j such that $x_\ell \notin \gamma_j$ (respectively, $\bar{x}_\ell \notin \gamma_j$). Finally, voter i^* approves of c .

Observe that we can interpret c winning in issue $(\ell, 2)$ as setting x_ℓ to true (and conversely for $\bar{c}$); we claim that this assignment satisfies φ if and only if c wins in this final issue. To see this, observe that if c wins the minimal satisfaction will be $s + 1$ (all voters approve of c). Now, consider a clause γ_j and its candidate c_j . If all literals in γ_j are unsatisfied, then the corresponding voters have all (up to this issue) satisfaction $s + 1$. Hence, if c_j would win, the minimal satisfaction will be at least $s + 1$, as all other voters (except i^*) approve of it, and i^* has satisfaction at least $s + 1$. Thus, c_j would win by tiebreaking. Conversely, if at least one literal in γ_j is satisfied, there is at least one voter with satisfaction s that does not vote for c_j : hence, c_j loses against c . Our claim follows.

To conclude, observe that c wins in $(\ell, 2)$ if and only if i^* free-rides in $(\ell, 1)$. This corresponds to setting x_ℓ to true, and hence a free-riding pattern for i^* maps to a truth

assignment. Now, observe that, if i^* were to always vote truthfully, her true satisfaction would be $4n$ (she would win three issues per quadruple, all issues from $4n + 1$ to $5n$, and lose the last issue, by the assumption that setting all variables to false does not satisfy φ). Observe also that, as we discussed before, she can only free-ride in the first issue of every quadruple. Therefore, the only way she can raise her satisfaction to $4n + 1$ is by winning the last issue (observe that if she free-rides in some quadruple, she still truly wins three issues). To do so, she has to force a satisfying assignment for φ by free-riding. It follows that i^* can manipulate by free-riding if and only if φ is satisfiable, and so we are done. $\square$

4.5.2 A Weaker Notion of Free-Riding

In this section, we study the complexity of a weaker, that is, more permissive, notion of manipulation. We call this *generalised free-riding*. This allows us to obtain results for a broader class of rules, in particular for the OWA family. Under this definition, a voter can free-ride if she can change her ballot so that the winning alternative is still one she truthfully approves of, even if this alternative differs from the original winner. Thus, we only require that the new winner remain in the manipulator's true approval set. At its core, generalised free-riding rests on the assumption that, for a given issue, voters are *truly* indifferent among the candidates they approve. That is, they are not merely listing all candidates they consider acceptable, possibly with different degrees of preference, but are instead reporting a set of candidates that are indistinguishable in terms of the "utility" they provide.

To formally define generalised free-riding, we replace " $o_\ell = o_\ell^*$ " in Definition 4.4 with " $o_\ell^* \in A_\ell(i)$ ". We report the full definition for convenience.

Definition 4.8. Fix a voting rule F and consider two elections $\mathcal{E} = (C, N, A)$ and $\mathcal{E}^* = (C, N, A^*)$, a voter $i \in N$, and a set of issues $L \subseteq [k]$. Suppose that $\mathcal{E}$ and $\mathcal{E}^*$ only differ in the ballots submitted by i in issues L , i.e.,

$$\begin{aligned} A_\ell(j) &= A_\ell^*(j) \text{ for all } j \in N \setminus \{i\} \text{ and } \ell \in [k], \quad \text{and} \\ A_\ell(i) &= A_\ell^*(i) \text{ for all } \ell \in [k] \setminus L. \end{aligned}$$

Consider outcomes

$$F(\mathcal{E}) = (o_1, \dots, o_k) \quad \text{and} \quad F(\mathcal{E}^*) = (o_1^*, \dots, o_k^*).$$

If, for every $\ell \in L$, we have

$$\begin{aligned} o_\ell^* &\in A_\ell(i), \\ o_\ell &\in A_\ell(i), \text{ and} \\ o_\ell &\notin A_\ell^*(i), \end{aligned}$$

then we say that i can perform *generalised free-riding* in $\mathcal{E}$ via $\mathcal{E}^*$ on L .

The definition of manipulation by generalised free-riding is analogous to the definition of manipulation by free-riding (Definition 4.5). Note that, if $|A_\ell(i)| = 1$ for all $\ell \in L$, then free-riding and generalised free-riding coincide. Furthermore, whether generalised free-riding is possible is still decidable in polynomial time, and this can be shown by arguments analogous to Observation 4.11. Finally, generalised free-riding can too lower the satisfaction of the manipulator: all results from Section 4.4 carry over.

We remark that the definition of generalised free-riding is *considerably* more permissive than the main definition (Definition 4.4), as we allow o_ℓ^* to belong to $A_\ell^*(i)$. Hence, generalised free-riding encompasses situations where the manipulator does not artificially decrease her satisfaction; that is why we do not adopt this as our main definition of free-riding. Nevertheless, we will briefly focus on this notion, as it is conceptually related to ours—it captures all scenarios where a voter attempts to manipulate the election by opposing a popular candidate she truly approves of—and it allows us to capture broader classes of rules.

The problem of SUCCESSFUL- F -GENERALISED-FREE-RIDING is then defined analogously to SUCCESSFUL- F -FREE-RIDING.

Theorem 4.14. SUCCESSFUL- F -GENERALISED-FREE-RIDING is NP-complete for every sequential ω -Thiele rule distinct from the utilitarian rule such that $\omega(\ell) > 0$ holds for every $\ell \in \mathbb{N}$.

Note that this properly includes the class covered by Theorem 4.12: it allows for rules that have multiple plateaus, not only on an initial segment of $\mathbb{N}$.

Proof of Theorem 4.14. Fix an ω -Thiele rule F satisfying the conditions of the theorem. First, notice that SUCCESSFUL- F -GENERALISED-FREE-RIDING is in NP, as we can guess an insincere ballot sequence for the manipulator and check whether it improves her satisfaction (and is an instance of generalised free-riding) in polynomial time.

Now we show hardness by a reduction from 3-SAT. In the following, recall that, in every issue ℓ , a voter i gives each of her approved candidates an extra score of

$$\omega(\text{sat}(i, (o, \dots, o_{\ell-1})) + 1),$$

where $(o, \dots, o_{\ell-1})$ are the candidates elected thus far.

Let φ be a 3-CNF with n variables and m clauses. We refer to the j -th clause as γ_j . We assume without loss of generality that φ is not satisfied by setting all variables to true. Furthermore, let $h \in \mathbb{N}$ be the smallest h such that $\omega(h) > \omega(h+1)$. As F is not the utilitarian rule, such a h indeed exists and is constant in the size of φ .

We construct a SUCCESSFUL- F -GENERALISED-FREE-RIDING instance with $h+n$ issues as follows: for every variable x_i there are 4 voters $v_i^1, v_i^2, \bar{v}_i^1$ and $\bar{v}_i^2$. Furthermore, we add nine voters v_0^1, v_0^2, w , and $\{u_1, \dots, u_6\}$. Finally, we add another voter i^* , who

will be the distinguished voter that tries to manipulate. In the first $h - 1$ issues, there are two candidates c and $\bar{c}$. In the following n issues, there are three candidates, c_0 , c and $\bar{c}$, plus one private candidate c_i for every voter $i \in N$. In issue $h + n$, we have $m + 1$ candidates $\{c, c_1, c_2, \dots, c_m\}$ plus one candidate c_i for every voter $i \in N$. We assume that if ties need to be broken between c_0 and another candidate, then c_0 is selected. If a tie between c and a candidate other than c_0 needs to be broken, then c wins.

In the first $h - 1$ issues, all voters approve both candidates. Then, we have n issues defined as follows. In all issues $\ell \in \{h, \dots, h + n - 1\}$, v_0^1 and v_0^2 approve c_0 . Moreover, v_ℓ^1 and v_ℓ^2 approve c , while $\bar{v}_\ell^1$ and $\bar{v}_\ell^2$ approve $\bar{c}$. Next, i^* approves both c and $\bar{c}$, while w approves c , $\bar{c}$ and c_0 . Finally, all other voters i only approve their candidate c_i , with the exception of issue h , where u_1 and u_2 additionally approve c_0 , c and $\bar{c}$. We represent the issues in $\{h, \dots, h + n - 1\}$ visually as follows:

Voter	Issue $\ell = h$	Issue $\ell \in \{h + 1, \dots, h + n - 1\}$
i^*	$\{c, \bar{c}\}$	$\{c, \bar{c}\}$
v_0^1, v_0^2	$\{c_0\}$	$\{c_0\}$
w	$\{c, \bar{c}, c_0\}$	$\{c, \bar{c}, c_0\}$
v_ℓ^1, v_ℓ^2	$\{c\}$	$\{c\}$
$\bar{v}_\ell^1, \bar{v}_\ell^2$	$\{\bar{c}\}$	$\{\bar{c}\}$
u_i with $i \in \{1, 2\}$	$\{c_0, c, \bar{c}, c_{u_i}\}$	$\{c_{u_i}\}$
Any other voter j	$\{c_j\}$	$\{c_j\}$

Finally, in issue $h + n$, voters i^* and $\{u_1, \dots, u_6\}$ approve c . Furthermore, for every candidate c_j with $1 \leq j \leq m$, voters $\bar{v}_i^1$ and $\bar{v}_i^2$ approve c_j if and only if variable x_i appears positively in clause γ_j and v_i^1 and v_i^2 approve c_j if and only if variable x_i appears negatively in γ_j . Additionally, w approves $\{c_1, \dots, c_m\}$. All other voters i approve only their candidate c_i .

Before proving that the reduction is correct, let us give an overview. The reduction proceeds in three stages. First, the initial $h - 1$ issues ensure that every voter accumulates the same satisfaction level $h - 1$. Second, the next n issues encode a truth assignment: in each such issue, the manipulator can force either c or $\bar{c}$ to win, corresponding to setting a variable to true or false, respectively. This has no impact on her true satisfaction with respect to this issue. Finally, the last issue checks whether the induced assignment satisfies the formula. The candidate c wins exactly when every clause contains a satisfied literal, so the manipulator can strictly improve her satisfaction if and only if the original 3-SAT instance is satisfiable.

We claim that in issues h to $h + n - 1$ (i) c wins if i^* approves c and $\bar{c}$ (i.e., i^* does not misrepresent her preferences), (ii) either c or $\bar{c}$ becomes the winner if i^* approves only one of these candidates, and (iii) c_0 wins if i^* approves neither c nor $\bar{c}$. We show the claim by induction.

Up until issue $h - 1$, all voters have gathered the same satisfaction $h - 1$, and hence in issue h each voter contributes to the score of their approved candidates with the

same value of $\omega(h)$. The only candidates that are approved by more than one voter are c_0 , c and $\bar{c}$, where c_0 is approved by five voters, while c and $\bar{c}$ are both approved by six voters. Therefore, by our assumption about tiebreaking, c is the winner in issue h . Furthermore, if i^* misrepresents her preferences and votes only for c (respectively, $\bar{c}$), then c (respectively, $\bar{c}$) is the unique winner in issue h . Finally, if i^* approves neither c nor $\bar{c}$, then c_0 wins by our assumption on tiebreaking. Observe that i^* is not allowed to make c_0 win, by the definition of generalised free-riding.

Now assume the claim holds until issue $\ell - 1$. Then, in issue ℓ the only candidates that are approved by more than one voter are, again, c_0 , c and $\bar{c}$. To be more precise, c_0 is approved by v_0^1, v_0^2 and w . Next, c is approved by v_ℓ^1, v_ℓ^2, w and i^* . Finally, $\bar{c}$ by $\bar{v}_\ell^1, \bar{v}_\ell^2, w$ and i^* . Crucially, voters in

$$\{v_0^1, v_0^2, v_\ell^1, v_\ell^2, \bar{v}_\ell^1, \bar{v}_\ell^2\}$$

have not won in any issues in $\{h, \dots, \ell - 1\}$, since they never approve neither c nor $\bar{c}$ in those issues. Thus, their satisfaction at this point is still $h - 1$. Now, let $s_{i^*}^{\ell-1}$ and $s_w^{\ell-1}$ be the satisfaction of i^* and w up to issue $\ell - 1$, respectively. Then, both c and $\bar{c}$ yield a score of $2\omega(h) + \omega(s_{i^*}^{\ell-1} + 1) + \omega(s_w^{\ell-1} + 1)$. Furthermore, c_0 yields a score of $2\omega(h) + \omega(s_w^{\ell-1} + 1)$, whereas all other alternatives can yield at most a score of $\omega(h)$. As we know that $\omega(s_{i^*}^{\ell-1} + 1) > 0$, this implies, by our assumption about tiebreaking, that c is the winner in issue ℓ . Furthermore, as before, if i^* misrepresents her preferences, she can make c (respectively, $\bar{c}$) the unique winner. If she approves neither c nor $\bar{c}$, then c_0 wins by our assumption on tiebreaking. Observe again that c_0 cannot win here, by definition of generalised free-riding.

We can interpret the winners in the issues $\{h, \dots, h + n - 1\}$ as a truth assignment by setting x_i to true if c wins in issue i and to false if $\bar{c}$ wins in issue i (observe that c_0 can never win by the previous arguments). Then, we claim that c wins in issue $h + n$ if and only if φ is satisfied by this truth assignment. All voter-candidates c_i are approved by at most one voter with satisfaction $h - 1$, and hence have a score of at most $\omega(h)$. The satisfaction of v_i^1 and v_i^2 is $h - 1$ if x_i is set to false in the assignment, h if x_i is set to true. Similarly, the satisfaction of $\bar{v}_i^1$ and $\bar{v}_i^2$ is $h - 1$ if x_i is set to true and h if x_i is set to false. Finally, the satisfaction of w is $h + n - 1$. Hence, the score of c_j is $6\omega(h) + \omega(h + n)$ if all literals in γ_j are set to false and at most $4\omega(h) + 2\omega(h + 1) + \omega(h + n)$ if at least one literal is set to true. On the other hand, the satisfaction of u_1 and u_2 is h , the satisfaction of voters in $\{u_3, \dots, u_6\}$ is $h - 1$, and the satisfaction of i^* is $h + n - 1$. Hence, the score of c is $4\omega(h) + 2\omega(h + 1) + \omega(h + n)$. As we assumed $\omega(h + 1) < \omega(h)$, we get that $4\omega(h) + 2\omega(h + 1) + \omega(h + n) < 6\omega(h) + \omega(h + n)$. Therefore, if there is a clause γ_j for which no literal is set to true, then c_j yields a higher score than c and hence c is not a winner in issue $h + n$. On the other hand, if for every clause at least one literal is set to true, then $\{c_1, \dots, c_m\}$ have at most the same score as c , and c wins by tiebreaking.

Now, the honest ballot of i^* leads to the victory of c in all issues in $\{h, \dots, h + n - 1\}$. Hence, the corresponding truth assignment sets every variable to true. By assumption, this assignment does not satisfy φ and hence c does not win in issue $h + n$. By

construction, the satisfaction of i^* equals $h + n - 1$ in this case. Moreover, as i^* cannot manipulate in the first $h - 1$ issues, the only way that i^* can gain more satisfaction is by forcing the winners in issues $\{h, \dots, h + n - 1\}$ to form a satisfying truth assignment without allowing c_0 to win any issue, thereby making c win in the final issue. Hence, i^* can manipulate by generalised free-riding if and only if φ is satisfiable. $\square$

Theorem 4.15. SUCCESSFUL- F -GENERALISED-FREE-RIDING is NP-complete for every sequential $\omega^{(n)}$ -OWA rule such that, for some $n^* \in \mathbb{N}$, $\omega^{(n)}(1) > \omega^{(n)}(n)$ holds for every $n \in \mathbb{N} \setminus [n^*]$.

The condition on n^* aims to capture rules that “for large enough electorate size n ” are distinct from the utilitarian rule.

Proof of Theorem 4.15. Fix a $\omega^{(n)}$ -OWA rule F satisfying the conditions of the theorem. First, notice that SUCCESSFUL- F -GENERALISED-FREE-RIDING is in NP, as we can guess an insincere ballot sequence for the manipulator and check whether it improves her satisfaction (and is an instance of generalised free-riding) in polynomial time.

Next, we show hardness by a reduction from 3-SAT. Let φ be a 3-CNF with n variables and m clauses. We refer to the j -th clause as γ_j . We assume without loss of generality that φ is not satisfied by setting all variables to false. We restrict ourselves to instances satisfying $2n + 5 > n^*$, where n^* is the constant from the theorem statement which depends only on F . This restriction does not affect the NP-hardness of 3-SAT, because the excluded instances contain only a constant number of variables, and therefore there are only finitely many such 3-CNF formulae.

We will construct an instance of SUCCESSFUL- F -GENERALISED-FREE-RIDING with $2n + 5$ voters. More specifically, there are two voters v_i and $\bar{v}_i$ for each variable x_i , four voters $\{u_1, u_2, u_3, u_4\}$, plus one distinguished voter i^* who will try to manipulate. Since the number of voters is fixed, we will simply refer to $\omega^{(2n+5)}$ as ω .

Given the weight vector $\omega = (\omega(1), \dots, \omega(2n + 5))$, we distinguish three (not necessarily exclusive) cases:

1. $\omega(2) > \omega(n + 5)$,
2. $\omega(n + 1) > \omega(2n + 5)$, and
3. $\omega(2) = \omega(2n + 5)$.

Since $\omega(1) > \omega(2n + 5)$, at least one case must be true. In the following, we will give a different reduction for each of the three cases.

First case: $\omega(2) > \omega(n + 5)$. We construct an instance with $n + 2$ issues as follows. In the first $n + 1$ issues, there are two candidates: c and $\bar{c}$. In the last issue, there are $m + 1$ candidates, namely $\{c, c_1, \dots, c_m\}$. We assume that, in the case of ties, c always loses.

In the first issue, everybody votes for $\bar{c}$ except for i^* and u_1 , who vote for c . Here, $\bar{c}$ wins by tiebreaking, and i^* cannot manipulate.

In each issue ℓ with $\ell \in \{2, \dots, n+1\}$, voter v_ℓ votes for candidate $\bar{c}$, voter $\bar{v}_\ell$ for candidate c . Everyone else, including i^* , votes for both candidates. We claim that i^* can manipulate in every such issue ℓ and force the win of either c or $\bar{c}$. We show so by induction. In issue 2, suppose that i^* votes only for c (the case where she votes for $\bar{c}$ is analogous). Then, if c were to win, the satisfaction vector would be of form

$$(1, 1, 1, 2, \dots, 2)$$

(everyone but v_1 wins). On the other hand, if $\bar{c}$ wins, then it would be of form

$$(0, 1, 1, 2, \dots, 2)$$

(everyone but $\bar{v}_1$ and i^* win). Hence, c wins. Observe that if i^* votes truthfully, $\bar{c}$ wins by tiebreaking. Now, suppose this holds up to a issue ℓ . Before issue $\ell+1$, i^* and u_1 won $\ell-1$ issues (all but the first one), whereas voters $\{u_2, u_3, u_4\}$, as well as any pair of voters v_j and $\bar{v}_j$ (with $j \geq \ell$) have won ℓ issues. Furthermore, for every pair v_j and $\bar{v}_j$ (with $j < \ell$), exactly one voter won $\ell-1$ issues while the other ℓ issues. Suppose again that i^* votes for c (the case where she votes for $\bar{c}$ is similar). Then, observe that, if c or $\bar{c}$ win, the satisfaction vectors (excluding i^*) would be completely symmetric (and every voter would have at least a score of ℓ). However, if $\bar{c}$ wins, i^* would have a satisfaction of $\ell-1$, whereas if c wins, she would get a satisfaction of ℓ . Hence, since $\omega(1) > 0$, candidate c wins. Observe again that if i^* votes truthfully, then $\bar{c}$ wins by tiebreaking.

Consider the final issue. Up to here, i^* and u_1 have won n issues (they lost the first issue), while u_2, u_3 , and u_4 won $n+1$ issues. Furthermore, every voter v_ℓ has won n issues if c won in issue $\ell+1$ and $n+1$ times otherwise (and conversely for $\bar{v}_\ell$ and $\bar{c}$). In this issue, voters $\{u_1, \dots, u_4\}$ approve of all candidates but c , voter v_ℓ (respectively, $\bar{v}_\ell$) approves of c and every candidate c_j such that $x_\ell \notin \gamma_j$ (respectively, $\bar{x}_\ell \notin \gamma_j$). Finally, voter i^* approves of c . Observe that we can interpret c winning in issue $\ell+1$ as setting x_ℓ to true, and $\bar{c}$ winning as setting x_ℓ to false. We claim that c wins in the last issue if and only if this assignment satisfies φ . To see this, consider that, if c were to win, the satisfaction vector would be:

$$(n, \underbrace{n+1, \dots, n+1}_{n+4 \text{ times}}, \underbrace{n+2, \dots, n+2}_{n \text{ times}}).$$

Let us call this vector s . Consider now a candidate c_j and its corresponding clause γ_j . If all three of its literals are unsatisfied, then the corresponding voters all have satisfaction $n+1$. Hence, if c_j were to win in this case, the satisfaction vector would again be exactly s . By our assumptions on tiebreaking, here c_j would win against c . Furthermore, suppose that either one, two, or three of the literals have been satisfied. Then, the vectors

are, respectively:

$$\begin{aligned} s_1 &= (n, n, \underbrace{n+1, \dots, n+1}_{n+2 \text{ times}}, \underbrace{n+2, \dots, n+2}_{n+1 \text{ times}}), \\ s_2 &= (n, n, n, \underbrace{n+1, \dots, n+1}_n, \underbrace{n+2, \dots, n+2}_{n+2 \text{ times}}), \text{ and} \\ s_3 &= (n, n, n, n, \underbrace{n+2, \dots, n+1}_{n-2 \text{ times}}, \underbrace{n+2, \dots, n+2}_{n+3 \text{ times}}). \end{aligned}$$

One can show that if $\omega(2) > \omega(n+5)$ then the dot product between ω and each of these three vectors would be strictly lower than the dot product between ω and s . For example:

$$\begin{aligned} s \cdot \omega > s_1 \cdot \omega &\iff \\ \omega(1)n + \left(\sum_{i=2}^{n+5} \omega(i)(n+1) \right) + \left(\sum_{i=n+6}^{2n+5} \omega(i)(n+2) \right) &> \\ (\omega(1) + \omega(2))n + \left(\sum_{i=3}^{n+4} \omega(i)(n+1) \right) + \left(\sum_{i=n+5}^{2n+5} \omega(i)(n+2) \right) &\iff \\ \omega(2) > \omega(n+5). \end{aligned}$$

The other two cases are similar. Hence, if γ_j is satisfied, candidate c_j cannot win against c . Consequently, if all clauses are satisfied, candidate c wins.

Now, if c wins in the last issue, then the satisfaction of i^* would be $n+1$; if c loses, it would be n . Notice also that i^* cannot raise her satisfaction by manipulating in the final issue. Furthermore, if i^* always submits her true preferences, then by tiebreaking $\bar{c}$ would win in every issue ℓ with $\ell \in \{2, \dots, n+1\}$. By assumption, this would not satisfy φ , and hence c would not win in the last issue. Therefore, i^* has an incentive to manipulate in these issues to try and choose a satisfying assignment for φ . It follows that i^* can manipulate by generalised free-riding if and only if φ is satisfiable, so we are done.

Second case: $\omega(n+1) > \omega(2n+5)$. We construct an instance with $n+2$ issues as follows. In the first $n+1$ issues, there are two candidates, c and $\bar{c}$. In the last issue, there are $m+1$ candidates, namely $\{c, c_1, \dots, c_m\}$. We assume that, in the case of ties, c always loses.

In the first issue, voters in $\{i^*, u_1, u_2, u_3, u_4\}$ approve of c , whereas everyone else approves of $\bar{c}$. In each issue ℓ with $\ell \in \{2, \dots, n+1\}$, voter v_ℓ votes for candidate c , voter $\bar{v}_\ell$ for candidate $\bar{c}$, and everyone else votes for both candidates. In the final issue, voters i^*, u_1 and u_2 approve of c , voter v_ℓ (respectively, $\bar{v}_\ell$) approves of candidate c_j if $x_\ell \in \gamma_j$ (respectively, $\bar{x}_\ell \in \gamma_j$). Finally, voters u_3 and u_4 approve of all candidates.

We claim that (i) voters in $\{i^*, u_1, u_2, u_3, u_4\}$ lose in the first issue, and that in the following n issues (ii) candidate $\bar{c}$ wins if i^* votes truthfully and (iii) i^* can force the win

of either c or $\bar{c}$ by manipulating. The arguments for this are essentially the same as in the first case.

Now, consider the last issue. Up to here, voters in $\{i^*, u_1, u_2, u_3, u_4\}$ won n issues (all but the first). Furthermore, every voter v_ℓ has won $n + 1$ issues if c won in issue ℓ and n times otherwise (and conversely for $\bar{v}_\ell$ and $\bar{c}$). We interpret again c winning in issue ℓ as setting x_ℓ to true, and $\bar{c}$ winning as setting x_ℓ to false. We claim that c wins in the last issue if and only if this assignment satisfies φ . To see this, consider that, if c were to win, the satisfaction vector would be:

$$s = (\underbrace{n, \dots, n}_{n \text{ times}}, \underbrace{n+1, \dots, n+1}_{n+5 \text{ times}}).$$

Consider now a candidate c_j and its corresponding clause γ_j . If all three of its literals are unsatisfied, then the corresponding voters all have satisfaction n . Hence, if c_j were to win in this case, the satisfaction vector would again be exactly s . By our assumptions on tiebreaking, here c_j would win against c . Furthermore, suppose that either one, two, or three of the literals have been satisfied. Then, the vectors are, respectively:

$$\begin{aligned} &(\underbrace{n, \dots, n}_{n+1 \text{ times}}, \underbrace{n+1, \dots, n+1}_{n+3 \text{ times}}, n+2), \\ &(\underbrace{n, \dots, n}_{n+2 \text{ times}}, \underbrace{n+1, \dots, n+1}_{n+1 \text{ times}}, n+2, n+2), \text{ and} \\ &(\underbrace{n, \dots, n}_{n+3 \text{ times}}, \underbrace{n+1, \dots, n+1}_{n-1 \text{ times}}, n+2, n+2, n+2). \end{aligned}$$

One can show that if $\omega(n+1) > \omega(2n+5)$ the dot product between ω and each of these three vectors would be strictly lower than the dot product between ω and s . The argument concludes similarly to the first case.

Third case: $\omega(2) = \omega(2n+5)$. We construct an instance with $n+3$ issues as follows. Issues in $\{2, \dots, n+1\}$ and the last issue are equal to the first case. The first issue is almost identical, save for the fact that u_1 also votes for $\bar{c}$.

Hence, we know that i^* loses in the first issue and that she can manipulate in all the following issues to force a win of either c or $\bar{c}$. Let us focus on issue $n+2$. We will design this issue to make sure that only i^* wins, and that she cannot manipulate by generalised free-riding. We distinguish two sub-cases:

1. $\omega(2) = \dots = \omega(2n+5) = 0$, that is, F is the egalitarian rule. Here, everyone votes for c , save for i^* , who votes for $\bar{c}$. There are no other candidates. Observe that, in case either c or $\bar{c}$ wins here, the minimal satisfaction will be n in both cases; $\bar{c}$ wins by tiebreaking. Furthermore, were i^* to vote for c , then c would win (as the minimal satisfaction for c winning would raise to $n+1$).
2. $\omega(2) = \dots = \omega(2n+5) > 0$. Here, there is one candidate c_i for every voter $i \in N$, and we assume that in case of ties c_{i^*} wins. Furthermore, we assume that all voters

$i \in N$ vote for their voter-candidate c_i . We show that here all candidates receive the same score. Consider any two candidates c_i and c_j . Let $\mathbf{i} = (i_1, \dots, i_{2n+5})$ and $\mathbf{j} = (j_1, \dots, j_{2n+5})$ be the satisfaction vectors corresponding to c_i and c_j winning, respectively. For both c_i and c_j , there is surely at least one voter with satisfaction n that disapproves of them; hence, $i_1 = j_1 = n$. Furthermore, there is exactly one voter approving each candidate, and hence

$$\sum_{\ell=2}^{2n+5} i_\ell = \sum_{\ell=2}^{2n+5} j_\ell.$$

These two facts, together with the fact that $\omega(2) = \dots = \omega(2n+5)$, imply that $\omega \cdot \mathbf{i} = \omega \cdot \mathbf{j}$. Hence, every two candidates receive the same score: by tiebreaking, c_{i^*} wins. Now, notice that, if i^* approves of any other candidate c_i distinct from c_{i^*} , then c_i will receive a strictly greater score than any other candidate (as now two voters approve of it).

Now, consider the last issue. Up to here, voters in $\{i^*, u_1, u_2, u_3, u_4\}$ won $n+1$ issues. Furthermore, every voter v_ℓ has won $n+1$ issues if c won in issue ℓ and n times otherwise (and conversely for $\bar{v}_\ell$ and $\bar{c}$). We interpret again c winning in issue ℓ as setting x_ℓ to true, and $\bar{c}$ winning as setting x_ℓ to false. We claim that c wins in the last issue if and only if this assignment satisfies φ . To see this, consider that, if c were to win, the satisfaction vector would be:

$$\mathbf{s} = (\underbrace{n+1, \dots, n+1}_{n+4 \text{ times}}, \underbrace{n+2, \dots, n+2}_{n+1 \text{ times}}).$$

Consider now a candidate c_j and its corresponding clause γ_j . If all three of its literals are unsatisfied, then the corresponding voters all have satisfaction $n+1$. Hence, if c_j were to win in this case, the satisfaction vector would again be exactly $\mathbf{s}$. By our assumptions on tiebreaking, here c_j would win against c . Furthermore, suppose that either one, two, or three of the literals have been satisfied. Then, the vectors are, respectively:

$$\begin{aligned} & (n, \underbrace{n+1, \dots, n+1}_{n+2 \text{ times}}, \underbrace{n+2, \dots, n+2}_{n+2 \text{ times}}), \\ & (n, n, \underbrace{n+1, \dots, n+1}_n, \underbrace{n+2, \dots, n+2}_{n+3 \text{ times}}), \text{ and} \\ & (n, n, n, \underbrace{n+1, \dots, n+1}_{n-2 \text{ times}}, \underbrace{n+2, \dots, n+2}_{n+4 \text{ times}}). \end{aligned}$$

One can show that if $\omega(1) > \omega(n+4)$ (which is implied by $\omega(1) > \omega(2n+5)$ and $\omega(2) = \omega(2n+5)$) the dot product between ω and each of these three vectors would be strictly lower than the dot product between ω and $\mathbf{s}$. The argument concludes analogously to the previous cases.

This concludes the proof. □

This concludes the technical part of this section. Since the section has been rather long, we close with a brief recap.

We have argued that deciding whether free-riding is possible is computationally easy. At the same time, we have shown that deciding whether a manipulator can improve her satisfaction by free-riding is generally hard for (sequential) Thiele rules, and at least for the egalitarian rule. We have also extended these results to a broader class of Thiele rules, as well as to essentially all [OWA](#) rules, in the setting of generalised free-riding. This more permissive manipulation strategy is closely related to free-riding and captures situations in which a voter may try to manipulate by voting against a popular candidate whom she actually approves of.

4.6 Numerical Simulations

So far, we have seen that sequential rules are generally susceptible to free-riding. On the positive side, we have also seen that free-riding can be harmful to the free-riders themselves, in the sense that their final satisfaction may decrease, and that determining in advance whether free-riding will be beneficial is computationally hard under many natural rules. This, however, does not say anything about its practical risk. If free-riding typically turns out to be beneficial in realistic instances, voters may still be tempted to try it in the hope of increasing their satisfaction, thereby circumventing the computational difficulty of the decision problem.

Hence, in this section, we will use numerical simulations to shed more light on the *possibility*, *effect*, and *risk* of free-riding with sequential rules. The idea is to generate a number of elections using a sampling algorithm, simulate the corresponding rules, and then simulate free-riding by a voter, as described below. Our code is publicly available online [[Lackner et al., 2025](#)].

We address three questions:

1. How often do voters have the possibility to increase their (final) satisfaction by free-riding? How many elections contain issues in which free-riding leads to a lower satisfaction for the free-rider?
2. What is the average risk of free-riding? That is, what is the likelihood of free-riding resulting in a negative outcome?
3. Is there a difference in the effect of free-riding if it is done by a voter belonging to a majority or minority?

The motivation behind the last question is that, intuitively, free-riding might appeal to minorities and majorities differently. On the one hand, it seems easier for voters belonging to a majority to free-ride, since they will more often vote for a popular candidate. On the other hand, minority voters might have more to gain from manipulating, since they will be generally less satisfied with the outcome.

To be able to give robust answers, we study these questions in a range of different statistical models (i.e., sampling algorithms). We describe them in the next section.

4.6.1 Model Setup and Parameter Values

We employ the so-called $2d$ -Euclidean model, a rather common statistical model of elections [Enelow and Hinich, 1990, Elkind et al., 2017, Brederbeck et al., 2019, Godziszewski et al., 2021]. Focusing on a single issue, we imagine that the space of possible opinions is captured by the square $[-1, 1] \times [-1, 1]$. Each voter is a point in this space; intuitively, Euclidean distance captures “opinion distance” between voters. As mentioned, Euclidean models are rather common in voting; we adapt them to the repeated setting.

We generate an election as follows. First, we sample n point from some distribution (we describe the exact distributions used below) on the $[-1, 1] \times [-1, 1]$ square. These represent the voters, and will remain fixed throughout the election. Then, for each issue, we sample independently and uniformly m points on the same square. These represent the candidates (which change in every issue). The idea is that the relative positions of voters remain fixed, but each issue captures a different political problem.

We consider three distributions for voters:¹⁵

- *Square*: Voters are drawn from a uniform distribution on the $[-1, 1] \times [-1, 1]$ square, i.e., voters and candidates are independently drawn from the same distribution.
- *Many groups*: Voters are split into four groups. The first three groups each contain 20% of voters, and the fourth contains the remaining 40%. These groups are centred at

$$(0.5, 0.5), \quad (0.5, -0.5), \quad (-0.5, 0.5), \quad \text{and} \quad (-0.5, -0.5).$$

For each voter, the x - and y -coordinates are drawn independently from a normal distribution centred at the corresponding group centre with standard deviation 0.2.

- *Unbalanced*: Voters are divided into two disjoint groups, one containing 20% of voters and the other 80%. The groups are centred at $(0.5, 0.5)$ and $(-0.5, -0.5)$, respectively. As above, each coordinate is drawn independently from a normal distribution, now with standard deviation 0.1.

Given a set of points (voters and candidates) representing an issue, we transform these points into approval ballots as follows. Every voter approves her closest candidate. Additionally, every voter approves all candidates that lie within 1.5 times the distance to her closest candidate. Intuitively, this captures a situation where a voter approves her favourite candidate as well as any candidate that are similarly close.

Our default choice are multi-issue elections with $n = 20$ voters, $k = 20$ issues, and 5 candidates per issue. Later, we will investigate the effect of this parameter choices. We

¹⁵The unbalanced and many groups scenarios have been proposed by Chandak et al. [2026]. We use the same parameter values to increase comparability, even though their election model is different from ours.

sample 2000 elections for each of the three voter distributions and choice of parameters. Figure 4.1 illustrates the different voter distributions as well as the approval generation.




Figure 4.1: Visualisation of the sampling strategy. The top row shows three collections of voters, sampled from the *square*, *many groups* and *unbalanced* distributions. The bottom row presents the approvals of a selected voter across three different issues. Observe that the positions of the voters remain fixed, while the candidates vary. The red line indicates the closest candidate, and candidates within the circle are approved.

4.6.2 Considered Voting Rules and Forms of Free-Riding

In our simulations, we consider specific subclasses of sequential Thiele and OWA rules. For better comparison, we parametrise both classes with a parameter x (albeit this parameter has a different interpretation in both classes). We consider Thiele rules with

$$\omega_{(x)} = \left(1, \frac{1}{2^x}, \frac{1}{3^x}, \frac{1}{4^x}, \dots\right) \quad \text{for } x \in \{0, 0.25, 0.5, 0.75, \dots\}.$$

For $x = 0$ this is the utilitarian rule, for $x = 1$ it is sequential [PAV](#), and for increasing x it approaches the sequential leximin rule, since $\omega_{(x)}$ diminishes quickly. Further, we consider [OWA](#) rules with

$$\omega_{(x)}^{(n)} = (\underbrace{1, \dots, 1}_{n-x \text{ many}}, \frac{1}{kn}, \frac{1}{k^2n^2}, \dots) \quad \text{for } x \in \{0, 1, \dots, n-1\}.$$

(Note that here we parametrise both by n , the number of voters, and by x .) For decreasing x , more and more voters receive full consideration (weight 1). For $x = 0$, all voters receive full consideration, and we obtain the utilitarian rule. By contrast, for $x = n - 1$, only the least satisfied voter receives full consideration; this is the leximin rule (Proposition [4.1](#)).

We distinguish two ways in which free-riding can occur. We speak of *single-issue free-riding* if a voter may free-ride once in a given election. By contrast, with *repeated free-riding* a voter free-rides whenever it is possible in a given election. Single-issue free-riding models situations where free-riding occurs rarely, e.g., if there is limited information available about expected outcomes, or if there is a high social cost for free-riding. Repeated free-riding models situations where free-riding is a viable strategy and a free-riding voter can always exploit manipulation opportunities.

As in the rest of the chapter, we assume that there is only one free-rider. We leave the study of coalitional free-riding as a topic for future work.

4.6.3 Results

Let us present and discuss our findings.

4.6.3.1 Possibility and Risk of Free-Riding: Simulation Study

In the single-issue free-riding model, we speak of *successful free-riding* for a given election, voter, and issue, if the voter can free-ride in the given issue and this increases her overall satisfaction (i.e., at the end of the election); we speak of *harmful free-riding* if the voter can free-ride but this decreases her satisfaction. Note that free-riding can also be neutral (with no change in satisfaction). In the repeated setting, we speak of *successful free-riding* for a given election and voter if the voter free-rides whenever possible and this increases her overall satisfaction; we speak of *harmful free-riding* if this decreases her satisfaction.

Single-issue free-riding. Let us first consider the single-issue free-riding model. That is, for each multi-issue election, we iterate over all voters and all issues and check whether free-riding is possible (Definition [4.4](#)). Figure [4.2](#) shows the proportion of instances where in at least one issue successful and harmful free-riding is possible (averaged over all voters and elections). As successful and harmful free-riding may be possible in the same election—albeit in different issues—the sum of these proportions may be larger than 1. In addition, it shows the risk of free-riding: we define the risk of a voter in an election as the number of issues where harmful free-riding occurs divided by the number of issues where either successful or harmful free-riding occurs. We average




Figure 4.2: Possibility and risk of single-issue free-riding. The left column of diagrams show sequential Thiele methods, the right sequential OWA methods. The three rows correspond to the *square*, *many groups*, and *unbalanced* distributions.

the risk over all elections where a voter might change her satisfaction by free-riding, i.e., where the denominator of the risk is positive.

Let us discuss the results from Figure 4.2. First, we see that the results for the *square* and *many groups* distributions are very similar. By contrast, the *unbalanced* distribution yields different results. For all distributions, we clearly see that rules closer to the utilitarian rule ($x = 0$) are less susceptible to free-riding than those closer to leximin (larger values of x). We also see that, as expected, the utilitarian rule is the only rule where free-riding is not possible (see Observation 4.2). We note that this increase in susceptibility (with distance to the utilitarian rule) has also been observed by Barrot et al. [2017] for arbitrary manipulations. Both the proportion of voters that can successfully free-ride and those with the possibility of harmful free-riding grow with parameter x .

The most important conclusion from this simulation is that the risk of single-issue free-riding is considerable. For the *square*, *many groups* and *unbalanced* distributions the risks are, respectively:

- 2.7%, 4.0% and 0.7% for PAV (Thiele, $x = 1$), and
- 16.5%, 16.2% and 8.4% for leximin (OWA, $x = 19$).

This shows that harmful free-riding is not a mere theoretical possibility. In particular for more egalitarian voting rules (i.e., larger values of x in Figure 4.2), this risk could conceivably decrease the temptation of free-riding.

Repeated free-riding. For repeated free-riding, we define risk as follows. Here, the risk of a voter in a given election is either 0 (free-riding increases the satisfaction), or 1 (free-riding decreases the satisfaction). This apparent difference with respect to the single-issue free-riding setting is motivated as follows. For a given election, in the single-issue setting, the manipulator can choose on which issue to free-ride, leading to a fine-grained notion of risk (i.e., fraction of issues that lead to a decrease in satisfaction). For the repeated model, there is no choice besides to free-ride or not to free-ride. This explains the binary (i.e., one-or-zero) nature of risk in this setting. Risk is averaged over all voters (for whom successful or harmful free-riding is possible) and over all elections.

Let us look at the results in Figure 4.3. We see that the risk is small compared to the single-issue model. In particular, we see much higher likelihood of successful free-riding with almost no risk involved. This is because the probability of successful free-riding in a single issue is generally larger than harmful free-riding (i.e., risk in the single-issue model is lower-or-equal than 0.5), the risk almost vanishes with sufficient repetitions. Recall, however, that this form of free-riding requires the capability to reliably identify free-riding possibilities and to repeatedly exploit them, as well as a low social cost for free-riding. One slight exception is the case of sequential OWA rules under the *unbalanced* distribution, at least for small values of x (i.e., rules close to leximin). In this setting, all measures are close to zero. One possible explanation is that these rules are strongly egalitarian; hence, voters might easily change the outcome by voting against the majority opinion (which leads to a situation that does not fit our notion of free-riding).




Figure 4.3: Results for repeated free-riding. The left column of diagrams show sequential Thiele methods, the right sequential OWA methods. The three rows correspond to the *square*, *many groups*, and *unbalanced* distributions. In contrast to Figure 4.2, we see a much higher likelihood of successful free-riding with very little risk.




Figure 4.4: Average change in satisfaction from free-riding, comparing different voter distributions as well as single-issue vs. repeated free-riding.

Quantifying the benefits of free-riding. So far, we have not taken into account how much a voter can benefit from free-riding. To investigate this, we show in Figure 4.4 the expected change in satisfaction for a free-riding voter. As one could expect, repeated free-riding leads to a more pronounced change in satisfaction. Furthermore, the change in satisfaction grows with parameter x , i.e., with increasing distance to the utilitarian rule. This aligns with the increased possibility to free-ride (as we have seen in Figures 4.2 and 4.3). Finally, note that in all cases, the change in satisfaction is not particularly large in comparison to the number of issues (20). Even with repeated free-riding when using the leximin rule, the expected number issues with an improvement is less than 1.6. This suggests that risk-averse voters may have little incentive to free-ride in practice, since the potential gains appear relatively limited.

4.6.3.2 Impact of Model Parameters

Let us now briefly describe the impact of our chosen model parameters. In the following, we use our default setting ($n = 20$ voters, $k = 20$ issues, 5 candidates per issue, *square* distribution, single-issue free-riding) and vary a single parameter to observe its impact. For the sake of space, we focus on two prominent rules, sequential PAV and sequential leximin.

Increasing the number of voters (Figure 4.5) decreases the chance of voters being pivotal. In line with this observation, we see an overall decrease in both successful and harmful free-riding with an increase in voters. For a larger number of voters, it would make sense to move to a model where groups of voters free-ride. This requires additional assumptions about voter coordination (see, e.g., the framework of iterative voting [Meir, 2017]).

In Figure 4.6, we see that varying the number of candidates (per issue) has moderate impact. The increase in the likelihood of both successful and harmful free-riding can be explained by the increase in heterogeneity: as voters' preferences become more diverse and like-minded groups become smaller, a single free-riding action is more likely to change the outcome.




Figure 4.5: Varying the number of voters (our default choice is $n = 20$): an increase in the number of voters decreases the effectiveness of free-riding by a single voter.




Figure 4.6: Varying the number of candidates (our default choice is $m = 5$): more candidates per issue increase the likelihood of free-riding.




Figure 4.7: Varying the number of issues (our default choice is $k = 20$): a longer sequence of decisions increases the potential impact of free-riding.

Finally, increasing the number of issues (Figure 4.7) also gives a slight increase in the possibility of both successful and harmful free-riding, as some effects may only materialise in the long run. Note that also the risk of free-riding increases in a longer sequence of issues, in particular for sequential leximin.




Figure 4.8: The possibility of successful and harmful free-riding for minority and majority voters.

4.6.3.3 Minority and Majority Free-Riders

As a final simulation experiment, we study the differences between free-riders belonging to a minority or to a majority. Here, we use the repeated free-riding model, as it is generally more beneficial for the free-rider (as we have discussed in Section 4.6.3.1) and thus magnifies differences between free-riding and truthful voters. Results are in Figures 4.8 and 4.9.

Let us first consider the many groups distribution. We randomly select one voter from the 40% group as “majority” voter, and one from a 20% group as “minority” voter. Figure 4.8 (upper part for *many groups*) shows that a majority voter is slightly more likely to successfully free-ride. This can be explained intuitively by the interpretation that free-riding majority voters form their own (one-voter) minority and thus are able to benefit from more egalitarian rules. However, when considering the average change in satisfaction for free-riding voters (Figure 4.9, upper part for *many groups*), we see that the impact of free-riding is small: the change in satisfaction for a free-riding voter (both minority and majority) is small in comparison to her total satisfaction.

The results for the unbalanced distribution are very similar in terms of satisfaction (Figure 4.9, lower part). A noteworthy difference is that—for Thiele rules—minority voters have a slightly larger change of successful free-riding. This is probably due to the small weight that a single majority voter has within the majority block (16 voters in total). A second interesting observation is that free-riding is here almost nonexistent for close-to-utilitarian OWA rules (Figure 4.8, lower part). This is due to the fact that for




Figure 4.9: The change in satisfaction for minority and majority voters in relation to the average satisfaction of voters in either groups *without* free-riding. We see that the change in satisfaction is similar for both groups (and relatively small).

these rules the majority almost always dictates the decision and consequently free-riding is without effect.

Overall, we can conclude that for highly structured instances (such as the many groups and unbalanced distributions) free-riding is not an overly important phenomenon, due to its very limited impact on satisfaction. This statement, however, only holds for single-voter free-riding. Here, we are not studying models for groups of free-riders. By contrast, we expect that group free-riding has significant impact in these scenarios (depending on the chosen model), in particular if a majority subgroup forms a “fake minority”. Note that a group free-riding model will require additional assumptions on information exchange between free-riders (is group free-riding even possible?) and group cohesiveness (are group preferences identical or merely similar?).

4.7 Beyond the Sequential Case

In this section, we will consider the case of multiple issues decided simultaneously i.e., the non-sequential case. For example, this happens with multiple referenda [Amanatidis et al., 2015], already mentioned several times in this chapter, which is a special case of our model where all issues have two candidates (yes or no).

In summary, in this section, we will show that, while free-riding is still inevitable when we want to guarantee some form of representation across issues, in the non-

sequential setting free-riding might be “safe”. Still, we show that deciding whether free-riding is possible is a hard computational problem.

4.7.1 Non-Sequential Rules

As explained in Section 4.3, sequentiality is captured in our model as a feature of the voting rule. Hence, to analyse the non-sequential case, we can just focus on different classes of rules which lack the sequentiality property. Concretely, instead of studying sequential Thiele and OWA rules, we will focus on their global optimisation counterparts, similarly to how we originally introduced these families in Chapter 2. To avoid misunderstandings, and in contrast to elsewhere in the thesis, we will explicitly call the rules discussed here as *opt- F* , meaning “global optimisation version of F ”. We will adopt a similar convention for families of rules (i.e., *opt-Thiele*, *opt-OWA*).

The optimisation-based rule corresponding to the OWA rule $\omega^{(n)}$ simply selects the complete outcome $\mathbf{o}$ maximising

$$s(\mathcal{E}, \mathbf{o}) \cdot \left(\omega^{(n)}(1), \dots, \omega^{(n)}(n) \right)$$

whereas the rule corresponding to the Thiele rule ω selects $\mathbf{o}$ maximising

$$\sum_{i \in N} \sum_{\ell=1}^{sat_{\mathcal{E}}(i, \mathbf{o})} \omega(\ell).$$

Again, we assume a fixed tiebreaking over the outcomes. In particular, we assume that each issue is associated with an ordering over its candidates, and outcomes are compared lexicographically. For example, assuming that each issue uses the alphabetical ordering, outcome (a, b, c) is preferred over (a, c, c) which is preferred, in turn, over (b, a, a) .

To illustrate these definitions, we recall some rules. The *opt-egalitarian* rule selects an outcome $\mathbf{o}$ maximising $\min_{i \in N} sat(i, \mathbf{o})$. This rule is NP-hard to compute [Amanatidis et al., 2015]. By contrast, the *opt-utilitarian* rule is equivalent to its sequential formulation, and thus is polynomial-time computable. To give another familiar example, *opt-leximin* is the *opt-OWA* rule with $\omega^{(n)} = (1, 1/kn, 1/k^2n^2, \dots)$ (Proposition 4.1 still holds in this setting). This rule selects any outcome that is maximal with respect to the leximin ordering $\mathcal{L}(\mathcal{E})$ described in Definition 4.2.

4.7.2 Possibility and Risk of Free-Riding with Non-Sequential Rules

We believe that, in this setting, free-riding is rather different from the sequential case. In the latter, issues are presented one by one, and at each step, a voter may choose to free-ride *in the current issue*. If a voter believes that a candidate she approves has a strong chance of victory (i.e., it has a large margin of support), she might be tempted to free-ride to immediately lower her (perceived) satisfaction. Our previous results show that it is computationally easy to free-ride but hard to determine whether it is possible

to improve one's satisfaction—even with full information. Moreover, in many settings, it is risky to do so with no information regarding future issues.

By contrast, free-riding in the non-sequential setting is intuitively more akin to classical, unrestricted manipulation—that is to say, submitting *any* vector of ballots that differs from your true preferences with the goal of increasing your satisfaction. Given the interconnections among issues and the large strategy space, each voter here must carefully consider how to change her ballots to decide whether free-riding is possible, let alone beneficial. Indeed, as we will see, the main barrier against free-riding in the non-sequential case is the hardness of determining the winner (or the consequences of changing one's ballot). Thus, it seems free-riding here loses its simple, “greedy” appeal. Under this view, restricting oneself to free-riding in this setting is not advantageous to a manipulator: while it reduces the strategy space, computationally it is not significantly easier than arbitrary manipulation.

Just as in the sequential case, we start by studying the possibility and risk of free-riding. Given that our axioms from Section 4.4 also apply to this setting, our results already indicate that, for essentially all rules that include some form of fairness towards minorities, free-riding is possible. However, not all the rules we study satisfy those axioms (again, this is the case of opt-egalitarian). Thus, similarly to Theorem 4.7, we integrate those results with the following.

Theorem 4.16. Every opt-Thiele and opt-OWA rule except the utilitarian rule can be manipulated by free-riding on only one issue.

Proof. Both constructions are the same as in the proof of Theorem 4.7.

Consider an opt-Thiele rule defined by ω and different from the utilitarian rule. Then, there exists a k such that $\omega(k) > \omega(k+1) \geq 0$. Consider an election with $k+1$ issues, four voters and with two candidates a and b for all issues. Suppose that for the first k issues all voters only approve candidate a . Moreover, on issue $k+1$ voters 1 and 2 approve b while voter 3 and 4 approve a . Assume further that a is prioritised over b in the tiebreaking order. Clearly, a wins in the first k issues. By the symmetry of issue $k+1$, together with the tiebreaking order, the final outcome is $(a, \dots, a)$.

We claim that voter 1 can manipulate by changing her vote in one of the first k issues to $\{b\}$. For ease of presentation, assume that she free-rides in issue k . Then, it is still clearly best to select a in all issues in $[k-1]$. Hence, all voters will have satisfaction at least $k-1$. We can compute the total score of the four possible complete outcomes:

$$\begin{aligned} (a, \dots, a, a, a) & \text{ leads to score } \left(4 \sum_{s \in [k-1]} \omega(s) \right) + 3\omega(k) + 2\omega(k+1), \\ (a, \dots, a, a, b) & \text{ leads to score } \left(4 \sum_{s \in [k-1]} \omega(s) \right) + 4\omega(k) + \omega(k+1), \end{aligned}$$

$$\begin{aligned}
(a, \dots, a, b, a) &\text{ leads to score } \left(4 \sum_{s \in [k-1]} \omega(s) \right) + 3\omega(k), \text{ and} \\
(a, \dots, a, b, b) &\text{ leads to score } \left(4 \sum_{s \in [k-1]} \omega(s) \right) + 2\omega(k) + \omega(k+1).
\end{aligned}$$

As $\omega(k) > \omega(k+1) \geq 0$, this implies that $(a, \dots, a, a, b)$ is the winning outcome. Therefore, voter 1 did manipulate by free-riding.

Consider an opt-OWA rule defined by $\omega^{(n)}$ that is not the utilitarian rule. Then there exists a n for which $\omega^{(n)}$ satisfies $\omega^{(n)}(1) > \omega^{(n)}(n)$. Consider an election with 2 issues and n voters. In both issues there are n candidates, namely $\{a_1, \dots, a_n\}$. In the first issue, voters 1 and 2 approve only a_1 . Every other voter $i \in \{3, \dots, n\}$ approves only a_i . In the second issue voter 1 approves a_1 , voter 2 approves a_2 and all other voters approve both a_1 and a_2 . We assume that candidates with a lower index are prioritised by the tiebreaking, which is applied lexicographically. It is straightforward to show that either a_1 or a_2 must win in the second issue. Selecting a_2 or a_i with $i \in \{3, \dots, n\}$ in the first issue, irrespectively to whether a_1 or a_2 is selected in the second issue, leads to sorted satisfaction vectors $(0, 1, \dots, 1, 1)$ or $(0, 1, \dots, 1, 2)$, respectively. On the other hand, selecting a_1 in the first issue leads to satisfaction vector $(1, 1, \dots, 1, 2)$ independently of whether a_1 or a_2 is selected in the second issue. This means (a_1, a_1) and (a_1, a_2) lead to the highest opt-OWA score. By tiebreaking, (a_1, a_1) wins.

Now, we claim that voter 2 can free-ride by approving a_2 instead of a_1 in the first issue. Again, one can easily show that the winner in issue 2 must be either a_1 or a_2 . Furthermore, one can show that outcomes (a_1, a_1) , (a_2, a_2) , (a_i, a_1) , and (a_i, a_2) (with $i \in [n] \setminus [2]$) all yield satisfaction vector $(0, 1, \dots, 1, 2)$. On the other hand, (a_1, a_2) and (a_2, a_1) yield $(1, \dots, 1)$. Since $\omega^{(n)}(1) > \omega^{(n)}(n)$, the latter two outcomes optimise the objective. By tiebreaking, (a_1, a_2) is the winning outcome. It follows that voter 2 did manipulate by free-riding. $\square$

We now show a crucial difference with respect to the sequential case. For all opt-Thiele rules, as well as opt-leximin, we find that the satisfaction of the free-riding voter can never decrease.

Proposition 4.17. Free-riding cannot reduce the satisfaction of the free-riding voter when an opt-Thiele rule is used, but it can increase the satisfaction of the free-riding voter.

Proof. Fix the opt-Thiele rule defined by ω . It follows directly from Theorem 4.16 that free-riding can increase the satisfaction of the free-riding voter. Let us show that it can never decrease it. Let $\mathcal{E}$ be an election, $\mathbf{o}$ be the outcome of $\mathcal{E}$ under the ω -Thiele rule, and consider a voter i^* such that i^* can free-ride in issues $L \subset [k]$. Finally, let $\mathcal{E}^*$ be the election after i^* free-rides and $\mathbf{o}^*$ the outcome of $\mathcal{E}^*$ under the same ω -Thiele rule. Now, as i^* free-rides, i.e., the winners in issues L are the same in $\mathbf{o}$ and $\mathbf{o}^*$, we know that i^*

approves the winners of issues L in her honest ballot and does *not* approve the winners of issues L in her free-riding ballot. It follows that the satisfaction of i^* with respect to $\mathbf{o}$ in $\mathcal{E}$ is $|L|$ points higher than in $\mathcal{E}^*$, and similarly for $\mathbf{o}^*$. Precisely,

$$\begin{aligned} \text{sat}_{\mathcal{E}^*}(i^*, \mathbf{o}) &= \text{sat}_{\mathcal{E}}(i^*, \mathbf{o}) - |L| \quad \text{and} \\ \text{sat}_{\mathcal{E}^*}(i^*, \mathbf{o}^*) &= \text{sat}_{\mathcal{E}}(i^*, \mathbf{o}^*) - |L|. \end{aligned} \quad (4.1)$$

All other voters submit the same ballot in $\mathcal{E}$ and $\mathcal{E}^*$. Hence, for all $i \in N \setminus \{i^*\}$ we have

$$\begin{aligned} \text{sat}_{\mathcal{E}^*}(i, \mathbf{o}) &= \text{sat}_{\mathcal{E}}(i, \mathbf{o}) \quad \text{and} \\ \text{sat}_{\mathcal{E}^*}(i, \mathbf{o}^*) &= \text{sat}_{\mathcal{E}}(i, \mathbf{o}^*). \end{aligned} \quad (4.2)$$

Now assume for the sake of a contradiction that $\text{sat}_{\mathcal{E}}(i^*, \mathbf{o}) > \text{sat}_{\mathcal{E}}(i^*, \mathbf{o}^*)$, i.e., free-riding led to a lower satisfaction for i^* with respect to her honest ballot.

As $\mathbf{o}^*$ is the winning outcome of $\mathcal{E}^*$, we know

$$\sum_{i \in N} \sum_{s=1}^{\text{sat}_{\mathcal{E}^*}(i, \mathbf{o}^*)} \omega(s) > \sum_{i \in N} \sum_{s=1}^{\text{sat}_{\mathcal{E}^*}(i, \mathbf{o})} \omega(s).$$

Note that the inequality is strict because, if the two sides were equal, it would mean that $\mathbf{o}^*$ is preferred over $\mathbf{o}$ by the tiebreaking. Moreover, by Equations (4.1) and (4.2), it would also imply that $\mathbf{o}$ and $\mathbf{o}^*$ would have the same score in $\mathcal{E}$. Hence, $\mathbf{o}^*$ would win in $\mathcal{E}$, contradicting our assumptions. Next, by the previous inequality and Equations (4.1) and (4.2),

$$\sum_{s=1}^{\text{sat}_{\mathcal{E}}(i^*, \mathbf{o}^*) - |L|} \omega(s) + \sum_{i \in N \setminus \{i^*\}} \sum_{s=1}^{\text{sat}_{\mathcal{E}}(i, \mathbf{o}^*)} \omega(s) > \sum_{s=1}^{\text{sat}_{\mathcal{E}}(i^*, \mathbf{o}) - |L|} \omega(s) + \sum_{i \in N \setminus \{i^*\}} \sum_{s=1}^{\text{sat}_{\mathcal{E}}(i, \mathbf{o})} \omega(s).$$

This can be rewritten as

$$\sum_{i \in N} \sum_{s=1}^{\text{sat}_{\mathcal{E}}(i, \mathbf{o}^*)} \omega(s) - \sum_{s=\text{sat}_{\mathcal{E}}(i^*, \mathbf{o}^*) - |L| + 1}^{\text{sat}_{\mathcal{E}}(i^*, \mathbf{o}^*)} \omega(s) > \sum_{i \in N} \sum_{s=1}^{\text{sat}_{\mathcal{E}}(i, \mathbf{o})} \omega(s) - \sum_{s=\text{sat}_{\mathcal{E}}(i^*, \mathbf{o}) - |L| + 1}^{\text{sat}_{\mathcal{E}}(i^*, \mathbf{o})} \omega(s),$$

which gives

$$\sum_{i \in N} \sum_{s=1}^{\text{sat}_{\mathcal{E}}(i, \mathbf{o}^*)} \omega(s) - \sum_{i \in N} \sum_{s=1}^{\text{sat}_{\mathcal{E}}(i, \mathbf{o})} \omega(s) > \sum_{s=\text{sat}_{\mathcal{E}}(i^*, \mathbf{o}^*) - |L| + 1}^{\text{sat}_{\mathcal{E}}(i^*, \mathbf{o}^*)} \omega(s) - \sum_{s=\text{sat}_{\mathcal{E}}(i^*, \mathbf{o}) - |L| + 1}^{\text{sat}_{\mathcal{E}}(i^*, \mathbf{o})} \omega(s). \quad (4.3)$$

Next, we know that the left-hand side of Equation (4.3) must be nonpositive, as $\mathbf{o}$, and not $\mathbf{o}^*$, is the winner of $\mathcal{E}$. Moreover, we assumed that $\text{sat}_{\mathcal{E}}(i^*, \mathbf{o}) > \text{sat}_{\mathcal{E}}(i^*, \mathbf{o}^*)$, which together with the fact that ω is nonincreasing implies that the right-hand side of Equation (4.3) must be nonnegative: a contradiction. Hence, we get $\text{sat}_{\mathcal{E}}(i^*, \mathbf{o}) \leq \text{sat}_{\mathcal{E}}(i^*, \mathbf{o}^*)$, completing the proof. $\square$

Proposition 4.18. Free-riding cannot reduce the satisfaction of the free-riding voter when opt-leximin is used, but it can increase the satisfaction of the free-riding voter.

Proof. This proof is similar to the proof of Theorem 4.12.

It follows directly from Theorem 4.7 that free-riding can increase the satisfaction of the free-riding voter under opt-leximin. Let us show that it can never decrease it. Let $\mathcal{E}$ be an election, $\mathbf{o}$ be the outcome of $\mathcal{E}$ under the opt-leximin rule, and consider a voter i^* such that i^* can free-ride in issues $L \subset [k]$. Finally, let $\mathcal{E}^*$ be the election after i^* free-rides and $\mathbf{o}^*$ the outcome of $\mathcal{E}^*$ under opt-leximin. We define

$$N_s^{\mathcal{E}}(\mathbf{o}) = \{i \in N : \text{sat}_{\mathcal{E}}(i, \mathbf{o}) = s\}.$$

Now, as i^* free-rides, i.e., the winners in issues L are the same in $\mathbf{o}$ and $\mathbf{o}^*$, we know that i^* approves the winners of issues L in her honest ballot and does *not* approve the winners of issues L in her free-riding ballot. It follows that the satisfaction of i^* with respect to $\mathbf{o}$ in $\mathcal{E}$ is $|L|$ points higher than in $\mathcal{E}^*$, and similarly for $\mathbf{o}^*$. Precisely,

$$\begin{aligned} \text{sat}_{\mathcal{E}^*}(i^*, \mathbf{o}) &= \text{sat}_{\mathcal{E}}(i^*, \mathbf{o}) - |L| \quad \text{and} \\ \text{sat}_{\mathcal{E}^*}(i^*, \mathbf{o}^*) &= \text{sat}_{\mathcal{E}}(i^*, \mathbf{o}^*) - |L|. \end{aligned} \quad (4.4)$$

All other voters submit the same ballot in $\mathcal{E}$ and $\mathcal{E}^*$. Hence, for all $i \in N \setminus \{i^*\}$ we have

$$\begin{aligned} \text{sat}_{\mathcal{E}^*}(i, \mathbf{o}) &= \text{sat}_{\mathcal{E}}(i, \mathbf{o}) \quad \text{and} \\ \text{sat}_{\mathcal{E}^*}(i, \mathbf{o}^*) &= \text{sat}_{\mathcal{E}}(i, \mathbf{o}^*). \end{aligned} \quad (4.5)$$

Now assume for the sake of a contradiction that $\text{sat}_{\mathcal{E}}(i^*, \mathbf{o}) > \text{sat}_{\mathcal{E}}(i^*, \mathbf{o}^*)$, i.e., free-riding led to a lower satisfaction for i^* with respect to her honest ballot.

First of all, observe that $\mathbf{o}^*$ and $\mathbf{o}$ cannot tie in the opt-leximin order in $\mathcal{E}^*$ (i.e., the victory of $\mathbf{o}^*$ over $\mathbf{o}$ cannot be due to tiebreaking). Were this not the case, then we would have $|N_s^{\mathcal{E}^*}(\mathbf{o}^*)| = |N_s^{\mathcal{E}^*}(\mathbf{o})|$ for every index s . However, $\text{sat}_{\mathcal{E}}(i^*, \mathbf{o}) > \text{sat}_{\mathcal{E}}(i^*, \mathbf{o}^*)$ and Equation (4.4) imply that $\text{sat}_{\mathcal{E}^*}(i^*, \mathbf{o}) > \text{sat}_{\mathcal{E}^*}(i^*, \mathbf{o}^*)$. This, together with Equation (4.5), implies that $|N_s^{\mathcal{E}^*}(\mathbf{o}^*)| \neq |N_s^{\mathcal{E}^*}(\mathbf{o})|$ for some index s .

Consequently, we know that $(\mathbf{o}^*, \mathbf{o}) \in \mathcal{L}(\mathcal{E}^*)$. In other words, there is a s^* such that $|N_{s^*}^{\mathcal{E}^*}(\mathbf{o}^*)| < |N_{s^*}^{\mathcal{E}^*}(\mathbf{o})|$ and $|N_s^{\mathcal{E}^*}(\mathbf{o}^*)| = |N_s^{\mathcal{E}^*}(\mathbf{o})|$ for all $s < s^*$. We claim that the deciding index s^* cannot be smaller than $\text{sat}_{\mathcal{E}^*}(i^*, \mathbf{o}^*)$ as for all smaller indices $s < \text{sat}_{\mathcal{E}^*}(i^*, \mathbf{o}^*)$ it follows from $\text{sat}_{\mathcal{E}}(i^*, \mathbf{o}) > \text{sat}_{\mathcal{E}}(i^*, \mathbf{o}^*)$ that

$$i^* \notin \left(N_s^{\mathcal{E}}(\mathbf{o}) \cup N_s^{\mathcal{E}^*}(\mathbf{o}) \cup N_s^{\mathcal{E}}(\mathbf{o}^*) \cup N_s^{\mathcal{E}^*}(\mathbf{o}^*) \right).$$

Therefore, it follows from Equation (4.5) that $|N_s^{\mathcal{E}}(\mathbf{o})| = |N_s^{\mathcal{E}^*}(\mathbf{o})|$ and $|N_s^{\mathcal{E}}(\mathbf{o}^*)| = |N_s^{\mathcal{E}^*}(\mathbf{o}^*)|$. Hence, $s^* < \text{sat}_{\mathcal{E}^*}(i^*, \mathbf{o}^*)$ would be a contradiction to the assumption that $\mathbf{o}$ is the opt-leximin outcome of $\mathcal{E}$ and hence preferred by opt-leximin to $\mathbf{o}^*$ in $\mathcal{E}$.

Thus, we know that $|N_s^{\mathcal{E}^*}(\mathbf{o})| = |N_s^{\mathcal{E}^*}(\mathbf{o}^*)|$ for all $s < \text{sat}_{\mathcal{E}^*}(i^*, \mathbf{o}^*) \leq s^*$ and

$$|N_{\text{sat}_{\mathcal{E}^*}(i^*, \mathbf{o}^*)}^{\mathcal{E}^*}(\mathbf{o}^*)| \leq |N_{\text{sat}_{\mathcal{E}^*}(i^*, \mathbf{o}^*)}^{\mathcal{E}^*}(\mathbf{o})|.$$

It follows also that, for all $s < \text{sat}_{\mathcal{E}^*}(i^*, o^*) \leq s^*$,

$$|N_s^{\mathcal{E}}(o^*)| = |N_s^{\mathcal{E}^*}(o^*)| = |N_s^{\mathcal{E}^*}(o)| = |N_s^{\mathcal{E}}(o)|.$$

Next, it follows from Equation (4.4) that

$$i^* \in N_{\text{sat}_{\mathcal{E}^*}(i^*, o^*)}^{\mathcal{E}^*}(o^*) \quad \text{but} \quad i^* \notin N_{\text{sat}_{\mathcal{E}^*}(i^*, o^*)}^{\mathcal{E}}(o^*).$$

Moreover, because we assumed $\text{sat}_{\mathcal{E}}(i^*, o) > \text{sat}_{\mathcal{E}}(i^*, o^*)$, we get

$$i^* \notin N_{\text{sat}_{\mathcal{E}^*}(i^*, o^*)}^{\mathcal{E}^*}(o) \quad \text{and} \quad i^* \notin N_{\text{sat}_{\mathcal{E}^*}(i^*, o^*)}^{\mathcal{E}}(o).$$

With the above two facts, together with Equation (4.5), we can finally infer

$$|N_{\text{sat}_{\mathcal{E}^*}(i^*, o^*)}^{\mathcal{E}}(o^*)| + 1 = |N_{\text{sat}_{\mathcal{E}^*}(i^*, o^*)}^{\mathcal{E}^*}(o^*)| \leq |N_{\text{sat}_{\mathcal{E}^*}(i^*, o^*)}^{\mathcal{E}^*}(o)| = |N_{\text{sat}_{\mathcal{E}^*}(i^*, o^*)}^{\mathcal{E}}(o)|.$$

However, this means that $(o^*, o) \in \mathcal{L}(\mathcal{E})$, which is a contradiction to the assumption that o is the outcome of $\mathcal{E}$. $\square$

It remains an open problem to generalise the latter result to other **opt-OWA** rules.

We conclude this section by noting that Propositions 4.17 and 4.18 do not hold for generalised free-riding (Definition 4.8). Indeed, consider a situation with two voters and two binary issues. The first voter has preferences $A_1(1) = \{a, b\}$ and $A_2(1) = \{a\}$. The second voter has preferences $A_1(2) = A_2(2) = \{b\}$. Here, most reasonable rules (with alphabetical tiebreaking) pick (b, a) . If voter 1 changes her ballot in the first issue to $\{a\}$, the outcome (again, for most reasonable rules) becomes (a, b) . This is an instance of generalised free-riding, but the final outcome is worse for voter 1.

4.7.3 Complexity of Free-Riding with Non-Sequential Rules

Propositions 4.17 and 4.18 seem to indicate that free-riding is safe in this setting, at least for many appealing rules. However, we will now show that free-riding is hard to perform *in the first place*. That is, it is hard to decide whether free-riding is possible even without checking whether it would raise one's satisfaction. Hence, a manipulator has no computational reason to restrict herself to free-riding instead of considering unrestricted manipulation.

We begin with a more fundamental problem: outcome determination. Although this problem is not identical to free-riding, it provides a natural starting point for assessing its computational difficulty and is interesting in its own right. A hypothetical free-rider must first identify the “truthful” outcome in order to determine which current winners could be affected, and then compare the satisfaction yielded by that truthful outcome with the satisfaction obtained by free-riding. Thus, outcome determination is a natural prerequisite for one common route to free-riding. If outcome determination is

computationally intractable, then it already provides a first barrier to free-riding.¹⁶ We therefore study the following problem:

Definition 4.9. *F-OUTCOME-DETERMINATION* is the following decision problem: “Given an election $\mathcal{E} = (C, N, A)$, an issue $\ell \in [k]$, and a candidate $o \in C_\ell$, does o win in $\mathcal{E}$ on issue ℓ under F ?”

The above is formulated as a decision problem. Clearly, if the corresponding optimisation problem of computing the full outcome o was easy, then the above decision problem would be easy too, and *vice versa*. Moreover, recall that, as stated in Sections 2.2.5 and 2.2.6, we assume that all weight functions ω and $\omega^{(n)}$ are computable in polynomial time.

As stated several times, it is clear that outcome determination is easy for sequential rules. Now we show that it is hard for opt-Thiele and opt-OWA rules.

Theorem 4.19. *F-OUTCOME-DETERMINATION* is NP-hard for every opt-Thiele rule distinct from the utilitarian rule.

To prove the above we reduce from CUBICVERTEXCOVER [Alimonti and Kann, 2000], similarly as in the proof of Theorem 3 by Skowron et al. [2016]. Their result is similar to ours: they show NP-hardness for a large class of Thiele rules in the ABC voting model. We note that our result is not a consequence of theirs, since ABC voting is not a special case of our model.

Moreover, we note that the NP-hardness of opt-PAV was already shown by Brill et al. [2022, Theorem 5.1], as the party-approval model they study is a special case of ours. This implies that the hardness of opt-PAV holds even in the special case where the alternatives and preferences of the voters are constant across issues. Our result, by contrast, requires different preferences and candidates for each issue, but holds for all opt-Thiele rules except, naturally, the utilitarian rule.

Proof of Theorem 4.19. Fix an opt-Thiele rule defined by ω which is distinct from the utilitarian rule. Let this rule be F . We show hardness by a reduction from CUBICVERTEXCOVER, a variant of VERTEXCOVER where every node has a degree exactly three [Alimonti and Kann, 2000]. We are given an undirected graph and a number $k \in \mathbb{N}$, and we ask whether there is a set of at most k nodes such that each edge in the graph has one of its endpoints in this set. In the following, let s be the smallest s such that $\omega(s) > \omega(s + 1)$ (such an s must exist by virtue of F not being the utilitarian rule).

¹⁶We stress that the two problems are different. Even if outcome determination was easy, free-riding could still be hard, since the number of ballot vectors a voter can submit is exponential in the size of the input. Conversely, even if outcome determination was hard, free-riding could still be easy in principle; for example, this could happen for rules under which free-riding is possible only on very specific and easily recognizable inputs. Nonetheless, we believe that outcome determination informs all the subsequent analysis.

Consider an instance (G, k) of CUBICVERTEXCOVER. Again, k is a natural number and $G = (V, E)$ is an undirected graph with n nodes and m edges where every node has a degree 3. We assume without loss of generality that $k < m$. Formally, an edge is an unordered pair of nodes, i.e., a subset $e \subseteq V$ with $|e| = 2$. We will construct an instance $(\mathcal{E}, s + k, c_{d_1})$ of F -OUTCOME-DETERMINATION, where $\mathcal{E} = (C, N, A)$ is an election with $s + k$ issues and $2m$ voters. Observe in particular that s does not depend on (G, k) .

We construct the instance as follows. We have one voter v_e for every edge $e \in E$, plus m extra dummy voters $\{d_1, \dots, d_m\}$. In the first $s - 1$ issues, there are two candidates, c and c' , and all voters approve of both. In the next k issues, there is one candidate c_η for every node $\eta \in V$, plus one candidate c_{d_i} for every dummy voter d_i . In all of these issues, every edge voter v_e approves of the two candidates c_η and $c_{\eta'}$ such that $e = \{\eta, \eta'\}$. Furthermore, every dummy voter d_i approves of only c_{d_i} . Finally, in the last issue, there is one candidate c_i for every voter $i \in N$, and every such voter only approves of c_i . Visually:

Voter	Issues $[s - 1]$	Issues $\{s, \dots, s + k - 1\}$	Issue $s + k$
v_e for $e = \{\eta, \eta'\} \in E$	$\{c, c'\}$	$\{c_\eta, c_{\eta'}\}$	$\{c_{v_e}\}$
d_i for $i \in [m]$	$\{c, c'\}$	$\{c_{d_i}\}$	$\{c_{d_i}\}$

To deal with ties, we assume that each issue ℓ is associated with a total ordering $\succ_\ell$ such that,

- if $\ell \in \{s, \dots, s + k - 1\}$, then node-candidates are preferred over other candidates, and $c_{d_n} \succ_\ell c_{d_{n-1}} \succ_\ell \dots \succ_\ell c_{d_1}$.
- If $\ell = s + k$, then all candidates c_{v_e} (with $e \in E$) are preferred over other candidates, and $c_{d_1} \succ_\ell c_{d_2} \succ_\ell \dots \succ_\ell c_{d_n}$.

As always, we compare outcomes lexicographically according to the above ordering. Given an outcome $\mathbf{o}$, we define

$$V[\mathbf{o}] = \{\eta \in V : c_\eta \in \mathbf{o}\}$$

as the set of nodes η that have a “node-candidate” c_η included in $\mathbf{o}$.

We want to show that (G, k) is a yes-instance if and only if $(\mathcal{E}, s + k, c_{d_1})$ is. First, note that all voters have a satisfaction of at least $s - 1$ (because of the first $s - 1$ issues). Next, let us show the following useful fact:

Fact 1. *Given $F(\mathcal{E}) = \mathbf{o}$, candidate c_{d_1} wins in the last issue if and only if $V[\mathbf{o}]$ corresponds to a vertex cover of G .*

Proof of Fact 1. First, assume $V[\mathbf{o}]$ is not a vertex cover. Assume also that c_{d_1} wins in the last issue. Then there is one “edge-voter” v_e with satisfaction $s - 1$. But then we can swap c_{d_1} with c_{v_e} in the last issue without decreasing the Thiele score, and this swap is

preferred by the tiebreaking. This is a contradiction, and hence c_{d_1} does not win in the last issue.

Next, assume $V[o]$ is a vertex cover. Then, every edge-voter v_e has a satisfaction of at least s . Moreover, as there are m dummy voters and $k < m$, at least two dummy voters will be satisfied with no issue in $\{s, \dots, s+k-1\}$ (as in at least one issue a node-candidate wins). Then, if in the last issue some candidate c_{v_e} (with $e \in E$) wins, we get a contradiction: since $\omega(s) > \omega(s+1)$, by the previous arguments, we could increase the Thiele score by replacing c_{v_e} with a candidate c_{d_i} for some dummy voter d_i that wins no issue in $\{s, \dots, s+k-1\}$.

Next, suppose that some candidate c_{d_i} with $i > 1$ wins in the last issue. First, observe that c_{d_i} cannot occur in the previous issues, as otherwise we could swap this candidate with a candidate corresponding to a dummy voter with satisfaction $s-1$ (which must exist as previously argued) and strictly increase the Thiele score. If c_{d_1} is also never selected, then we can replace c_{d_i} in the last issue with c_{d_1} without changing the Thiele score. Otherwise, by the symmetries in the election, we can swap every occurrence of c_{d_i} and c_{d_1} in o and obtain an outcome with the same Thiele score. Both swaps are preferred by the tiebreaking, showing that c_{d_1} must be selected in the last issue. This concludes the proof of Fact 1. $\diamond$

Fix $F(\mathcal{E}) = o$. It remains to be shown that if there is a vertex cover of G with at most k vertices, then $V[o]$ is a vertex cover for the winning outcome o . Together with Fact 1, this is enough to show our claim. Then, assume towards a contradiction that a vertex cover of G of size at most k exists and that $V[o]$ is not a vertex cover.

If $V[o]$ is not a vertex cover, we claim that for every issue $\ell \in \{s, \dots, s+k-1\}$ the winner must be a node-candidate. Assume towards a contradiction that there is an issue $\ell \in \{s, \dots, s+k-1\}$ where a candidate c_{d_i} wins. Let $e = \{\eta, \eta'\}$ be an edge that is not covered by $V[o]$. There are two cases. First, assume that c_{v_e} wins in the last issue. Then we can change the outcome so that c_η wins in issue ℓ and c_{d_i} wins in the last issue. This cannot decrease the Thiele score, and is preferred by the tiebreaking. Second, assume that c_{v_e} does not win in the last issue. Then, v_e has satisfaction $s-1$. But then if we change the outcome so that c_η wins in issue ℓ we cannot decrease the Thiele score, and this is preferred by the tiebreaking. In both cases we have reached a contradiction with the fact that o is the outcome.

Now, let $o_{[-1]}$ be the partial outcome obtained by truncating the last issue from o . Moreover, let

$$E_o^s = \left\{ v_e \in N : e \in E \text{ and } \text{sat}_{\mathcal{E}}(v_e, o_{[-1]}) \geq s \right\}$$

be the set of edge-voters with satisfaction at least s with respect to $o_{[-1]}$. We define $E_o^{s+1} \subseteq E_o^s$ analogously. We observe

$$\sum_{v_e \in E_o^s} \left(\text{sat}_{\mathcal{E}}(v_e, o_{[-1]}) - (s-1) \right) = 3k.$$

The above holds because every time we select a node-candidate, we give additional satisfaction to its edges. Since every node has 3 edges, this quantity must equal $3k$.

Hence,

$$\sum_{v_e \in E_o^s} \left(\text{sat}_{\mathcal{E}}(v_e, \mathbf{o}_{[-1]}) - (s-1) \right) = |E_o^s| + \sum_{v_e \in E_o^{s+1}} \left(\text{sat}_{\mathcal{E}}(v_e, \mathbf{o}_{[-1]}) - s \right) = 3k.$$

This holds because voters in $E_o^s \setminus E_o^{s+1}$ have $\text{sat}_{\mathcal{E}}(v_e, \mathbf{o}_{[-1]}) - s = 0$. We can finally compute the Thiele score of $\mathbf{o}$:

$$\sum_{i \in N} \left(\sum_{r \in [s-1]} \omega(r) \right) + \omega(s) + |E_o^s| \omega(s) + \underbrace{\sum_{v_e \in E_o^{s+1}} \sum_{r=s+1}^{\text{sat}_{\mathcal{E}}(v_e, \mathbf{o}_{[-1]})} \omega(r)}_{3k - |E_o^s| \text{ addends}}.$$

Let us break this down. The first summation comes from the fact that all voters are satisfied with the first $s-1$ issues. The $\omega(s)$ contribution comes from the fact that in the last issue some voter that is never satisfied in issues from s to $s+k-1$ must win, since $\omega(s) > \omega(s+1)$. The remaining contributions come from the definition of E_o^s and E_o^{s+1} and our previous arguments.

Since $\omega(s) > \omega(s+1)$, this function is maximised by maximising $|E_o^s|$. As a vertex cover exists, we know that we can reach $|E_o^s| = m$. Hence, the outcome maximising the Thiele score must do so, which means that $V[\mathbf{o}]$ must be a vertex cover. We have reached the desired contradiction, completing the proof. $\square$

Theorem 4.20. *F-OUTCOME-DETERMINATION* is NP-hard for every opt-OWA rule defined by $\omega^{(n)}$ such that, for some $n^* \in \mathbb{N}$, $\omega^{(n)}(1) > \omega^{(n)}(n)$ holds for every $n \in \mathbb{N} \setminus [n^*]$.

Note that our result is related to Theorems 2 and 3 by [Amanatidis et al. \[2015\]](#). While their results hold even for the special case of binary elections, they only focus on a special subclass of OWA rules, i.e., given by vectors of the form $(1, \dots, 1, 0, \dots, 0)$.

In proving Theorem 4.20, as well as other results concerning opt-OWA rules, we will make use of the following useful lemma, which states that if everyone's satisfaction weakly increases, then the OWA score weakly increases.

Lemma 4.21. Consider two outcomes $\mathbf{o}$ and $\mathbf{o}'$ for $\mathcal{E}$. If $\text{sat}_{\mathcal{E}}(i, \mathbf{o}) \geq \text{sat}_{\mathcal{E}}(i, \mathbf{o}')$ for all $i \in [n]$, then in $\mathcal{E}$ outcome $\mathbf{o}$ has OWA score at least as large as $\mathbf{o}'$.

Proof. Let $\mathbf{s} = \mathbf{s}(\mathcal{E}, \mathbf{o})$ and $\mathbf{s}' = \mathbf{s}(\mathcal{E}, \mathbf{o}')$ denote the sorted satisfaction vectors. Since $\text{sat}_{\mathcal{E}}(i, \mathbf{o}) \geq \text{sat}_{\mathcal{E}}(i, \mathbf{o}')$ for all $i \in [n]$, we have $s_j \geq s'_j$ for all $j \in [n]$. Therefore,

$$\sum_{j \in [n]} \omega^{(n)}(j) s_j \geq \sum_{j \in [n]} \omega^{(n)}(j) s'_j.$$

Hence $\mathbf{o}$ has OWA score at least as large as $\mathbf{o}'$. $\square$

Proof of Theorem 4.20. Fix a rule F satisfying the condition of the theorem. We again show hardness by a reduction from CUBICVERTEXCOVER. Consider an instance (G, k) of this problem. Here, $G = (V, E)$ is a graph with n nodes and m edges where each node has a degree exactly three, and we assume without loss of generality that $k < n$. We construct an instance of F -OUTCOME-DETERMINATION with $k + 1$ issues and $3m$ voters.

We restrict ourselves to instances satisfying $3m > n^*$, where n^* is the constant from the theorem statement which depends only on F . This restriction does not affect the NP-hardness of CUBICVERTEXCOVER, because the condition $3m \leq n^*$ implies that both m and n are bounded by constants. Hence, only finitely many instances occur in the excluded range.

In the following, since the number of voters is fixed, we will simply refer to $\omega^{(3m)}$ as ω . As $\omega(1) > \omega(3m)$, there must be some $p \in [3m - 1]$ such that $\omega(1) = \omega(p) > \omega(p + 1)$. There are two cases: either $p \in [2m]$ or $2m < p < 3m$. We treat these cases separately.

First case. We construct an instance $(\mathcal{E}, k + 1, c_{d_1})$ of F -OUTCOME-DETERMINATION. Here, we have one voter v_e for each edge $e \in E$, and two sets of dummy voters, $\{d_1, \dots, d_p\}$ and $\{w_1, \dots, w_{2m-p}\}$. In the first k issues, there is one candidate c_η for each node $\eta \in V$, plus one dummy candidate c_{d_i} for each dummy voter d_i . Here, each edge-voter v_e with $e = \{\eta, \eta'\}$ approves of the two node-candidates c_η and $c_{\eta'}$. Moreover, each dummy voter d_i approves only of dummy candidate c_{d_i} , and all dummy voters w_i approve of all candidates. In the last issue, there is one candidate c_i for all voters $i \in N \setminus \{w_1, \dots, w_{2m-p}\}$, and every such i only approves of c_i . Finally, here, all voters in $\{w_1, \dots, w_{2m-p}\}$ approve of all candidates. Visually:

Voter	Issues $[k]$	Issue $k + 1$
v_e for $e = \{\eta, \eta'\} \in E$	$\{c_\eta, c_{\eta'}\}$	$\{c_{v_e}\}$
d_i for $i \in [p]$	$\{c_{d_i}\}$	$\{c_{d_i}\}$
w_i for $i \in [2m - p]$	all candidates	all candidates

We use the same tiebreaking mechanism as in the proof of Theorem 4.19, verbatim.

We want to show that (G, k) is a yes-instance if and only if $(\mathcal{E}, k + 1, c_{d_1})$ is. Suppose that there exists a vertex cover for G with size at most k .

First, observe that, regardless of the outcome, the last $2m - p$ entries of the satisfaction vector will be $k + 1$ (as the w_i voters always approve all candidates). We claim that any optimal outcome gives satisfaction at least 1 to every edge-voter and to at least one dummy voter d_i . Assume the contrary holds. There are two cases.

- Some dummy voter d_i has satisfaction at least one. Then, some edge-voter v_e has satisfaction zero. Take some issue where d_i (and only her) approves of the outcome, and replace the winner of this issue by some candidate approved by v_e . This cannot decrease the opt-OWA score, as we are transferring satisfaction from a more-satisfied voter to a least-satisfied voter. Moreover, this replacement would be preferred by the tiebreaking.

- All dummy voters d_i have satisfaction zero, and hence p many voters have satisfaction zero. Then, the total satisfaction of the edge-voters and dummy voters d_i will be $3k + 1$ (since all node-candidates are approved by three edge-voters each, plus the last issue), irrespectively of how this satisfaction is distributed. But then we can reach a higher opt-OWA score by giving satisfaction one to every edge-voter within the first k issues—and this is possible, for a vertex cover exists—and by selecting some candidate c_{d_i} in the last issue (and in that issue only). This would assign zero satisfaction to at most $p - 1$ voters while keeping the total satisfaction (ignoring voters w_i) at $3k + 1$. Since $\omega(p) > \omega(p + 1)$, this is an improvement.

Next, we argue that this implies that some dummy candidate c_{d_i} wins in the last issue. Indeed, if c_{v_e} (with $e = \{\eta, \eta'\}$) wins in the last issue and c_{d_i} in some issue $\ell \in [k]$, we can let c_{d_i} win in the final issue and some candidate approved by v_e win in issue ℓ . This cannot decrease the score (Lemma 4.21) and is preferred by the tiebreaking. Finally, by tiebreaking order, together with the symmetry of the dummy voters d_i , we get that in the last issue c_{d_1} wins.

For the reverse implication, suppose that there exists no vertex cover for G with size at most k . Then there is an edge-voter v_e that does not win any issue in $[k]$ (otherwise, some vertex cover would exist). By the tiebreaking rule, c_{v_e} prevents c_{d_1} from winning the last issue.

Second case. Here, we have $\omega(1) = \omega(p) = 1 > \omega(p + 1)$. We construct an instance $(\mathcal{E}, k + 1, c)$ of F -OUTCOME-DETERMINATION. Here, we have one voter v_e for each edge $e \in E$, and three sets of dummy voters: $\{d_1, \dots, d_{p-2m+1}\}$, $\{a_1, \dots, a_m\}$, and $\{w_1, \dots, w_{3m-p-1}\}$. In the first k issues, there is one candidate c_η for each node $\eta \in V$, plus one dummy candidate c_{d_i} for each dummy voter d_i . Here, each edge-voter v_e approves of every node-candidate c_η where $\eta \notin e$. Furthermore, each dummy voter d_i approves only of dummy candidate c_{d_i} . Every other dummy voter approves of all candidates. In the last issue, there are two candidates c and c' . All dummy candidates d_i and w_i approve both candidates, every edge-voter approves only of c , while every dummy voter a_i approves only of c' . Visually:

Voter	Issues $[k]$	Issue $k + 1$
v_e for $e \in E$	$\{c_\eta : \eta \in V \setminus e\}$	$\{c\}$
d_i for $i \in [p - 2m + 1]$	$\{c_{d_i}\}$	$\{c, c'\}$
a_i for $i \in [m]$	all candidates	$\{c'\}$
w_i for $i \in [3m - p - 1]$	all candidates	$\{c, c'\}$

We use the same tiebreaking rule as in the previous case, except that in the last issue c loses to c' .

First, observe that, since all dummy voters w_i approve of all candidates, the last $3m - p - 1$ entries of the sorted satisfaction vector will always be $k + 1$. The remaining $p + 1$ entries are the only ones that determine the outcome. Notice that all these entries

are associated with the same weight, $\omega(p)$, save for the last entry where the associated weight is $\omega(p + 1)$. Moreover, no dummy candidate c_{d_i} can be selected. If that was the case, we could swap this dummy candidate with some node-candidate. This swap would increase the satisfaction of more than one voter, while only the dummy voter d_i would lose one point of satisfaction, leading to a higher opt-OWA score.

We know that the first $p - 2m + 1$ positions of the sorted satisfaction vector will be 1 (given by the dummy voters d_i). The last $3m - p - 1$ entries will be $k + 1$, given by the w_i voters. The “central part” is occupied by the edge-voters and the dummy voters a_i . It is easy to see that the total satisfaction of these two groups is the same for every outcome, namely $k(m - 3) + km + m$.

Indeed, over the first k issues, the edge-voters receive total satisfaction exactly $k(m - 3)$, since each chosen node-candidate is approved by exactly $m - 3$ edge-voters. In the same issues, a_i voters receive always km total satisfaction. The extra m comes from the last issue. It follows that to optimise the opt-OWA score we must minimise the satisfaction of the voter in position $p + 1$ of the sorted satisfaction vector. This satisfaction is either k or $k + 1$, since k is the minimum satisfaction of the voters a_i and there are $k + 1$ issues.

Suppose that a vertex cover of size at most k exists. Then we can maximize the score by selecting c in the final issue and ensuring that no edge-voter wins more than $k - 1$ of the first k issues, which is possible because a vertex cover of size k exists. If no vertex cover of size at most k exists, then choosing either c or c' in the final issue yields the same score, regardless of how the first k issues are resolved.

This concludes the second case and, with it, the proof. $\square$

Since it is hard to know whether a candidate wins in a given issue (and thus whether we can free-ride on its victory) one might be tempted to conclude that free-riding is unfeasible with non-sequential rules. However, sometimes the fact that a certain candidate wins can still be known (or guessed). For example, when a candidate receives an extremely disproportionate support, or when some external source (i.e., a polling agency having the computational power to solve F -OUTCOME-DETERMINATION) communicates the projected winners. In this case, the manipulator would need to solve a slightly different, potentially easier, problem: *Given that some candidate that I approve of wins in this specific issue, can I deviate from my honest approval ballot without making this candidate lose?* That is, *can I free-ride?* Motivated by this, we study the following problem:¹⁷

Definition 4.10. F -FREE-RIDING-RECOGNITION is the following decision problem: “Given an election $\mathcal{E} = (C, N, A)$, an issue $\ell \in [k]$, a voter $i \in N$, and a candidate $o \in C_\ell$ that wins in issue ℓ under F in $\mathcal{E}$, can i free-ride on issue ℓ ?”

We define GENERALISED- F -FREE-RIDING-RECOGNITION analogously. Luckily, the picture does not change: this problem is still computationally hard for essentially

¹⁷Technically speaking, the following is a *promise problem* [Goldreich, 2006], because we are “promising” that the input is well-formed, i.e., that o truly wins in issue ℓ . This is not relevant for us, as all the machinery of reductions stays the same.

the same families of rules. The next two reductions build on the reductions used in Theorems 4.19 and 4.20.

Theorem 4.22. *F-FREE-RIDING-RECOGNITION and GENERALISED-F-FREE-RIDING-RECOGNITION are NP-hard for every opt-Thiele rule distinct from the utilitarian rule.*

Proof. We again show the hardness of *F-FREE-RIDING-RECOGNITION* by a reduction from *CUBICVERTEXCOVER*. Let s be the smallest s where $\omega(s) > \omega(s+1)$ holds, and consider an instance (G, k) of *CUBICVERTEXCOVER*. Here, $k \in [m-1]$, and $G = (V, E)$ is an undirected graph with n nodes and m edges where every node has degree of 3. We construct an instance $(\mathcal{E}, s+k, d_2, c_{d_1})$ of *F-FREE-RIDING-RECOGNITION*, where $\mathcal{E}$ is an election with $s+k$ issues and $2m$ voters. We use a construction similar to the one in the proof of Theorem 4.19, except for the fact that, on issue $s+k$, voter d_2 approves only of c_{d_1} .

First, suppose that (G, k) is a yes-instance. By the same arguments used in the proof of Theorem 4.19, we know that c_{d_1} wins in issue $s+k$ (the fact that now d_2 also approves of it is irrelevant). Moreover, if d_2 votes for c_{d_2} in the last issue, we obtain the same election constructed in the proof of Theorem 4.19. We have already shown that under these conditions c_{d_1} wins the final issue: hence, d_2 can free-ride.

Now, suppose that (G, k) is a no-instance. Let us first show that c_{d_1} is still selected on issue $s+k$. Clearly, at least one edge-voter does not win any issue in $\{s, \dots, s+k\}$ (otherwise, a vertex cover would exist). Towards a contradiction, suppose that in $F(\mathcal{E})$ some candidate c_{d_i} is winning some issue in $\ell \in \{s, \dots, s+k-1\}$. Then, at least one edge-voter must have satisfaction $s-1$ (for if all edge-voters were to have satisfaction at least s , we could cover all but one edge with $k-1$ nodes). So let v_e be some edge-voter with satisfaction $s-1$ in $F(\mathcal{E})$. If we select c_η (for some $\eta \in e$) in ℓ instead of c_{d_i} , we would increase the total score by at least $\omega(s)$ (contributed by v_e) and decrease it by $\omega(r)$ for some $r \geq s$ (contributed by d_i). Since $\omega(r) \leq \omega(s)$, this new outcome cannot have a lower score than $F(\mathcal{E})$, and would be preferred to it by the tiebreaking. Contradiction: we conclude that no candidate c_{d_i} can win in issues in $\{s, \dots, s+k-1\}$. This implies, in particular, that neither d_1 nor d_2 are satisfied with any of those issues. Therefore, selecting c_{d_1} for issue $s+k$ contributes $2\omega(s)$ to the total score, whereas selecting any other candidate can contribute at most $\omega(s)$. Since $\omega(s) > \omega(s+1) \geq 0$, c_{d_1} must win in the final issue.

It remains to be shown that d_2 cannot free-ride. Suppose that d_2 does not approve of c_{d_1} . Consider some edge-voter v_e that never wins in $\{s, \dots, s+k-1\}$ (which, as argued above, must exist). Choosing c_{v_e} or c_{d_1} in the final issue contributes at least $\omega(s)$ to the total score, whereas choosing c_{d_1} can contribute at most $\omega(s)$. By tiebreaking, c_{d_1} does not win in the final issue, that is, d_2 cannot free-ride in this issue.

To conclude, observe that the same construction can be used to show hardness for *GENERALISED-F-FREE-RIDING-RECOGNITION*. Indeed, here d_2 only approves of c_{d_1} , and hence free-riding and generalised free-riding coincide. $\square$

Hardness for **OWA** rules is split in two theorems, concerning two (overlapping) classes; essentially all reasonable **OWA** rules fall into one of the two. For example, opt-leximin only appears in the first class, whereas the opt-egalitarian rule only in the second one.

Theorem 4.23. *F-FREE-RIDING-RECOGNITION and GENERALISED-F-FREE-RIDING-RECOGNITION are NP-hard for every opt-OWA rule defined by $\omega^{(n)}$ for which there is a $n^* \geq 3$ such that, for every $m \in \mathbb{N}$, there is a number of voters $n \in \mathbb{N}$ with $3m \leq n \leq n^* \cdot m$ such that $\omega^{(n)}(1) > \omega^{(n)}(n)$ and $\omega^{(n)}(3m) > 0$.*

The n^* constant in the statement of the theorem is a technical condition. It ensures that, in the following reduction, the number of voters does not blow up exponentially.

Proof of Theorem 4.23. We show the hardness of *F-FREE-RIDING-RECOGNITION* by a reduction from *CUBICVERTEXCOVER*. Consider an instance (G, k) of this problem. Here, $G = (V, E)$ is a graph with n nodes and m edges where each node has a degree of exactly three, and $k \in \mathbb{N}$. By the condition of the theorem, we know there is an $\ell \geq 3m$ (polynomial in the size of m) such that $\omega^{(\ell)} = (\omega^{(\ell)}(1), \dots, \omega^{(\ell)}(\ell))$ contains at least $3m$ nonzero entries and $\omega^{(\ell)}(1) > \omega^{(\ell)}(\ell)$.

We will construct an instance of *F-FREE-RIDING-RECOGNITION* with $(k + 1)$ issues and ℓ voters. Since ℓ is fixed, we will simply indicate $\omega^{(\ell)}$ as ω . Moreover, as $\omega(1) > \omega(\ell)$ and $\omega(3m) > 0$, there are two cases:

1. There is a $p \in [2m]$ such that $\omega(1) = \omega(p) > \omega(p + 1)$ and $\omega(p + m) > 0$, or
2. There is a $p \in \{2m + 1, \dots, \ell - 1\}$ such that $\omega(1) = \omega(p) > \omega(p + 1)$.

We treat them separately.

First case. We construct an instance $(\mathcal{E}, k + 1, v_{e^*}, c_{d_1})$ of *F-FREE-RIDING-RECOGNITION* (here, $e^* \in E$ is some edge, it does not matter which). The construction is similar to the one shown in the first case of the proof for Theorem 4.20. However, here, in issue $k + 1$ voter v_{e^*} approves only of c_{d_1} , and we have $\ell - m - p$ dummy voters w_i instead of $2m - p$. The latter change makes no difference in our construction: in both cases, the w_i voters force all entries after the $(p + 1)$ -th one to be exactly $k + 1$.

First, note that $(\mathcal{E}, k + 1, v_{e^*}, c_{d_1})$ is indeed a well-formed instance of *F-FREE-RIDING-RECOGNITION*, as surely c_{d_1} wins in issue $k + 1$. If (G, k) is a yes-instance then we have already shown that this candidate wins in the proof of Theorem 4.20, and here it is only receiving increased support. If it is a no-instance, then one can argue (as done in the proof of Theorem 4.20, first case) that d_1 cannot be satisfied with any issue in $[k]$. Thus, c_{d_1} will be supported by one voter that never won in the first k issues (namely, d_1), as well as by v_{e^*} . Hence c_{d_1} must be selected in the last issue: it is the only candidate that is supported by two voters (excluding the w_i voters), and one of these two voters is not satisfied with any other issue. Since $\omega(p + m) > 0$ and since the edge-voters together

with the dummy voters d_i occupy at most the first $p + m$ positions of the satisfaction vector, any other candidate would yield a strictly lower opt-OWA score.

Now, if (G, k) is a yes-instance of CUBICVERTEXCOVER, then v_{e^*} can free-ride in the last issue: if she votes for her voter-candidate, then we have an election essentially identical to the one constructed in the first case of the proof of Theorem 4.20, and we have already shown there that c_{d_1} wins if (G, k) has a vertex cover.

If (G, k) is a no-instance, we show that v_{e^*} cannot free-ride in the last issue. There are two cases: either v_{e^*} is satisfied with some issue in $[k]$ or not. If she is, there will be at least one voter v_e (with $e \in E \setminus \{e^*\}$) that is not. Hence selecting c_{v_e} in the last issue yields at least the same score as selecting c_{d_1} , since v_{e^*} does not approve of the latter when she free-rides. Hence, by tiebreaking, c_{d_1} is not the winner in the final issue. If, on the other hand, v_{e^*} is never satisfied in issues $[k]$, there are again two cases: either v_{e^*} approves in issue $k + 1$ some c_{v_e} (where $e \in E$) or some dummy candidate c_{d_i} (with $i > 1$). In the former case, selecting c_{v_e} would yield at least the same opt-OWA score as selecting c_{d_1} : by tiebreaking, c_{d_1} loses. In the latter case, assume towards a contradiction that c_{d_1} is selected in the last issue. Hence, there are edge-voters with zero satisfaction. Surely, then, no candidate c_{d_j} can win within the first $[k]$ issues, because replacing this candidate with a node-candidate approved by an edge-voter with zero satisfaction would not decrease the opt-OWA score and be preferred by the tiebreaking. Hence, replacing c_{d_1} with c_{d_i} in the last issue would increase the total opt-OWA score by $\omega(p)$. This holds because c_{d_i} is approved by two voters that are never satisfied in issues $[k]$, while c_{d_1} only by one. Contradiction: c_{d_1} is not selected in the final issue. Hence no free-riding is possible in the final issue.

Second case. We construct another instance $(\mathcal{E}, k + 1, a_1, c)$ of the problem F -FREE-RIDING-RECOGNITION. The construction is similar to the one shown in the second case of the proof of Theorem 4.20. However, here, in issue $k + 1$ voter a_1 approves only of c , and we have $\ell - p - 1$ dummy voters w_i instead of $3m - p - 1$. The latter change makes no difference in our construction.

First, note that $(\mathcal{E}, k + 1, a_1, c)$ is a legal instance of F -FREE-RIDING-RECOGNITION. Indeed, for any outcome, the final $\ell - p - 1$ positions in the sorted satisfaction vector are all $k + 1$ (due to the w_i voters approving all candidates). Additionally, in the final issue, ignoring the voters that approve all candidates, c' receives the support of $m - 1$ voters, whereas c receives the support of $m + 1$ voters. Hence, selecting c over c' in the last issue increases the opt-OWA score of at least

$$\omega(p + 1) + m\omega(p) - (m - 1)(\omega(p)) > 0.$$

Thus, c wins in the last issue.

If a_1 free-rides in $k + 1$, she must vote only for c' . Here, we obtain a construction essentially identical to the one shown in the second case of the proof Theorem 4.20, and we have already shown there that c wins if and only if (G, k) is a yes-instance.

This concludes the case analysis. Finally, observe that, in both cases, the same construction can be used to show hardness for GENERALISED- F -FREE-RIDING-RECOGNITION. Indeed, here the manipulator only approves of one alternative, and hence free-riding and generalised free-riding coincide. This concludes the proof. $\square$

The next theorem shows CONP-hardness rather than NP-hardness. The class CONP consists of all decision problems whose complements are in NP [Arora and Barak, 2009], that is,

$$\text{CONP} = \{\overline{\text{PROB}} : \text{PROB} \in \text{NP}\}.$$

Here, the bar denotes complementation, meaning that yes-instances become no-instances and *vice versa*. We note that CONP-hardness is also regarded as evidence of intractability.

Theorem 4.24. F -FREE-RIDING-RECOGNITION and GENERALISED- F -FREE-RIDING-RECOGNITION are CONP-hard for every opt-OWA rule defined by $\omega^{(n)}$ for which there is a $n^* \geq 2$ such that, for every $m \in \mathbb{N}$, there is a number of voters $n \in \mathbb{N}$ with $m < n \leq n^* \cdot m$ such that $\omega^{(n)}(1) > \omega^{(n)}(n)$ and $\omega^{(n)}(n - m + 1) = 0$.

Proof. We show the hardness of F -FREE-RIDING-RECOGNITION by a reduction from VERTEXCOVER [Garey and Johnson, 1979]. Hence, here, we do not assume that every node has degree 3. Consider an instance of this problem, (G, k) , where G has n nodes and m edges. By the condition of the theorem, we know there is an $\ell > m$ (polynomial in the size of m) such that $\omega^{(\ell)} = (\omega^{(\ell)}(1), \dots, \omega^{(\ell)}(\ell))$ contains at least m zeros and at least one nonzero value.

We will construct an instance $(\mathcal{E}, k + 1, d_1, c_{d_1})$ of F -FREE-RIDING-RECOGNITION with ℓ voters and $k + 1$ issues. Since ℓ is fixed, we will simply indicate $\omega^{(\ell)}$ as ω . Here, let p be the unique value such that $\omega(p) > \omega(p + 1) = 0$.

In $\mathcal{E}$, there is one voter v_e for each edge $e \in E$, p dummy voters $\{d_1, \dots, d_p\}$, and $\ell - m - p$ dummy voters $\{w_1, \dots, w_{\ell - m - p}\}$. In all issues, all voters w_i approve of all candidates (so they always have satisfaction $k + 1$, and occupy the last $\ell - m - p$ positions of the satisfaction vector). In the first k issues, there is one candidate c_η for each node $\eta \in V$, plus one candidate c_{d_i} for every dummy voter d_i . Here, each edge-voter v_e approves of the two node-candidates c_η and $c_{\eta'}$ such that $e = \{\eta, \eta'\}$, while every d_i approves of c_{d_i} . In the last issue, there is one candidate c , plus one candidate c_i for all voters $i \in N$, and any such i approves only of c_i . In the case of ties, we assume that in the last issue c_{d_1} dominates all candidates and that c_{d_1} is dominated by all other candidates in all other issues. Visually:

Voter	Issues $[k]$	Issue $k + 1$
v_e for $e = \{\eta, \eta'\} \in E$	$\{c_\eta, c_{\eta'}\}$	$\{c_{v_e}\}$
d_i for $i \in [p]$	$\{c_{d_i}\}$	$\{c_{d_i}\}$
w_i for $i \in [\ell - m - p]$	all candidates	all candidates

We will show that $(\mathcal{E}, k+1, d_1, c_{d_1})$ is a legal instance of *F-FREE-RIDING-RECOGNITION* (i.e., that c_{d_1} wins in issue $k+1$). Furthermore, we will show that (G, k) is a yes-instance if and only if $(\mathcal{E}, k+1, d_1, c_{d_1})$ is a no-instance. This shows that the complement of *F-FREE-RIDING-RECOGNITION* is NP-hard, that is, *F-FREE-RIDING-RECOGNITION* is coNP-hard.

Suppose (G, k) is a no-instance of *VERTEXCOVER*. We show that c_{d_1} wins in issue $k+1$ and that d_1 can free-ride here. First, notice that any outcome will give satisfaction 0 to at least p voters. Consider an outcome where we only select node-candidates in the first k issues, and focus on voters in $N \setminus \{w_1, \dots, w_{\ell-m-p}\}$. Since (G, k) is a no-instance, this outcome can give positive satisfaction to at most m of these voters (including the last issue). Any other outcome can only satisfy fewer voters, since dummy voters d_i only approve of their corresponding candidates. Since $\omega(p+1) = 0$, this means that every outcome gives the same opt-OWA score (namely, 0), and by tiebreaking c_{d_1} wins in the last issue. Observe that if d_1 votes for c in the last issue nothing changes. Hence, d_1 can free-ride.

Suppose (G, k) is a yes-instance of *VERTEXCOVER*. We first show that c_{d_1} wins in issue $k+1$. First, observe that at least $m+1+\ell-m-p$ voters need to have positive satisfaction. Were this not the case, as argued above, the total opt-OWA score would be zero. However, the outcome where all edge-voters are satisfied at least once in the first k issues and some dummy candidate wins in the last issue has a greater score—and this is possible, because a vertex cover of size k exists. Then, if d_1 has satisfaction zero, she occupies one of the first $p-1$ entries in the sorted satisfaction vector. Hence, we can increase the opt-OWA score of exactly $\omega(p)$ by selecting c_{d_1} in the last issue. If, on the other hand, c_{d_1} wins in some issue $j \in [k]$, but loses to some candidate c_i (with $i \in N \setminus \{d_1\}$) in issue $k+1$, we could obtain an outcome with a greater or equal opt-OWA score by making c_{d_1} win in issue $k+1$ and making some candidate approved by i win in issue j (Lemma 4.21). Since this would be preferred in the tiebreaking, c_{d_1} wins in $k+1$.

Now we show that d_1 cannot free-ride in the last issue. Suppose that d_1 does not approve of c_{d_1} in issue $k+1$. We have argued above that any outcome must yield a sorted satisfaction vector where at most

$$\ell - (m+1+\ell-m-p) = p+1$$

entries are zero. Hence, suppose towards a contradiction that c_{d_1} wins in the last issue. If in this outcome there is a voter i with satisfaction zero, we could improve the total score by electing c_i in the last issue, because $w(p) > 0$. If all voters have positive satisfaction, it means that d_1 also does. Hence, c_{d_1} is selected in some issue $j \in [k]$. Then, we can obtain an outcome with equal or higher opt-OWA score by selecting any node-candidate approved by at least one voter in issue j and a candidate that d_1 approves of in the last issue (Lemma 4.21). This would be preferred by the tiebreaking. Therefore, d_1 cannot free-ride here.

Finally, observe that the same construction can be used to show hardness for *GENERALISED-F-FREE-RIDING-RECOGNITION*. Indeed, here d_1 only approves of c_{d_1} , and hence free-riding and generalised free-riding coincide. This concludes the proof. $\square$

Theorems 4.22 to 4.24 strengthen our observations from Theorems 4.19 and 4.20. We conclude that free-riding is generally hard for non-sequential rules, since the manipulator cannot even decide efficiently whether free-riding is possible. Thus, in contrast to the sequential setting, free-riding here seems to lose its flavour of simplicity.

4.8 Conclusion

In this chapter, we have studied a specific form of manipulation in multi-issue decision-making: free-riding.

Summary. We summarise the main lessons of this chapter.

- We began by observing that multi-issue voting rules may aim to distribute voter satisfaction fairly across issues, for example by guaranteeing minorities some form of representation. However, especially when issues are decided sequentially, such fairness considerations create the possibility of a particularly simple form of manipulation: free-riding. This occurs when a voter untruthfully opposes a popular candidate in order to gain more influence in subsequent decisions.
- In Section 4.4, we have shown that even minimal representation guarantees in repeated elections lead to the possibility of free-riding. However, we have also shown that, under many natural rules, free-riding can backfire. That is, it can lead to a final satisfaction that is lower than what the manipulator would have enjoyed had she voted earnestly.
- In Section 4.5 we strengthened the above observation—about the risk of free-riding—with a computational argument. We showed that, while it is easy to decide whether free-riding is possible, it is often hard to predict if it will be beneficial. We extended these results to a related form of manipulation, generalised free-riding (Section 4.5.2).
- To assess how serious this risk is in practice, in Section 4.6 we conducted several simulations. We found that harmful free-riding arises with non-negligible probability, especially for voters who free-ride only occasionally rather than systematically. Such parsimony in the use of free-riding might be realistic given the social context of repeated decision-making: in small groups, free-riding may be observable to others and thus carry social costs. Moreover, although free-riding is beneficial in expectation, the expected gains are often limited, suggesting that risk-averse voters may prefer not to engage in it.
- Finally, in Section 4.7, we considered the case where issues are decided simultaneously, and not over time. We argue that free-riding is a rather different phenomenon here. Indeed, while free-riding remains possible, we found that, under several natural rules, it is safe. In other words, it can never lead to a lower final satisfaction than truthful voting. However, we have shown that in this case deciding

whether free-riding is possible is computationally hard, even if the manipulator knows the outcome of the issue where she wants to free-ride on.

Taken together, our results paint an optimistic picture. Free-riding appears to be an unavoidable consequence of representation guarantees in multi-issue decision-making. However, it might be less attractive than it looks. In sequential settings, it can even backfire; in simultaneous settings, it is safe under many rules but computationally difficult to perform. While further research is warranted, our findings suggest that free-riding is not as alluring as one might initially think.

Future work. We conclude by highlighting specific open problems and research directions.

- Many of our hardness proofs rely on more than two candidates per issue. Do all of these results still hold for binary elections, e.g., multiple referenda?
- Our classification of sequential [OWA](#) rules with potentially harmful free-riding is not complete. Are there sequential [OWA](#) rules, other than the utilitarian rule, where free-riding is never harmful?
- Theorem [4.13](#) shows hardness of [SUCCESSFUL-F-FREE-RIDING](#) for the sequential egalitarian rule. Can this result be extended to a larger class of [OWA](#) rules?
- We have shown that under opt-leximin free-riding is always safe; does this extend to other opt-[OWA](#) rules?
- All our results apply to resolute rules, i.e., rules returning exactly one outcome. Many of our results hinge in particular on the fact that ties are broken according to some fixed ordering.¹⁸ This assumption can be relaxed by allowing comparisons between sets of outcomes, as done by [Barrot et al. \[2017\]](#). An interesting direction for future work is to determine whether this relaxation affects our conclusions.
- Another natural extension would be to consider further voting rules, such as rules based on Phragmén’s ideas [[Phragmén, 1895](#), [Janson, 2018](#), [Brill et al., 2024](#)].
- Finally, coordinated free-riding by groups of voters may be much more impactful than single-voter free-riding. Studying this phenomenon, especially in simulations, provides an attractive opportunity to analyse group dynamics in voting, for example regarding coordination and strategy selection.

¹⁸In particular, our results do not capture, e.g., random tiebreaking or mechanisms that use the preferences of specific voters to break ties.

CHAPTER 5

Deliberation

Do I contradict myself?
Very well then I contradict myself,
(I am large, I contain multitudes.)

— Walt Whitman [1882]

5.1 Introduction

In the previous chapters, we discussed two types of multi-winner elections and specific challenges associated with each. In this chapter, we will use multi-winner voting as a formal framework to achieve representation in a different setting: large-scale discussions.

Looking back at the history outlined in Chapter 2, multi-winner elections have always been tied to deliberation. Indeed, the prototypical multi-winner election is that of a representative body, whose job is to engage in debate while representing the electorate. For a long time, this has been the main way to approximate a nation-scale debate: by proxy. Yet the relationship between representation and democracy has long been contested. While representation is today considered an essential feature of democracy, historically the two concepts have often been seen as being in tension [Dovi, 2018]. A famously critical stance was that of Jean-Jacques Rousseau, who claimed that “the moment a people allows itself to be represented, it is no longer free: it no longer exists” [Rousseau, 1762].

Regardless of the validity of such critical views, they highlight a persistent appeal of more direct forms of participation in democratic institutions. In recent years, this appeal has been reflected in a “deliberative wave” that has swept through many democracies [OECD, 2020], bringing with it many new discursive participation formats, including several online discussion platforms. These formats allow citizens to discuss often divisive political issues, and thus offer a way to surface the views of well-informed

citizens more directly. Polis (pol.is), Your Priorities (yrpri.org), and Decidim (decidim.org) are notable examples of the rapidly expanding array of tools used around the world [Fuller, 2022].¹ These platforms allow citizens to submit comments, but also to approve (and sometimes disapprove) comments others have posted, and these discussions can be used to inform political decision-making.




Figure 5.1: Screenshot of an example discussion on Polis about a Canadian electoral reform. Source: <https://pol.is/5mnfxu85hx> (accessed: 2026-06-01).

However, the spectre of representation still looms over these platforms. While digital platforms enable these direct debates, the resulting discussions can be lengthy, disorderly, and even chaotic. Interpreting and summarising them meaningfully in a way that *represents* the actual opinions (and their support) expressed in the discussions poses a challenge in itself. In a sense, this inverts the large-scale deliberation problem: from selecting a (proportionally) representative body to deliberate *for us*, to obtaining a (proportionally) representative image of *our* discussion. This raises a natural question: can we use what multi-winner voting theorists have learned about proportionality to tackle this challenge?

We emphasise that we are not “inventing” this issue artificially. We are not posing the problem for the sake of studying it. Some of the aforementioned platforms themselves are currently researching this. For example, Polis uses machine learning and statistics to find clusters of users with similar opinions, as well as a set of comments representing their joint opinion [Small et al., 2023]. With this method, however, the comments used to represent a cluster of users might be inconsistent (i.e., they might disagree with one another). Additionally, the obtained results are

¹All links accessed on 2026-06-04.

generally not explainable, which is highly problematic for processes that inform political decision-making. A different approach one might consider is to use modern natural language processing techniques, such as large language models. Here, one would feed the discussion into such models to generate a summary. These techniques also suffer from a lack of explainability and, more importantly, typically struggle to produce *representative* summaries, that is, summaries that fairly represent the various positions expressed in the text [Olabisi and Agrawal, 2024, Zhang et al., 2024].

In this chapter, we propose an explainable method for generating representative summaries that are *extractive*, that is, consisting of pieces of text extracted from the original discussion. Our method selects *consistent* viewpoints that represent the opinions of the users as expressed by their approvals, and is grounded in multi-winner voting. Indeed, as we will demonstrate, selecting a representative set of viewpoints based on the approvals of the users is an aggregation problem closely related to ABC voting.

Importantly, because every comment is but one argument in a larger discussion, we suggest that it is insufficient to pick a single unstructured set of representative comments. Instead, we want to identify several consistent *sets of comments*, each jointly expressing a distinct viewpoint. Indeed, viewpoints should be coherent and—to a degree—complete if we want them to be as effective as possible in informing policymakers. To judge whether a set of comments is coherent, we need additional semantic information; for example, whether two comments are in disagreement. We adopt the approach of abstract argumentation [Dung, 1995], a theory of debates and their semantics, which views discussions as sets of arguments that may “attack” one another. This allows us to identify consistent sets of comments that form coherent viewpoints. We can then select a representative subset of these viewpoints based on users’ approvals using tools from ABC voting. An overview of our approach is sketched in Figure 5.2.

We choose abstract argumentation for several reasons. In general, it is one of the most studied models in formal argumentation theory [Baroni et al., 2018]. More importantly, it is particularly well-suited to our setting because it requires only minimal semantic annotation. Moving from a plain-text discussion to a richer representation requires that semantic information be either added by a moderator, crowdsourced from participants, or extracted from natural-language text using argument mining techniques [Palau and Moens, 2009, Cabrio and Villata, 2012, Baumann et al., 2020, Goffredo et al., 2023]. Sticking with abstract argumentation, a lightweight model, might help reduce this annotation effort. In particular, recent developments in large language models suggest that such annotation may become feasible in practice for abstract argumentation in the near future [Gorur et al., 2025, Li et al., 2025].

Our contributions are as follows: we formally introduce **Approval-Based Social Argumentation Frameworks (ABSAFs)** to model online discussions in which citizens can approve arguments. We study the problem of selecting a small but representative set of so-called *preferred extensions* (viewpoints), which are maximally consistent and defensible sets of arguments. We do so from two perspectives. First, if the discussion




Figure 5.2: Simplified representation of our approach. The usernames are fictional, but comments and approvals are taken from a real discussion on Canadian electoral reform from Polis (see Example 5.2).

is small, we might want to pick a set of extensions that perfectly represent every user. Whether this is possible depends, as we shall show, on how we conceptualise perfect representation. Secondly, for larger discussions, we might be more interested in selecting a fixed number of extensions that represent user approvals *as well as possible* among outcomes of this size. We propose several methods for doing so, mostly based on the [OWA](#) operators; this is the part in which we draw most heavily on the connection to the [ABC](#) voting literature. We compare our methods with respect to their computational

complexity, their axiomatic properties, and their performance in numerical simulations; the axioms we consider are inspired by the JR-style axioms discussed in Section 2.2.8. Based on our analysis, we discuss which methods may be more suitable under different application requirements.

5.2 Related Work

Let us first review the related work. There are two main branches of literature to consider. First, we will discuss other approaches that use social choice theory in deliberative contexts. Secondly, since we combine techniques from social choice theory with abstract argumentation, we will also review prior work in this general direction.

Social choice theory and deliberation. Social choice theory has been used before to analyse discussions in participation platforms (such as Polis). For example, see the work of Halpern et al. [2023], in which they formulate the task of choosing a representative set of comments as an ABC voting problem. Moreover, Fish et al. [2026] and Boehmer et al. [2025b] combine ideas from committee elections and machine learning to generate new statements to show users on these platforms. Next, Revel et al. [2025] use JR-style notions to make comment-ranking systems (such as those used in social media) representative. De et al. [2026] likewise use a quantitative analogue of JR to assess the representatives of a slate of statements in the context of deliberative processes (for example, deliberative polls). Another prominent example is the Habermas Machine by Tessler et al. [2024], a system that incorporates large language models and voting to generate consensus statements among groups in democratic deliberation. Finally, similarly to our approach, Maly et al. [2026] present a tool that uses ABC elections (as well as possibly other multi-winner voting models) to select a fixed number of comments from a Polis discussion.²

Notably, none of the above approaches uses formal argumentation models to capture the relationships among comments, nor are they concerned with selecting a consistent subset of them. Hence, they are technically quite distinct from ours. Indeed, to the best of our knowledge, our work is the first to consider user representation in formal argumentation.

Social choice theory and formal argumentation. We focus now on the relationship between social choice and argumentation theory. There is a significant amount of literature on combining the two [Bodanza et al., 2017, Baumeister et al., 2021]. One strain of literature [Caminada and Pigozzi, 2011, Awad et al., 2017] applies ideas from judgment aggregation [Grossi and Pigozzi, 2022] to argumentation. Here, the objective is to aggregate individual judgments about the acceptability of arguments. Crucially, these works assume that individuals submit judgments subject to strict rationality constraints,

²The web-tool is available at <https://simon-rey.github.io/ProportionalityPress/> (accessed: 2026-06-01).

which is unrealistic in online discussions, where users are often not even aware of all the arguments when expressing their preferences.

Another related approach considers the object of study of abstract argumentation—[Argumentation Frameworks \(AFs\)](#)—and adds a social dimension. For example, in weighted AFs [Dunne et al., 2011, Bistarelli and Santini, 2021] attacks between arguments have weights that can be established by having people vote on arguments. Similarly, in social AFs [Leite and Martins, 2011], individuals can approve (or disapprove) arguments. The strength of an argument is then related to its support and the support of its attacking arguments. Beyond abstract argumentation, in the *relational reasoning model* [Ganzer et al., 2023] individuals can judge the acceptability of a set of arguments and their relationships (by assigning them weights). The goal here is to aggregate such judgments and evaluate a set of target arguments collectively. While all three approaches share conceptual similarities with ours, they each redefine argumentation semantics in light of votes (e.g., by considering the strength of arguments or attacks). By contrast, we rely on standard semantics for AFs and use the additional approval information to select a representative subset of extensions. Most importantly, none of the papers mentioned above study the problem of representation, which is the main objective of our work.

The works cited above model scenarios in which there is a single AF and people vote on its arguments. A different approach is one in which everyone is endowed with a distinct AF, and the goal is to merge them into a single collective one. There exists a significant body of work focused on this [Coste-Marquis et al., 2007, Tohmé et al., 2008, Dunne et al., 2012, Delobelle et al., 2016, Chen and Endriss, 2019, Müller et al., 2024]. This differs significantly from our approach, where the AF is given, and we have to decide which extensions to select.

Other works have considered preferences in abstract argumentation without the social dimension. Here, we have a preference order among the arguments (encoding, e.g., argument strength or source reliability). This preference is typically resolved either by selecting a set of most favoured extensions [Brewka et al., 2010, Alfano et al., 2022, 2023] or by modifying the attack relation of the given AF [Modgil, 2009, Atkinson and Bench-Capon, 2021, Kaci et al., 2021, Bernreiter et al., 2022, 2023]. The former solution is more similar to ours.

As a final remark, since the publication of the work on which this chapter is based [Bernreiter et al., 2024a], Rossie et al. [2024] have considered a generalisation of our approach to settings where users can approve, disapprove, or be neutral towards comments. By contrast, we only allow for approval or neutrality.

5.3 Preliminaries

We will now introduce the formal preliminaries necessary for our analysis. We highlight that the formal object we introduce, [ABSAFs](#), is novel and is a contribution in itself.

Before delving into the details, let us give a high-level view. The basic setup considered in this chapter is the following. We have a discussion consisting of a set of interrelated comments. Moreover, each comment can be approved (or ignored) by

the users in the system. The problem we are facing, namely selecting representative comments based on the approvals of users, has a similar input as the standard ABC voting model: a finite set of voters N (also called *users* in this chapter), a finite set of candidates C (also called *comments* or *arguments*), and a profile of approval ballots A . The difference from ABC voting lies in the outcome: we want to select a *set of subsets* of C , i.e., a subset of 2^C , typically denoted by Ω . This is in contrast to ABC elections, where we want to select a subset of C directly. Often, we will study the problem of selecting exactly k subsets, i.e., we will stipulate $|\Omega| = k$. We will not constrain the cardinality of the selected subsets in Ω ; instead, we impose consistency constraints on them using the toolbox of abstract argumentation. We next introduce AFs as they are typically studied; this is the formal model we use to represent discussions. We will then enrich it with user preferences (approvals over the comments) in Section 5.3.2.

5.3.1 Abstract Argumentation

Abstract argumentation [Dung, 1995, Baroni et al., 2018] is a field that formally studies argumentative scenarios, with the goal of identifying argumentative outcomes (sets of arguments) that are “rational”. It does so by focusing on the relationship between arguments, not their structure (hence the word “abstract”). This is a well-studied formalism in which discussions can be represented and reasoned about. We will provide concrete examples of how to represent real-world discussions in this model.

The main object of study is that of **Argumentation Frameworks (AFs)**, introduced first by Dung [1995]. Arguments in AFs have no inner structure, and are typically denoted by lower-case letters, e.g., a , b , c , and so on.³ An argument a can attack another argument b , which implies that a and b are in conflict and cannot be jointly accepted. To accept b , it must be defended against the attack of a , i.e., either b attacks a or there is another argument c which attacks a and can be accepted alongside b . We will indicate the set of arguments, which we sometimes refer to as comments, as C .

Definition 5.1. An AF is a tuple $\text{Fram} = (C, \rightarrow)$ where C is a finite set of *arguments* (or *comments*) and $\rightarrow \subseteq C^2$ is an *attack relation* between arguments. We write $a \rightarrow b$ to mean $(a, b) \in \rightarrow$ and $a \not\rightarrow b$ to mean $(a, b) \notin \rightarrow$.

A set $S \subseteq C$ *attacks* $b \in C$ if $a \rightarrow b$ for some $a \in S$. Moreover, $S \subseteq C$ *defends* $b \in C$ if S attacks every $a \in C$ with $a \rightarrow b$.

AFs are usually visualised as directed graphs, where nodes are arguments and edges attacks between arguments (see Example 5.1). Note that self-attacking arguments are allowed in this model.

Readers unfamiliar with abstract argumentation might be perplexed by the fact that we represent only one possible relationship among arguments, namely, attacks. While other approaches exist (perhaps the most natural extension would be that of bipolar

³Elsewhere in this thesis, we indicated candidates with lower-case letters. This is not an abuse of notation: as we shall see, the sets of arguments and candidates, in this chapter, coincide.

AFs, where arguments can attack or support each other [Cayrol and Lagasquie-Schiex, 2005]), the one we present here is the standard one, and most understood. Nonetheless, we maintain that this restriction is reasonable for our application, as outlined in the introduction. Indeed, discussions arise naturally as unstructured collections of comments, and relationships, as it currently stands, must be annotated, either automatically or manually. We believe that striving for semantic minimality, as we do here, makes the task easier. Finally, we observe that nothing in our approach hinges critically on the underlying discussion formalism, meaning that our framework should be extensible to richer models if needs be.


Let us go back to **AFs**. How to make sense of such an object? The “outcome” of an **AF** is captured through argumentation semantics. Each semantics is defined as a function sem assigning a set $sem(Fram) \subseteq 2^C$ of sets of arguments to every $Fram = (C, \rightarrow)$. In the jargon of argumentation theory, such sets are called *extensions*. The extensions of an **AF** are the set of arguments (comments) that a rational agent can simultaneously accept or believe in. What “rational” means is encoded by the semantics; that is, each semantics represents a different view on what the rational outcome of a framework should be. Let us review some standard ones (see the chapter by Baroni et al. [2018] for alternative semantics and more details). *Conflict-free* semantics, denoted as cf , select all sets $S \subseteq C$ where no two arguments attack each other (i.e., the independent sets of $Fram$). *Admissible* semantics, denoted as adm , select conflict-free sets that defend themselves. *Preferred* semantics, denoted as prf , select subset-maximal admissible sets.

Definition 5.2. Let $Fram = (C, \rightarrow)$ be an **AF**. We define

$$\begin{aligned} cf(Fram) &= \{S \subseteq C : a \not\rightarrow b \text{ for all } a, b \in S\}, \\ adm(Fram) &= \{S \in cf(Fram) : \text{every } a \in S \text{ is defended by } S\}, \text{ and} \\ prf(Fram) &= \{S \in adm(Fram) : S \not\subseteq T \text{ for all } T \in adm(Fram)\}. \end{aligned}$$

Observe that $prf(Fram) \subseteq adm(Fram) \subseteq cf(Fram)$ holds by definition. We will see later how this can capture an actual discussion (Example 5.2). For now, let us consider an example that focuses on the formal machinery of abstract argumentation.

Example 5.1. Let $Fram$ be the following **AF**:



Here, we have that

$$cf(\text{Fram}) = \{\emptyset, \{a\}, \{b\}, \{c\}, \{d\}, \{e\}, \{a, b\}, \{a, d\}, \{a, e\}, \{b, e\}, \{c, e\}, \{a, b, e\}\}.$$

Note that the empty set is always conflict-free. Moreover, a singleton set of form $\{a\}$ is conflict-free whenever $a \in C$ is not self-attacking. Of the above, the admissible (i.e., self-defending) sets are

$$adm(\text{Fram}) = \{\emptyset, \{a\}, \{b\}, \{a, b\}, \{b, e\}, \{a, b, e\}\}.$$


Note that, again, the empty set is always admissible. For example, here, $\{b, e\}$ defends itself: all its attackers (namely, c and d) are attacked back by $\{b, e\}$ because b attacks both c and d . On the other hand, $\{c\}$ is not admissible because c is not defending itself from the attack of a (although it is defending itself from b). Finally, among the extensions in $adm(\text{Fram})$, the only preferred extension is $prf(\text{Fram}) = \{\{a, b, e\}\}$, as it is the only inclusion-maximal admissible extension.

Preferred extensions are, as we discuss at the end of Section 5.3.2, particularly suitable for our purposes. Indeed, they encode a rather strong notion of rationality, and represent a set of maximally “consistent and defensible” arguments. We will thus mostly focus on these semantics.

One might now ask: why then, is argumentation theory not enough for our purposes? Why not simply use the preferred extensions of a discussion as its summary? The problem lies in the fact that, in general, there might be too many preferred extensions to concisely summarise a discussion—exponentially many. A trivial bound is the following:

Observation 5.1. An AF with m arguments can have $2^{\lfloor m/2 \rfloor}$ preferred extensions.

Proof. Assume m is even. Consider the following AF:



Here, the preferred extensions are exactly all sets of form $\{\star_1, \dots, \star_{m/2}\}$ with $\star_i \in \{a_i, b_i\}$ for all $i \in [m/2]$. Indeed, any argument here defends itself. By maximality we must include at least one argument for every $i \in [m/2]$, and by conflict-freeness, at most one. Hence, we must choose for every $i \in [m/2]$ either a_i or b_i . There are $2^{m/2}$ such choices.

If m is not even, we can just add an isolated argument which will be featured in all preferred extensions. The construction and the computation remain unchanged. $\square$

For finer bounds on the number of extensions one can build out of a framework with m arguments, see the work of [Baumann and Strass \[2014\]](#) and [Dunne et al. \[2015\]](#).

We will introduce next a preference-based way to select a representative set of extensions.

5.3.2 Approval-Based Social Argumentation Frameworks

Let us now introduce our main object of study, [Approval-Based Social Argumentation Frameworks \(ABSAFs\)](#), which model discussions where users can approve arguments that they find convincing. Mechanically, this model is not a big leap from what we have seen till now. It is essentially an [AF](#) paired with a profile of approval ballots over the arguments as seen in the [ABC](#) voting model.

Definition 5.3. An [ABSAF](#) $\text{Disc} = (\text{Fram}, N, A)$ consists of an [AF](#) $\text{Fram} = (C, \rightarrow)$, a finite set $N = [n]$ of *voters* (also called *users*), and an approval profile $A = (A(1), \dots, A(n))$ with $\emptyset \subset A(i) \subseteq C$ for all in $i \in N$.

We do not stipulate any constraints for the submitted ballots (not even conflict-freeness) as ballots containing conflicting arguments appear in real-world examples, as the next example shows.

Example 5.2. We consider a discussion on an election reform proposed in Canada, taken from the “Computational Democracy Project”.⁴ This is the same discussion as in Figures 5.1 and 5.2. For succinctness, we consider only two arguments in favour of the first-past-the-post (FPTP) electoral system (f_1, f_2), three arguments in favour of proportional representation (PR) (p_1, p_2, p_3), two arguments in favour of the single transferable vote (STV) system (s_1, s_2), and one argument in favour of a mixed-member proportional (MMP) system (m_1). The `cid` refers to the `comment-id` field in the original dataset.

- f_1 (`cid` = 16): “If we used a system with PR, the opinions would be too strongly divided and nothing would ever get done.”
- f_2 (`cid` = 15): “The current electoral system results in more stable governance (ie: more majority governments which can serve a full term).”

⁴See <https://compdemocracy.org/> (accessed: 2026-04-17). The specific dataset is available at <https://github.com/compdemocracy/openData/tree/master/canadian-electoral-reform> (accessed: 2026-04-17), and the Polis discussion at <https://pol.is/5mnfxu85hx> (accessed: 2026-06-01).

- p_1 (cid = 21): “I believe a party’s share of seats in the House of Commons should reflect its share of the popular vote.”
- p_2 (cid = 42): “PR will reduce hyper-partisanship and promote greater collaboration and cooperation between parties.”
- p_3 (cid = 43): “PR will reduce wild policy swings and result in more long-lasting policies with a view to the future.”
- s_1 (cid = 162): “STV’s advantage over MMP is that it doesn’t explicitly enshrine political parties in our electoral system.”
- s_2 (cid = 168): “I like the simplicity of stating my favourite candidate, and who I hope will win if my first choice does not receive enough votes.”
- m_1 (cid = 163): “MMP has a better chance of obtaining public support than STV due to its relative simplicity.”

For the sake of this example (and with no pretence of ultimate truth), we manually added attacks between arguments. The resulting [AF](#) is

$$p_1 \longleftarrow f_1 \begin{array}{c} \longrightarrow \\ \longleftarrow \end{array} p_2 \quad f_2 \begin{array}{c} \longrightarrow \\ \longleftarrow \end{array} p_3 \quad s_1 \longleftarrow m_1 \begin{array}{c} \longrightarrow \\ \longleftarrow \end{array} s_2$$

There are 8 preferred extensions, namely

$$\text{prf}(\text{Fram}) = \{\{p_1, p_2, p_3, s_1, s_2\}, \{p_1, p_2, p_3, m_1\}, \{f_2, p_1, p_2, s_1, s_2\}, \\ \{f_2, p_1, p_2, m_1\}, \{f_1, p_3, s_1, s_2\}, \{f_1, p_3, m_1\}, \{f_1, f_2, s_1, s_2\}, \{f_1, f_2, m_1\}\}.$$

The following table shows the ballots along with how many users voted for the given ballot, as extracted from the original dataset. Note that we are abstracting away the identities of the voters, and just considering the frequency of each ballot.

$\{p_1\}$	33	$\{p_1, p_3\}$	9	$\{f_2, p_2, p_3\}$	3	$\{f_2, p_3\}$	1
$\{p_1, p_2, p_3\}$	31	$\{p_1, p_2\}$	8	$\{f_1, p_1\}$	2	$\{m_1\}$	1
$\{p_2\}$	16	$\{f_2, p_1, p_2, p_3\}$	7	$\{f_2, p_1, p_2\}$	2	$\{m_1, p_2\}$	1
$\{p_3\}$	16	$\{s_2\}$	6	$\{f_2, p_1, p_3\}$	2	$\{p_1, p_2, p_3, s_1, s_2\}$	1
$\{f_2\}$	11	$\{f_1, f_2\}$	4	$\{p_1, p_2, p_3, s_1\}$	2	$\{p_2, p_3, s_2\}$	1
$\{p_2, p_3\}$	10	$\{m_1, p_1, p_2, p_3, s_1\}$	4	$\{f_1, f_2, p_3\}$	1	$\{p_2, s_1\}$	1
$\{f_2, p_1\}$	9	$\{f_1\}$	3	$\{f_1, p_1, p_2, p_3\}$	1	$\{s_1\}$	1

Most voters approve arguments in favour of PR. Some approve both arguments in favour of PR and FPTP, e.g., the ballot $\{f_1, p_2\}$. Note that this ballot is conflict-free. However, not all ballots are conflict-free, e.g., $\{f_2, p_1, p_2, p_3\}$, which was approved by seven voters. Moreover, the ballot with the most votes, namely $\{p_1\}$, is not admissible. On the other hand, $\{p_1, p_2\}$ is admissible, since p_2 defends p_1 .

The goal is to select a (usually small) set of coherent viewpoints representing the approvals of the users. Importantly, we model viewpoints as extensions of the given [AF](#) (sets of comments). Hence, our goal is to select a small set of extensions that represent the approvals. We refer to the set of viewpoints we return as an *outcome*. Note that, at this stage, we do not concern ourselves with the size of each extension (i.e., the total number of comments in the outcome).

Definition 5.4. Given semantics sem , an *outcome* for an [ABSAF](#) $Disc = (Fram, N, A)$ is a set of extensions $\Omega \subseteq sem(Fram)$. We call each $\pi \in \Omega$ a *viewpoint*.

Note that this says nothing about what makes a viewpoint or an outcome “representative”; the following definition captures this. Intuitively, the representation of a voter by a viewpoint is the fraction of the voter’s approved arguments that the viewpoint covers. The representation of a voter by an outcome is then given by the viewpoint in the outcome that represents the voter best.

Definition 5.5. Let $Disc = (Fram, N, A)$ be an [ABSAF](#) and let $\alpha \in [0, 1]$. A viewpoint $\pi \in sem(Fram)$ is said to α -represent voter $i \in N$ if

$$rep_i(\pi) = \frac{|\pi \cap A(i)|}{|A(i)|} \geq \alpha.$$

Moreover, an outcome $\Omega \subseteq sem(Fram)$ is said to α -represent voter $i \in N$ if

$$rep_i(\Omega) = \max_{\pi \in \Omega} rep_i(\pi) \geq \alpha.$$

We will justify the use of the max operator, as well as our focus on preferred extensions, after the following technical example.

Example 5.3. Consider again the [ABSAF](#) from Example 5.2. Fix some voter $i \in N$ with $A(i) = \{f_2, p_1, p_2\}$ and consider the preferred extensions (viewpoints)

$$\begin{aligned} \pi_1 &= \{f_1, f_2, m_1\}, \\ \pi_2 &= \{p_1, p_2, p_3, m_1\}, \text{ and} \\ \pi_3 &= \{f_2, p_1, p_2, m_1\}. \end{aligned}$$

Then,

$$\begin{aligned} rep_i(\pi_1) &= 1/3, \\ rep_i(\pi_2) &= 2/3, \text{ and} \\ rep_i(\pi_3) &= 1. \end{aligned}$$

Regarding outcomes, of size 2 we have

$$\begin{aligned} \text{rep}_i(\{\pi_1, \pi_2\}) &= 2/3 \text{ and} \\ \text{rep}_i(\{\pi_1, \pi_3\}) &= \text{rep}_i(\{\pi_2, \pi_3\}) = 1. \end{aligned}$$

Let us now discuss this definition, and, at the same time, justify our focus on preferred extensions. In deliberative democracy, one usually assumes that voters are rational agents [Cohen, 2007] and thus, in principle, would arrive at a consistent and defensible viewpoint when considering all arguments for a sufficient amount of time, which we could call their *ideal* viewpoint. However, in practice, the vast majority of participants in an online discussion will not carefully consider all comments. Hence, we do not assume that ballots are complete or consistent. In light of this, we interpret $\text{rep}_i(\pi)$ as a measure of how consistent the observed voting behaviour of i is with the assumption that π is her ideal viewpoint.⁵ Due to our rationality assumption, ideal viewpoints should at least be admissible. As we aim to represent voters with a small number of extensions, we can focus on preferred extensions, *prf*, since this gives us subset-maximal coherent viewpoints. If an admissible viewpoint π contains all approved arguments of a voter, so does the preferred viewpoint π' with $\pi \subseteq \pi'$. Finally, as two different preferred extensions must be conflicting,⁶ each voter's ideal viewpoint can only coincide with one preferred extension. Thus, we define $\text{rep}_i(\Omega)$ as $\max_{\pi \in \Omega}(\text{rep}_i(\pi))$.

Notably, our use of the maximum operator (which implicitly assigns each voter to the viewpoint that best represents her) has an important precedent in the context of committee elections under linear preferences. In the rule proposed by Chamberlin and Courant [1983], each voter is assigned to the candidate she most prefers among the selected committee, and her “level of representation” is evaluated solely based on her assigned representative, analogously to our model of representation. More recently, and in a context closer to ours (finding a representative slate of statements in a deliberative process), De et al. [2026] also use an analogous model. The representation (in their terminology: utility) enjoyed by a participant for a slate is given by the maximum representation yielded by a statement in the slate.

This concludes the modelling part of this chapter. Let us now analyse the features of our objects of interest.

⁵Interestingly, Dodgson [1884] proposed similar views in the context of committee elections: according to Monroe [1995] he believed that “a representation system should find the coalitions in the electorate that would have formed if the voters had the necessary time and information”. However, he later abandoned the idea due to difficulties stemming from the instability of coalition formation. A closely related account of representation was put forth by Rogowski [1981].

⁶Indeed, if there were two preferred extensions π and π' where $\pi \cup \pi'$ is conflict-free, then $\pi \cup \pi'$ would also be admissible, because this set defends itself. To see why, consider some $b \in C$ which attacks some argument in π or in π' . By admissibility of π and π' there must be some argument $a \in (\pi \cup \pi')$ with $a \rightarrow b$, showing that $\pi \cup \pi'$ defends itself. However, this cannot hold, as it contradicts the maximality of π and π' .

5.4 Perfect Representation

Ideally, we would like to find an outcome that perfectly represents everyone, that is, an outcome that 1-represents all voters. We begin by studying this problem. It differs slightly from the problems we have considered so far in this thesis, in that we are not trying to select a fixed number of viewpoints, but rather to determine whether there exists *any* perfectly representative set of viewpoints.

Nonetheless, we believe this problem fits within our multi-winner theme. First, we are still selecting a subset of items from a larger set according to the approval preferences of a population. Second, and more directly, this analysis helps us build fundamental intuition about representability in our model. As we shall see, it also motivates casting this problem as a more traditional [ABC](#) election—namely, one in which k , the number of viewpoints to be selected, is fixed in advance.

5.4.1 When Can We Achieve Perfect Representation?

Right off the bat, we can observe that perfectly representing all voters is not always possible. Indeed, Example 5.2 clearly demonstrates that real votes cannot be assumed to be conflict-free, in which case there can be no outcome that 1-represents everyone.

But even if ballots are conflict-free, it may not be possible to represent all voters fully. Consider the [ABSAF](#) Disc = (Fram, N , A) from Figure 5.3:




Figure 5.3: Voter 2 approves indefensible arguments.

Here, we have $\text{prf}(\text{Fram}) = \{\{a, b\}, \{a, c\}\}$ but $\text{rep}_2(\{a, b\}) = 1/4$ and $\text{rep}_2(\{a, c\}) = 1/2$. The issue here is that voter 2 approves two indefensible arguments, namely d and e , that are featured in no admissible extension.

Not only it is not always possible to perfectly represent all voters: deciding *when* it is possible is NP-complete. This follows from the NP-completeness of deciding credulous acceptance with preferred semantics [[Dvořák and Dunne, 2018](#)]. An argument a is credulously accepted (for semantics sem) if and only if there exists $S \in \text{sem}(\text{Fram})$ with $a \in S$.

Definition 5.6. 1-REPRESENTABILITY is the following decision problem: “Given an [ABSAF](#) Disc = (Fram, N , A) and $k \in [n]$, is there $\Omega \subseteq \text{prf}(\text{Fram})$ with $|\Omega| \leq k$ such that $\text{rep}_i(\Omega) = 1$ for all $i \in N$?”

Note that assuming $k \in [n]$ is not a restriction, as the representation of a voter only depends on the viewpoint that best represents her.

Proposition 5.2. 1-REPRESENTABILITY is NP-complete. NP-hardness holds even if there is only one voter approving a single argument.

Proof. We first show membership. Given $\text{Disc} = (\text{Fram}, N, A)$, guess a set of viewpoints $S \subseteq 2^C$ with $|S| = k$. Note that $|S| \leq n$ and, for any $\pi \in S$, $|\pi| \leq |C|$. Thus, the size of S is polynomial in the size of Disc . First, verify that all $\pi \in S$ have $\pi \in \text{adm}(\text{Fram})$; this is possible in polynomial time [Dvořák and Dunne, 2018]. Next, for each $i \in N$, verify that there is some $\pi \in S$ such that $\text{rep}_i(\pi) = 1$ (trivially possible in polynomial time). If this is the case, then we are done. Indeed, if $\pi \in \text{adm}(\text{Fram})$, then there is $\pi' \in \text{prf}(\text{Fram})$ with $\pi \subseteq \pi'$ that 1-represents all voters whom were 1-represented by π . This follows from the relevant definitions immediately.

Let us now show NP-hardness. Let (Fram, a) be an arbitrary instance of credulous acceptance for preferred semantics. To reiterate, this NP-hard problem consists of deciding whether there is at least one preferred extension of Fram containing a [Dvořák and Dunne, 2018]. We construct the ABSAF $\text{Disc} = (\text{Fram}, N, A)$ with $N = \{1\}$ and $A(1) = \{a\}$. Then a is credulously accepted if and only if there is $\pi \in \text{prf}(\text{Fram})$ such that $a \in \pi$, which is equivalent to $\{a\} = A(1) \subseteq \pi$. The latter is equivalent, in turn, to the existence of some $\Omega \subseteq \text{prf}(\text{Fram})$ with $|\Omega| \leq 1$ that 1-represents the unique voter. $\square$

The above result also suggests that finding a straightforward, “simple” characterisation for representability is not possible in the general case.

In light of this, a natural question is then whether there are scenarios for which 1-representability is always possible. It turns out that this is the case. First, assuming conflict-free votes, we can give a lower-bound for α -representation by differentiating between arguments that defend themselves and arguments that do not.

Proposition 5.3. Let $\text{Disc} = (\text{Fram}, N, A)$ be an ABSAF. Let

$$D(i) = \{a \in A(i) : b \rightarrow a \text{ implies } a \rightarrow b\}$$

be the self-defending arguments approved by $i \in N$. If for all $i \in N$ we have that $A(i)$ is conflict-free and

$$\frac{|D(i)|}{|A(i)|} \geq \alpha,$$

then there is $\Omega \subseteq \text{prf}(\text{Fram})$ such that $\text{rep}_i(\Omega) \geq \alpha$ for all $i \in N$.

Proof. Consider $i \in N$. Observe that $D(i) \in \text{adm}(\text{Fram})$ since $A(i)$ is conflict-free by assumption and since each argument in $D(i)$ defends itself. Thus, there is $\pi_i \in \text{prf}(\text{Fram})$ such that $D(i) \subseteq \pi_i$. Observe that

$$\text{rep}_i(\pi_i) = \frac{|\pi_i \cap A(i)|}{|A(i)|} \geq \frac{|D(i) \cap A(i)|}{|A(i)|} = \frac{|D(i)|}{|A(i)|} \geq \alpha.$$

Thus, $\Omega = \{\pi_i : i \in N\}$ α -represents every voter in N . $\square$

Additionally, consider symmetric AFs [Coste-Marquis et al., 2005, Dunne, 2007], a natural and well-studied subclass of AFs where $a \rightarrow b$ implies $b \rightarrow a$ for every $a, b \in C$. For example, the framework on the left is symmetric, the framework on the right is not:



In symmetric AFs, every voter only approves self-defending arguments by definition. Thus, by Proposition 5.3, if we are given conflict-free ballots and a symmetric AF Fram, there is an outcome $\Omega \subseteq \text{prf}(\text{Fram})$ that perfectly represents all voters.

Corollary 5.4. Let $\text{Disc} = (\text{Fram}, N, A)$ be an ABSAF. If $A(i)$ is conflict-free for every voter $i \in N$ and if Fram is symmetric, then there is an outcome $\Omega \subseteq \text{prf}(\text{Fram})$ such that $\text{rep}_i(\Omega) = 1$ for all $i \in N$.

5.4.2 Relaxing the Notion of Representation

So far, we attempted to represent *all* approved arguments for each voter. However, even if we assume votes to be conflict-free, some voters might approve of arguments that cannot occur together in a preferred extension, or even arguments that cannot be defended altogether. As we already pointed out, voter 2 in Figure 5.3, for instance, approves two indefensible arguments (d and e). Thus, $\text{rep}_2(\{a, c\}) = 1/2$ even though $\{a, c\}$ contains *all* arguments approved by voter 2 that can *ever* occur in a viewpoint—namely, a and c . Hence, one could argue that voter 2 is already represented as well as possible by this viewpoint. We therefore introduce an alternative notion of representation that we call *core-representation*. The term “core” stems from the fact that we are trying to represent the “defensible core” of a voter.

Definition 5.7. Let $\text{Disc} = (\text{Fram}, N, A)$ be an ABSAF. For $i \in N$ we let

$$\mu(i) = \max_{\pi \in \text{sem}(\text{Fram})} |\pi \cap A(i)|$$

be the maximum number of arguments approved by i that jointly appear in an extension.⁷ We define the *core-representation* of a viewpoint π for voter i as

$$\text{core}_i(\pi) = \begin{cases} \frac{|\pi \cap A(i)|}{\mu(i)} & \text{if } \mu(i) > 0, \\ 1 & \text{otherwise.} \end{cases}$$

We say that $\pi \in \text{sem}(\text{Fram})$ α -core-represents i if and only if $\text{core}_i(\pi) \geq \alpha$. Moreover, $\Omega \subseteq \text{sem}(\text{Fram})$ α -core-represents i if and only if $\text{core}_i(\Omega) = \max_{\pi \in \Omega} \text{core}_i(\pi) \geq \alpha$.

Using core-representation for the **ABSAF** from Figure 5.3, for voter 1 we get $\text{core}_1(\{a, b\}) = 1$ and $\text{core}_1(\{a, c\}) = 1/2$ while for voter 2 we get $\text{core}_2(\{a, b\}) = 1/2$ and $\text{core}_2(\{a, c\}) = 1$.

In contrast to regular representation, it is always possible to find an outcome that perfectly core-represents every voter—this is true by definition.

Observation 5.5. Given an **ABSAF** $\text{Disc} = (\text{Fram}, N, A)$, the outcome $\Omega = \text{prf}(\text{Fram})$ 1-core-represents every voter $i \in N$. Thus, there is an outcome $\Omega' \subseteq \text{prf}(\text{Fram})$ of size $|\Omega'| = n$ that 1-core-represents every $i \in N$.

However, deciding whether there is a *small* outcome that perfectly core-represents every voter is harder than in the case of regular representation, for which this problem is NP-complete (Proposition 5.2). Intuitively, this depends on the additional computational burden of determining $\mu(i)$ for every $i \in N$. We define 1-CORE-REPRESENTABILITY analogously to 1-REPRESENTABILITY (Definition 5.6), except that we ask for an outcome of size $|\Omega| \leq k$ that 1-core-represents every voter $i \in N$.

Definition 5.8. α -CORE-REPRESENTABILITY is the following decision problem: “Given an **ABSAF** $\text{Disc} = (\text{Fram}, N, A)$ and a number $k \in [n - 1]$, is there an α -core-representative outcome $\Omega \subseteq \text{prf}(\text{Fram})$ of size $|\Omega| \leq k$?”

Note that we assume $k \in [n - 1]$ because the case $k = n$ is trivial (always true).

Theorem 5.6. 1-CORE-REPRESENTABILITY is Θ_2^P -complete. Moreover, Θ_2^P -hardness holds even if $n = 2$.

Recall that, as stated in the preliminaries (Section 2.2), Θ_2^P is the class of decision problems solvable in polynomial time with access to $O(\log n)$ calls to an NP-oracle [Wagner, 1990]. Problems that are hard for this class are believed to be harder than NP-complete problems.

Proof of Theorem 5.6. We begin with Θ_2^P -membership. Let $\text{Disc} = (\text{Fram}, N, A)$ be the input **ABSAF** and $k \in [n - 1]$. We make use of the fact that Θ_2^P coincides with $P_{||}^{\text{NP}}$ [Buss and Hay, 1991], which is the class of problems solvable in polynomial time with polynomially many *nonadaptive* queries to NP-oracles. This means that the queries cannot depend on the result of previous ones. Recall that $|C| = m$.

First, we compute $\mu(i)$ for all voters i . For every $i \in N$ and $s \in [m]$, we use an NP-oracle to decide whether there is $\pi \in \text{adm}(\text{Fram})$ such that $|\pi \cap A(i)| = s$. This problem is in NP, since we can guess $\pi \subseteq C$ and check in polynomially time whether

⁷As a mnemonic, $\mu(i)$ uses symbol μ , the Greek letter corresponding to “m”, to allude to the *maximum* operator in its definition.

$\pi \in \text{adm}(\text{Fram})$ and $|\pi \cap A(i)| = s$. Then, for each voter $i \in N$, we let $\mu(i)$ being the largest $s \in [m]$ for which the above query is true; this is legitimate, because

$$\mu(i) = \max_{\pi \in \text{prf}(\text{Fram})} |\pi \cap A(i)| = \max_{\pi \in \text{adm}(\text{Fram})} |\pi \cap A(i)|.$$

Next, we use a single NP-oracle query to execute the following decision-procedure. First, we guess argument sets $S = \{\pi_1, \dots, \pi_k\} \subseteq 2^C$. Then, we check that $\pi_s \in \text{adm}(\text{Fram})$ for every $\pi_s \in S$. Finally, we check that, for every voter $i \in N$, there is some $\pi_s \in S$ such that $|\pi_s \cap A(i)| = \mu(i)$. The query succeeds whenever such an S exists; this is clearly in NP as all the checks on S are in P. We claim that the input instance is a yes-instance if and only if the this query is successful.

To see this, consider the following. If the input instance is a no-instance, then the above query will fail, because there is no set of k admissible viewpoints that can perfectly core-represent all voters (and thus, no set of k preferred viewpoints can do the same). Suppose then the input is a yes-instance. We construct the outcome

$$\Omega = \{\pi'_s \in \text{prf}(\text{Fram}) : \pi_s \in S \text{ and } \pi_s \subseteq \pi'_s\}$$

obtained by extending every admissible π_s to a preferred extension π'_s (this is possible by definition). By the above query, for every $i \in N$ we have a corresponding $\pi_s \in S$ with $|\pi_s \cap A(i)| = \mu(i)$. This implies

$$|\pi_s \cap A(i)| \leq |\pi'_s \cap A(i)| \leq \mu(i) = |\pi_s \cap A(i)|,$$

which implies $|\pi'_s \cap A(i)| = \mu(i)$ and therefore $\text{core}_i(\pi'_s) = 1$. Thus, Ω does 1-core-represent all voters and we are done.

We now show Θ_2^P -hardness. Let (φ, x^*) be an instance of the Θ_2^P -hard problem **CARDMAXSAT** [Creignou et al., 2018], where we are given a propositional formula φ in **CNF** and a variable x^* , and ask whether there is a cardinality-maximal model of φ that contains x^* . Let X denote the set of variables contained in φ , and $\Gamma = \{\gamma, \dots\}$ the set of clauses. Note that φ can be assumed to be satisfiable, since φ can be replaced by $\varphi' = \varphi \vee (\bigwedge_{x \in X} (\neg x))$, which can further be transformed into **CNF** if needed. Indeed, φ' is always satisfied by interpretation $\emptyset$ and (φ, x^*) is a yes-instance of **CARDMAXSAT** if and only if (φ', x^*) is a yes-instance of **CARDMAXSAT**.

We construct an instance of **1-CORE-REPRESENTABILITY** (Disc, k) with $k = 1$ and $\text{Disc} = (\text{Fram}, N, A)$ with $\text{Fram} = (C, \rightharpoonup)$ where

$$C = X \cup \{\neg x : x \in X\} \cup \Gamma \cup \Phi \quad \text{where} \quad \Phi = \{\varphi_1, \dots, \varphi_{|X|+1}\}$$

and

$$\begin{aligned} \rightharpoonup = & \{(x, \neg x), (\neg x, x) : x \in X\} \cup \\ & \{(x, \gamma) : x \in X, \gamma \in \Gamma, \text{ and } x \in \gamma\} \cup \\ & \{(\neg x, \gamma) : x \in X, \gamma \in \Gamma, \text{ and } \neg x \in \gamma\} \cup \\ & \{(\gamma, \gamma), (\gamma, \varphi_\ell) : \gamma \in \Gamma \text{ and } \varphi_\ell \in \Phi\}. \end{aligned}$$

Moreover, $N = \{1, 2\}$ and A is given by $A(1) = X \cup \Phi$ and $A(2) = \{x^*\}$. As an example, Figure 5.4 shows the AF constructed from the formula containing clauses $\gamma_1 = (x_1 \vee \neg x_2 \vee x_3)$, $\gamma_2 = (\neg x_1 \vee \neg x_3 \vee x_4)$, and $\gamma_3 = (x_2 \vee x_3 \vee \neg x_4)$. If $x^* = x_1$, then the approval ballots in the same example would be $A(1) = \{x_1, \dots, x_4, \varphi_1, \dots, \varphi_5\}$ and $A(2) = \{x_1\}$.




Figure 5.4: Construction used in the proof of Theorem 5.6.

We show that (φ, x^*) is a yes-instance of CARDMAXSAT if and only if there is an outcome Ω of size $|\Omega| = 1$ that 1-core-represents all $i \in N$.

We begin with the “if” direction. Assume $x^* \in I$ for a cardinality-maximal model $I \subseteq X$ of φ . Let $\pi = I \cup \{\neg x : x \in X \setminus I\} \cup \Phi$. We claim that outcome $\{\pi\}$ does 1-core-represent both voters. Since I satisfies every clause in φ , every clause-argument $\gamma \in \Gamma$ is attacked by some propositional-variable-argument in $I \cup \{\neg x : x \in X \setminus I\}$. Thus, every argument in C is either in π or attacked by π , i.e., π defends itself against all outside attackers. Since $\pi \in cf(\text{Fram})$ by construction, we have $\pi \in adm(\text{Fram})$. Moreover, π is a subset-maximal admissible set, since it is in conflict with all arguments outside π . Thus, $\pi \in prf(\text{Fram})$. Consider now any $\pi' \in prf(\text{Fram})$. There are two cases:

1. If $\pi' \cap \Phi = \emptyset$, then $|\pi' \cap A(1)| \leq |X| < |\Phi| \leq |\pi \cap A(1)|$.
2. Otherwise, $\pi' \cap \Phi \neq \emptyset$. Then π' defends some $\varphi_\ell \in \Phi$ against attack $\gamma \rightarrow \varphi_\ell$ for every $\gamma \in \Gamma$. Thus, $I' = \pi' \cap X$ satisfies all clauses in φ . Since I is a cardinality-maximal model of φ , $|I'| \leq |I|$. Therefore, $|\pi' \cap A(1)| \leq |I' \cup \Phi| \leq |I \cup \Phi| = |\pi \cap A(1)|$.

We conclude that π overlaps maximally (among preferred extensions) with the approvals of voter 1. That is, $\mu(1) = |\pi \cap A(1)|$. Hence $\{\pi\}$ does 1-core-represent voter 1. For voter 2, since $x^* \in I$ we have $x^* \in \pi$ and therefore $|\pi \cap A(2)| = 1$. Since $\mu(2) = 1$, $\{\pi\}$ also 1-core-represents voter 2, concluding the “if” direction.

We continue with the “only if” direction. Assume there is an outcome $\Omega = \{\pi\} \subseteq prf(\text{Fram})$ that 1-core-represents all voters $i \in N = \{1, 2\}$. Let $I = \pi \cap X$.

Let $I' \subseteq X$ be an arbitrary model of φ . Such I' exists since φ is satisfiable. Let $\pi' = I' \cup \{\neg x : x \in X \setminus I'\} \cup \Phi$. By the same argument as above, $\pi' \in prf(\text{Fram})$. Since $\{\pi\}$ is an outcome that 1-core-represents voter 1, it must hold that

$$|X| + 1 = |\Phi| \leq |I' \cup \Phi| = |\pi' \cap A(1)| \leq |\pi \cap A(1)| = |\pi \cap (X \cup \Phi)| \leq |I \cup \Phi|.$$

By $|X| + 1 \leq |\pi \cap A(1)|$ we have $\varphi_\ell \in \pi$ for at least one $\varphi_\ell \in \Phi$. Thus, π defends φ_ℓ from all clause-arguments $\gamma \in \Gamma$, meaning that I satisfies all clauses in φ . Consequently I models φ . Furthermore, $|I' \cup \Phi| \leq |I \cup \Phi|$ implies $|I'| \leq |I|$. Since I' was chosen as an arbitrary model of φ , we conclude that I is a cardinality-maximal model of φ .

To finish the proof, observe that since voter 2 is 1-core-represented by $\{\pi\}$ and since $\mu(2) = 1$, it must be that $x^* \in \pi$. Therefore $x^* \in I$. $\square$

If the number of arguments or extensions in our given **ABSAF** is small, however, we can deal with representability more efficiently. Formally, this problem is **FPT** with respect to the number of arguments, meaning that, on inputs where this parameter is small, it is tractable.

Proposition 5.7. 1-REPRESENTABILITY and 1-CORE-REPRESENTABILITY are **FPT** with respect to the number of arguments in the given **ABSAF**.

Proof. Let $\text{Disc} = (\text{Fram}, N, A)$ and $k \in [n]$ and recall that $|C| = m$.

For regular 1-representation, we enumerate explicitly all outcomes $\Omega \subseteq \text{prf}(\text{Fram})$ of size k and check whether any such outcome 1-represents all voters. Letting $\ell = |\text{prf}(\text{Fram})|$, there are $\binom{\ell}{k}$ of these outcomes. Since $\binom{\ell}{k} \in O(\ell^k)$, this can be done in $O^*(\ell^k)$ -time. Moreover, $k \leq \ell$ (since $\Omega \subseteq \text{prf}(\text{Fram})$ for any outcome Ω) and furthermore $\ell \in O^*(3^{m/3})$ [Dunne et al., 2015]. Combining all, we get a runtime in

$$O^* \left(\left(3^{m/3} \right)^{\left(3^{m/3} \right)} \right)$$

from which **FPT** in m follows.

For core-representation, we first need to determine $\mu(i)$ for all $i \in N$ by enumerating every $\pi \in \text{prf}(\text{Fram})$ and computing $|\pi \cap A(i)|$. This can be done in $O^*(3^{m/3})$ time [Gaspers and Li, 2019]. Then, $\mu(i)$ is the size of the largest intersection. The rest of the algorithm is as in the previous case. $\square$

We have seen that we generally cannot 1-represent all voters, but we can 1-core-represent them. If the **AF** is small enough, we can even do so efficiently. On the other hand, when dealing with larger instances, we may require a huge number of preferred extensions to 1-(core-)represent all voters which does not help to understand the discussion. In this case, we might instead aim for an optimal representation with a fixed number of preferred extensions. Selecting a small, representative set of extensions is a problem with a distinct **ABC**-voting flavour, and we will make use of this connection. This will be the theme of the next sections.

Note that, while we will typically focus on regular representation (Definition 5.5), the theoretical results in the remainder of the chapter also apply to core-representation (Definition 5.7).

5.5 Optimising Representation

Let us now consider how to optimally represent voters in an **ABSAF** with a fixed number of preferred extensions. Concretely, the task is to select an outcome Ω of size k based on the representation it induces over N .

5.5.1 Global-Optimisation Methods

Looking back at the **ABC** voting rules introduced in Chapter 2 (Thiele and **OWA** methods), we can think of adapting them to this setting, as we have done in Chapter 4 for multi-issue elections. Here, we shall focus on **OWA** rules. This is motivated by the specific representation model that we use. Indeed, $rep_i(\Omega)$ is a real-valued quantity. Extending **OWA** rules to this setting is immediate, as **OWA** operators were originally introduced to aggregate real-valued vectors [Yager, 1988]. On the other hand, adapting Thiele methods is less straightforward and will therefore be postponed to future work.

We now introduce a class of **OWA** methods to select a set of k representative extensions. As argued above, the definition is essentially the same as Definition 2.3, save for the different “utility” (representation) model, and for the outcome we want to select.

Definition 5.9. Given an **ABSAF** $\text{Disc} = (\text{Fram}, N, A)$ and the outcome $\Omega \subseteq \text{sem}(\text{Fram})$, we define $s(A, \Omega)$ as the vector $(rep_1(\Omega), \dots, rep_n(\Omega))$ sorted in nondecreasing order. Given a family of weight functions $\omega^{(n)}$, the **OWA rule** corresponding to $\omega^{(n)}$ is then

$$\text{OWA}_{\omega^{(n)}}(A) \in \arg \max_{\Omega \subseteq \text{sem}(\text{Fram}) : |\Omega| \leq k} \left(s(A, \Omega) \cdot \left(\omega^{(n)}(1), \dots, \omega^{(n)}(n) \right) \right)$$

where k is the number of outcomes we want to select.

Note that these rules, as their **ABC** voting counterparts, are naturally designed to allow for ties. For the purpose of this chapter, we ignore this issue and assume that each rule is equipped with a tiebreaking mechanism, just as we did in Chapter 4.

Let us give some names to specific members of this family. In all these cases, the weight function is independent of n .⁸ For the first two, we use similar names as in the previous chapter, where we studied analogous **OWA** rules.

- The *utilitarian rule* is given by $\omega^{(n)} = (1, 1, 1, \dots)$ and picks the outcomes maximising average representation.
- The *egalitarian rule* is given by $\omega^{(n)} = (1, 0, 0, \dots)$ and picks the outcome that maximises the representation of the least-represented voter.

⁸In fact, nowhere in this chapter does this parametrisation of $\omega^{(n)}$ on n play a role. We keep this notation for the sake of consistency with Chapters 2 and 4.

- The *harmonic rule* is given by $\omega^{(n)} = (1, 1/2, 1/3, \dots)$ and can be seen as lying between utilitarian and egalitarian.

The utilitarian rule is fully majoritarian, while the egalitarian pledges to represent even minorities consisting of single individuals. The harmonic rule, defined by the harmonic sequence, is reminiscent of the [PAV](#) rule in the [ABC](#) voting model (Section 2.2.5), and has indeed similar conceptual underpinnings—giving less and less weight to voters as they get more and more represented. However, [PAV](#) is a Thiele rule, while the harmonic rule we consider here is an [OWA](#) rule.

A plausible alternative to the [OWA](#) approach is to simply maximise the number of voters that are 1-represented. We call this rule [MAXCOV](#), defined as:

$$\text{MAXCOV}(\text{Disc}) \in \arg \max_{\Omega \subseteq \text{sem}(\text{Fram}): |\Omega| \leq k} |\{i \in N: \text{rep}_i(\Omega) = 1\}|.$$

Again, we will assume that some tiebreaking mechanism is in place. We will later show that this rule is not a part of the [OWA](#) class (see the discussion after Theorem 5.16). For now, let us see two examples.

Example 5.4. Consider an [ABSAF](#) with voters $N = [4]$ and extensions $\text{prf}(\text{Fram}) = \{\pi_1, \pi_2, \pi_3\}$. Say that we want to pick an outcome of size $k = 1$ and that the representation scores are as follows:

Ω	$\text{rep}_1(\Omega)$	$\text{rep}_2(\Omega)$	$\text{rep}_3(\Omega)$	$\text{rep}_4(\Omega)$
$\{\pi_1\}$	0	1	1	1
$\{\pi_2\}$	0.4	0.5	1	1
$\{\pi_3\}$	0.5	0.5	0.5	0.5

It can be verified that utilitarian and [MAXCOV](#) select $\{\pi_1\}$, harmonic selects $\{\pi_2\}$, and egalitarian selects $\{\pi_3\}$.

Example 5.5. When selecting outcomes of size $k = 2$ for the Canada electoral reform data (Example 5.2), utilitarian and [MAXCOV](#) pick, e.g.,

$$\Omega_1 = \{\{p_1, p_2, p_3, s_1, s_2\}, \{f_2, p_1, p_2, m_1\}\}$$

while harmonic and egalitarian pick, e.g.,

$$\Omega_2 = \{\{p_1, p_2, p_3, s_1, s_2\}, \{f_1, f_2, m_1\}\}.$$

The selected outcomes are the same for core-representation. The minimum and average (core-)representation values across all voters are:

	min	avg	min (core)	avg (core)
Ω_1	0	0.9378	0	0.9706
Ω_2	0.5	0.9360	0.5	0.9697

The average representation for utilitarian is only marginally larger than for egalitarian and harmonic in this instance.

Note that all these rules achieve perfect representation if possible.

Proposition 5.8. Every OWA rule, as well as MAXCOV, returns an outcome of size k in which every voter is 1-represented whenever such an outcome of size k exists. In particular, the egalitarian rule returns an outcome of size k in which every voter is α -represented whenever such an outcome of size k exists.

Proof. Suppose that there is an outcome Ω of size k that α -represents all voters (for some $\alpha \in [0, 1]$). Assume towards a contradiction that egalitarian chooses an outcome Ω' that does not α -represent all voters. Thus, the least-satisfied voter will have satisfaction $\beta < \alpha$. But with outcome Ω the least-satisfied voter would have higher satisfaction (at least α). Contradiction: egalitarian must choose an outcome that α -represents all voters.

Next, suppose that there is an outcome Ω of size k that 1-represents all voters. It follows from definition that MAXCOV must select this outcome or any other outcome that 1-represents all voters. Now, consider a generic rule $\text{OWA}_{\omega^{(n)}}$. Consider a candidate outcome Ω' that does not 1-represent all voters. Let

$$s(A, \Omega) = (s_1, \dots, s_n) \quad \text{and} \quad s(A, \Omega') = (s'_1, \dots, s'_n).$$

be the respective sorted representation vectors and observe

$$s(A, \Omega) = (1, \dots, 1).$$

We get

$$s(A, \Omega) \cdot (\omega^{(n)}(1), \dots, \omega^{(n)}(n)) > s(A, \Omega') \cdot (\omega^{(n)}(1), \dots, \omega^{(n)}(n)) \iff \sum_{i \in [n]} (1 - s'_i) \omega^{(n)}(i) > 0.$$

Recall that $\omega^{(n)}(\ell)$ is nonnegative for all ℓ , that $\omega^{(n)}(1) > 0$, and that $s'_i \in [0, 1]$ for all $i \in N$. Furthermore, since Ω' does not 1-represent all voters, we have $s'_1 < 1$. These facts imply that the latter inequality holds. Hence, Ω has higher OWA score than Ω' , meaning that $\text{OWA}_{\omega^{(n)}}$ cannot pick Ω' . This concludes the proof. $\square$

Moreover, all of the rules considered here are computationally intractable.⁹ This follows from Proposition 5.2 and Proposition 5.8.

⁹Note that the problem of computing an optimal outcome for a given rule is an optimisation problem, not a decision problem. We argue that these optimisation problems are hard by showing that we can not solve them in polynomial time unless $P = NP$.

Proposition 5.9. If $P \neq NP$, it is not possible to compute an outcome that is optimal with respect to a given **OWA** rule or **MAXCOV** in polynomial time.

Proof. Recall that deciding whether all voters in a given **ABSAF** can be 1-represented is NP-hard even if there is only a single voter that approves a single argument, i.e., $A(1) = \{a\}$ (see the proof of Proposition 5.2). An outcome Ω either 1-represents the voter (if $a \in \pi$ for some $\pi \in \Omega$) or 0-represents the voter (otherwise). By Proposition 5.8, all considered rules will return an outcome that 1-represents the voter if such an outcome exists. Thus, the rules cannot be computed in polynomial time unless $P = NP$. $\square$

While intractable, **OWA** rules and **MAXCOV** are **FPT**.

Proposition 5.10. Computing the outcome of an **OWA** rule or of **MAXCOV** is **FPT** with respect to the number of arguments in the given **ABSAF**.

Proof. Analogously to Proposition 5.7, we can first enumerate all preferred extensions in $O^*(3^{m/3})$ time (where $|C| = m$) and then enumerate all outcomes of size k in $O^*(\ell^k)$ time (where $\ell = |\text{prf}(\text{Fram})|$), allowing us to choose the optimal outcome for the rule. $\square$

In practice, however, a runtime of $O^*(\ell^k)$ can still be prohibitive. Thus, we introduce sequential variants for all of the above rules, as done in Chapters 2 and 4.

5.5.2 Sequential Methods

We have already defined the sequential formulation of **OWA** rules in Chapters 2 and 4; we repeat it here expressed in the current notation.

Definition 5.10. Assuming we already have selected ℓ viewpoints $\pi_1, \dots, \pi_\ell$, the *sequential OWA rule* picks the $(\ell + 1)$ -th viewpoint from the set maximising the following objective:

$$\pi_{\ell+1} \in \arg \max_{\pi \in \text{sem}(\text{Fram}) \setminus \{\pi_1, \dots, \pi_\ell\}} s(A, \{\pi_1, \dots, \pi_\ell, \pi\}) \cdot \left(\omega^{(n)}(1), \dots, \omega^{(n)}(n) \right).$$

We stop as soon as k viewpoints have been selected. Again, some tiebreaking is assumed.

The sequential variant of **MAXCOV**, called **SEQCOV**, is defined analogously. We pick, at every step, the viewpoint that (together with the previously selected ones) 1-represents as many voters as possible, i.e., a viewpoint $\pi_{\ell+1}$ with

$$\pi_{\ell+1} \in \arg \max_{\pi \in \text{sem}(\text{Fram}) \setminus \{\pi_1, \dots, \pi_\ell\}} |\{i \in N : \text{rep}_i(\{\pi_1, \dots, \pi_\ell, \pi\}) = 1\}|.$$

Observe that **SEQCOV** approximates **MAXCOV** with a factor of $1 - 1/e \approx 0.63$, which follows directly from the approximation guarantee of the greedy algorithm for the

maximum coverage problem [Hochbaum, 1997].

Although we introduced sequential rules to gain a computational advantage, disturbingly, just like their non-sequential variants, these rules are also intractable—the proof for Proposition 5.9 works verbatim to show the following.

Proposition 5.11. If $P \neq NP$, it is not possible to compute an outcome that is optimal with respect to a given sequential OWA rule or SEQCOV in polynomial time.

Why, then, do we bother introducing these sequential formulations? Importantly, we can improve upon the FPT algorithm of Proposition 5.10 since we do not need to enumerate all outcomes. Rather, it suffices to enumerate all preferred extensions whenever we pick a new viewpoint in the procedure. Concretely, we can first pre-compute the preferred extensions in $O^*(3^{m/3})$ time [Dvořák and Dunne, 2018]. Then, we can pick k of them sequentially in $O^*(|prf(\text{Fram})| \cdot k)$ time. What is more, if we assume that we can perform the pre-computation efficiently in practice (for example, by relying on highly performant argumentation solvers [Egly et al., 2008, Niskanen and Järvisalo, 2020]), we may abstract over the $O^*(3^{m/3})$ cost. Under this assumption, the procedure can be done in $O^*(|prf(\text{Fram})| \cdot k)$ time, which is polynomial in the number of preferred extensions (since this quantity bounds k). This is better than the algorithm of Proposition 5.10, which had runtime $O^*(|prf(\text{Fram})|^k)$.

While this abstraction treats the solver runtime as constant and is therefore coarse from a theoretical standpoint, it reflects practical performance (this will be confirmed via simulations in Section 5.7) and still yields a meaningful complexity-theoretic improvement over non-sequential rules. In particular, assuming $P \neq NP$, the same polynomial algorithm in the number of extensions does not exist for the non-sequential OWA rules we named earlier, even under the same assumption about the cost of pre-computing the preferred extensions.

Proposition 5.12. Assuming $P \neq NP$, there is no algorithm running in polynomial time with respect to $|prf(\text{Fram})|$ such that, given an ABSAF Disc = (Fram, N , A) and set $prf(\text{Fram})$ as inputs, returns an outcome for Disc that is optimal with respect to the utilitarian, egalitarian, or harmonic rule.

Proof. For the utilitarian and harmonic rules we reduce from the problem of computing the outcome of the CC rule in the standard ABC voting model, which is NP-hard [Procaccia et al., 2008]. Consider an instance of this problem, given by candidates C (with $|C| = m$), voters N , and approval ballots A . Furthermore, k is the target committee size. We construct the ABSAF Disc = (Fram, N , A^*) (where Fram = (C^* , $\rightarrow$)) with

- $C^* = C \cup D$, where $D = \{c' : c \in C\}$;
- $\rightarrow = \{(c, d) : (c, d) \in C^2 \text{ and } c \neq d\} \cup \{(c', c') : c' \in D\}$;

- $A^*(i) = A(i) \cup \{c' : c \in C \setminus A(i)\}$ for every $i \in N$.

In other words, we turn the candidate set C in a complete graph (all candidates attacking all other candidates, save for themselves), and we add a self-attacking clone c' of every candidate c . Moreover, voters extend their approvals with the self-attacking clones of all the candidates they originally did not approve.

All the members of D are self-attacking, and hence cannot be featured in any preferred extensions. Moreover, all members of C attack all the other members of C . Consequently $\text{prf}(\text{Fram}) = \{\{c\} : c \in C\}$ and thus $|\text{prf}(\text{Fram})| = m$. Moreover, $|A^*(i)| = m$ for every $i \in N$. Thus, for every $\Omega \subseteq \text{prf}(\text{Fram})$, if there is at least one $\{c\} \in \Omega$ such that $c \in A^*(i)$, we have

$$\text{rep}_i(\Omega) = \max_{\pi \in \Omega} \text{rep}_i(\pi) = \frac{|\pi \cap A(i)|}{|A(i)|} = \max_{c \in C} \frac{\mathbb{1}[c \in A^*(i)]}{m} = \frac{1}{m}$$

and analogously $\text{core}_i(\Omega) = 1$. Otherwise, $\text{rep}_i(\Omega) = \text{core}_i(\Omega) = 0$. Hence, $W \subseteq C$ with $|W| = k$ maximises $|\{i \in N : A(i) \cap W \neq \emptyset\}|$ (the **CC** objective) if and only if $\Omega = \{\{c\} : c \in W\}$ is optimal with respect to the utilitarian and harmonic rules. Now, if there was an algorithm that could compute Ω in polynomial time with respect to $|C| = |\text{prf}(\text{Fram})|$, then we could compute the **CC** outcome in polynomial time, which would imply $P = NP$.

For the egalitarian rule, let $(U, S_1, \dots, S_\ell, k)$ be an instance of the prototypical NP-complete problem hitting set [Karp, 1972]. Specifically, U is a set of s elements, $S_i \subseteq U$ for every $i \in [s]$, and k is a natural number. Set $H \subseteq U$ is a hitting set if and only if $H \cap S_i \neq \emptyset$ for every $i \in [s]$. We must decide whether there is a hitting set H of size $|H| \leq k$.

The reduction is similar to the previous one. We construct the **ABSAF** Disc = (Fram, N, A) with $\text{Fram} = (C, \rightarrow)$ where

- $C = U \cup T$, where $T = \{c' : c \in U\}$;
- $\rightarrow = \{(c, d) : (c, d) \in U^2 \text{ and } c \neq d\} \cup \{(c', c') : c' \in T\}$;
- $A(i) = S_i \cup \{c' : c \in U \setminus S_i\}$ for every $i \in N$.

As the previous case, one can show that $\text{prf}(\text{Fram}) = \{\{c\} : c \in U\}$ and that, for every $\Omega \subseteq \text{prf}(\text{Fram})$, if there is $\{c\} \in \Omega$ such that $c \in A(i)$ then we have $\text{rep}_i(\Omega) = 1/s$ and $\text{core}_i(\Omega) = 1$; otherwise, $\text{rep}_i(\Omega) = \text{core}_i(\Omega) = 0$.

This implies that $H \subseteq U$ is a hitting set if and only if, for $\Omega = \{\{c\} : c \in H\}$, we have $\min_{i \in N} \text{rep}_i(\Omega) = 1/s$ (respectively, $\min_{i \in N} \text{core}_i(\Omega) = 1$). Thus, if there was an algorithm that could compute an optimal outcome for the egalitarian rule in polynomial time with respect to $s = |\text{prf}(\text{Fram})|$, then we could compute a minimal hitting set in polynomial time, which would imply $P = NP$. $\square$

We have now introduced several rules that aim at optimising representation under the constraint of picking k viewpoints. Next, we will compare them axiomatically and in numerical simulations.

5.6 Justified Representation

In this section, we will compare the rules we introduced via the axiomatic method, a standard approach in social choice theory.

Recall the well-known **JR** axiom from **ABC** voting (Definition 2.6). This axiom requires each sufficiently large group of voters with similar opinions to be represented in the outcome. This is also a natural desideratum for **ABSAFs**; we will thus translate it to this setting. To formalise it, we first need to define what it means for a group to have similar opinions. The following definition is analogous to that of cohesive groups in the **ABC** voting model (Definition 2.5).

Definition 5.11. Let $\text{Disc} = (\text{Fram}, N, A)$ be an **ABSAF**. We say that a group of voters $N' \subseteq N$ is *1-representable* if and only if there is some $\pi \in \text{prf}(\text{Fram})$ such that all $i \in N'$ have $\text{rep}_i(\pi) = 1$.

Observe that Definition 5.11 could be written using the notion of core-representation (Definition 5.7) instead. Again, all the definitions and results that we present in this section work in either case.

Let us illustrate Definition 5.11 with an example.

Example 5.6. Let us consider again the **ABSAF** from Example 5.2, derived from real-world data. For convenience, we restate here the ballots and their frequencies:

$\{p_1\}$	33	$\{p_1, p_3\}$	9	$\{f_2, p_2, p_3\}$	3	$\{f_2, p_3\}$	1
$\{p_1, p_2, p_3\}$	31	$\{p_1, p_2\}$	8	$\{f_1, p_1\}$	2	$\{m_1\}$	1
$\{p_2\}$	16	$\{f_2, p_1, p_2, p_3\}$	7	$\{f_2, p_1, p_2\}$	2	$\{m_1, p_2\}$	1
$\{p_3\}$	16	$\{s_2\}$	6	$\{f_2, p_1, p_3\}$	2	$\{p_1, p_2, p_3, s_1, s_2\}$	1
$\{f_2\}$	11	$\{f_1, f_2\}$	4	$\{p_1, p_2, p_3, s_1\}$	2	$\{p_2, p_3, s_2\}$	1
$\{p_2, p_3\}$	10	$\{m_1, p_1, p_2, p_3, s_1\}$	4	$\{f_1, f_2, p_3\}$	1	$\{p_2, s_1\}$	1
$\{f_2, p_1\}$	9	$\{f_1\}$	3	$\{f_1, p_1, p_2, p_3\}$	1	$\{s_1\}$	1

Here, the 33 voters who approve $\{p_1\}$ form a 1-representable group, as they can all be perfectly represented by the preferred extension $\{p_1, p_2, p_3, s_1, s_2\}$ (among other extensions). However, this group is not maximal, as the same preferred extension perfectly represents other voters. Indeed, the largest (in terms of cardinality) 1-representable group is the group of

$$33 + 31 + 16 + 16 + 10 + 9 + 8 + 6 + 2 + 1 + 1 + 1 + 1 = 135$$

voters who are perfectly represented by the preferred extension $\{p_1, p_2, p_3, s_1, s_2\}$. On the other hand, there is no 1-representable group that contains both the 33 voters approving $\{p_1\}$ and the 3 voters approving $\{f_1\}$, as no preferred extension can contain those two (conflicting) arguments.

If a 1-representable group contains at least n/k voters, with k being the number of viewpoints we want to select, then this group, arguably, deserves to be fully represented. Indeed, if n' is the size of this group, then $n' \geq n/k$ means $n'/n \geq 1/k$. In other words, this group proportionally deserves one “seat” in the output, and can unanimously propose an extension to occupy this seat: the one that perfectly represents them all.

The following axiom captures this. A warning about terminology: we use the term “deliberative” to distinguish this axiom from the **JR** axiom in the **ABC** voting model.¹⁰ Moreover, the role of the term “strong” will be clear later.

Definition 5.12. An outcome $\Omega \subseteq \text{prf}(\text{Fram})$ of size $|\Omega| = k$ satisfies *Strong Deliberative Justified Representation (SDJR)* if for every 1-representable group $N' \subseteq N$ of size $|N'| \geq n/k$ there is a viewpoint $\pi \in \Omega$ such that $\text{rep}_i(\pi) = 1$ for every $i \in N'$.

As can be seen quite easily, if the approval sets of all voters are preferred extensions, then **SDJR** can be satisfied.

Proposition 5.13. Let $\text{Disc} = (\text{Fram}, N, A)$ be an **ABSAF**. If $A(i) \in \text{prf}(\text{Fram})$ for all $i \in N$, then for any $k \in \mathbb{N}$ with $k \leq |\text{prf}(\text{Fram})|$ we can find an outcome Ω of size $|\Omega| = k$ that satisfies **SDJR**.


Proof. Fix an **ABSAF** $\text{Disc} = (\text{Fram}, N, A)$. For every two preferred extensions $\{\pi, \pi'\} \subseteq \text{prf}(\text{Fram})$ we know that $\pi \not\subseteq \pi'$ and *vice versa*. Hence, if the approval set $A(i)$ of a voter $i \in N$ is a preferred extension π , then i is only 1-represented by π . Let ℓ be the number of subset-maximal 1-representable groups of size at least n/k . As every voter can only be 1-represented by one specific extension, these groups must be disjoint; consequently, $\ell \leq k$, because $(k+1) \cdot n/k > n$. It follows that we can just select for every such group the preferred extension that 1-represents all its voters. $\square$

Unfortunately, once we drop this unreasonably strong assumption, the result does not hold. Indeed, it ceases to hold even if we assume admissibility in the approvals (which is already a rather strong assumption).

Proposition 5.14. It cannot be guaranteed that, given an **ABSAF** $\text{Disc} = (\text{Fram}, N, A)$ and $k \in [n-1]$, there is an outcome Ω of size k that satisfies **SDJR**, even if votes are assumed to be admissible.

¹⁰In the paper on which this chapter is based [Bernreiter et al., 2024a], we did not use the term “deliberative”, we just called this axiom “strong justified representation” (and similar for the axiom we will introduce later). In the paper, we did not need to distinguish our setting from the **ABC** voting setting.

Proof. Let $\text{Disc} = (\text{Fram}, N, A)$ be the [ABSAF](#) with framework



and approvals

$$A = \{\{a\}, \{a, c\}, \{a, d\}, \{a, e\}\}.$$

There are three preferred extensions, namely $\text{prf}(\text{Fram}) = \{\{a, c\}, \{a, d\}, \{a, e\}\}$. However, we cannot pick an outcome $\Omega \subseteq \text{prf}(\text{Fram})$ of size $k = 2$ that satisfies [SDJR](#). Indeed, the voter groups $\{1, 2\}$, $\{1, 3\}$, and $\{1, 4\}$ are all 1-representable and have a size of $2 = n/k$. However, if we pick the outcome $\{\{a, c\}, \{a, d\}\}$ we have that $\text{rep}_4(\pi) = 1/2$ despite voter 4 being part of a 1-representable group of size 2. Something analogous holds even if we pick $\{\{a, c\}, \{a, e\}\}$ or $\{\{a, d\}, \{a, e\}\}$. $\square$

This motivates us to introduce the following weakening of [SDJR](#), where we only require that *at least one* voter of every sufficiently large 1-representable group must be 1-represented. This is analogous to what is required in the [JR](#) axiom for the [ABC](#) voting model.

Definition 5.13. An outcome $\Omega \subseteq \text{prf}(\text{Fram})$ of size k satisfies [Deliberative Justified Representation \(DJR\)](#) if and only if for every 1-representable group $N' \subseteq N$ with $|N'| \geq n/k$ there is a $\pi \in \Omega$ such that $\text{rep}_i(\pi) = 1$ for some $i \in N'$.

Note that, for an outcome Ω of size k to satisfy [DJR](#), the amount of voters in any 1-representable group that are not fully represented by Ω must be smaller than n/k . Indeed, if n/k or more voters in N' were not 1-represented by Ω , then

$$\{i \in N' : \text{rep}_i(\pi) < 1 \text{ for all } \pi \in \Omega\}$$

would be a 1-representable group of size n/k that violates [DJR](#). Thus, even though this axiom only requires one voter per 1-representable group to be fully represented, the above argument places an upper-bound on the number of underrepresented voters in 1-representable groups.

More importantly, in contrast to [SDJR](#), this axiom can always be satisfied. Indeed, it can be satisfied by just maximising the number of 1-represented voters.

Theorem 5.15. [MAXCOV](#) and [SEQCOV](#) satisfy [DJR](#).

Proof. We begin with some notation. Given any set of viewpoints $\Omega \subseteq \text{prf}(\text{Fram})$, we write

$$N_\Omega = \{i \in N : \text{rep}_i(\pi) = 1 \text{ for some } \pi \in \Omega\}$$

to indicate the set of voters perfectly represented by Ω .

Now, let Ω be the outcome of either MAXCOV and SEQCOV with size $|\Omega| = k$. Assume that there is a 1-representable group N' of size at least n/k that is represented by some viewpoint $\pi^* \in \text{prf}(\text{Fram}) \setminus \Omega$ such that no voter in N' is 1-represented by Ω . Thus, for the set of 1-represented voters N_Ω we have $|N_\Omega| \leq n - |N'| \leq n - n/k$. Now, fix an enumeration $\pi_1, \pi_2, \dots, \pi_k$ of the elements of Ω . In the case of MAXCOV, let the enumeration be arbitrary, while for SEQCOV, let this be the order in which the elements of Ω are picked. Now partition N_Ω into the disjoint sets $N_\Omega = N_{\pi_1} \cup \dots \cup N_{\pi_k}$ with

$$N_{\pi_j} = \{i \in N_\Omega : \text{rep}_i(\pi_j) = 1 \text{ and } \text{rep}_i(\pi_\ell) < 1 \text{ for all } \ell \in [j-1]\}.$$

In other words, N_{π_j} is the set of voters for which j is the minimal index for which π_j 1-represents them. As this is a partition we have

$$\sum_{j=1}^k |N_{\pi_j}| = |N_\Omega| \leq n - \frac{n}{k} < n$$

and therefore there must be $\ell \leq k$ for which $|N_{\pi_\ell}| < n/k \leq |N'|$ holds. Fix this value ℓ .

Assume first that Ω was chosen by MAXCOV. Observe that

$$|N_\Omega| = \sum_{j=1}^k |N_{\pi_j}| < \left(\sum_{j=1}^k |N_{\pi_j}| \right) - |N_{\pi_\ell}| + |N'| \leq |N_{\Omega \cup \{\pi^*\} \setminus \{\pi_\ell\}}|$$

where the last inequality follows from the fact that no voter in N' is represented by Ω . However, this contradicts the assumption that Ω was chosen by MAXCOV. Assume now that Ω was chosen by SEQCOV. Then, in round ℓ , we have


$$|N_{\{\pi_1, \dots, \pi_\ell\}}| = \sum_{j=1}^{\ell} |N_{\pi_j}| < \left(\sum_{j=1}^{\ell-1} |N_{\pi_j}| \right) + |N'| = |N_{\{\pi_1, \dots, \pi_{\ell-1}, \pi^*\}}|$$

where the last inequality follows again from the fact that no voter in N' is represented by Ω . This contradicts the assumption that π_ℓ was chosen by SEQCOV in round ℓ . $\square$

By contrast, rules that do not explicitly maximise the number of 1-represented voters may not satisfy DJR.

Theorem 5.16. No OWA or sequential OWA rule is guaranteed to return an outcome that satisfies DJR, irrespective of the tiebreaking used, and even if all ballots are admissible.

Proof. Fix a weight function $\omega^{(n)}$ and let $k = 2$. Consider the **ABSAF** with framework



and approvals

$$A = \{\{a, b, c, d\}, \{a, b, c, d\}, \{e, f, g, h\}, \{e, i, j, k\}\}.$$

One can check that the preferred extensions are

$$\begin{aligned}\pi_1 &= \{a, b, c, d\}, \\ \pi_2 &= \{a, b, c, e, f, g, h\}, \text{ and} \\ \pi_3 &= \{a, b, c, e, i, j, k\}.\end{aligned}$$

Note that there is a 1-representable group of size at least $n/k = 2$, namely $N' = \{1, 2\}$, which we can perfectly represent only with π_1 . Indeed, $\pi_1 = A(1) = A(2)$. Thus, to satisfy **DJR**, we need to pick an outcome containing π_1 .

We first show that the outcome returned by the **OWA** rule defined by $\omega^{(n)}$ does not satisfy **DJR**. By considering all outcomes, it can be verified that the possible sorted representation vectors are

$$\begin{aligned}s(A, \{\pi_1, \pi_2\}) &= s(A, \{\pi_1, \pi_3\}) = \left(\frac{1}{4}, 1, 1, 1\right) \text{ and} \\ s(A, \{\pi_2, \pi_3\}) &= \left(\frac{3}{4}, \frac{3}{4}, 1, 1\right).\end{aligned}$$

Thus, for the **OWA** score of $\{\pi_1, \pi_2\}$ (equivalently, $\{\pi_1, \pi_3\}$) to be lower than the **OWA** score of $\{\pi_2, \pi_3\}$ we must have

$$\begin{aligned}\frac{1}{4}\omega^{(n)}(1) + \omega^{(n)}(2) + \omega^{(n)}(3) + \omega^{(n)}(4) &< \frac{3}{4}\omega^{(n)}(1) + \frac{3}{4}\omega^{(n)}(2) + \omega^{(n)}(3) + \omega^{(n)}(4) \\ &\iff \omega^{(n)}(2) < 2\omega^{(n)}(1).\end{aligned}$$

The last inequality holds by the fact that $(\omega^{(n)}(1), \dots, \omega^{(n)}(n))$ is a nonincreasing vector and $\omega^{(n)}(1) > 0$. Thus, outcome $\{\pi_2, \pi_3\}$ is chosen, failing **DJR**.

We now show that the sequential formulation of the same rule also fails [DJR](#). In the first round, we have the sorted representation vectors

$$\begin{aligned} s(A, \{\pi_1\}) &= (0, 0, 1, 1) \text{ and} \\ s(A, \{\pi_2\}) &= s(A, \{\pi_3\}) = \left(\frac{1}{4}, \frac{3}{4}, \frac{3}{4}, 1\right). \end{aligned}$$

Consequently,

$$\begin{aligned} s(A, \{\pi_1\}) \cdot (\omega^{(n)}(1), \dots, \omega^{(n)}(n)) &< s(A, \{\pi_2\}) \cdot (\omega^{(n)}(1), \dots, \omega^{(n)}(n)) \iff \\ &\omega^{(n)}(3) < \omega^{(n)}(1) + 3\omega^{(n)}(2). \end{aligned}$$

Again, the last inequality holds by definition of $\omega^{(n)}$. Hence, the rule will either pick $\{\pi_2\}$ or $\{\pi_3\}$. Assume that $\{\pi_2\}$ is chosen (the other case is analogous). Then, in the second (and final) round, the rule will only compare outcomes containing π_2 , namely $\{\pi_1, \pi_2\}$ and $\{\pi_2, \pi_3\}$. We have shown while treating the non-sequential case that the latter has a higher [OWA](#) score than the former. Thus, $\{\pi_2, \pi_3\}$ is chosen, failing [DJR](#). $\square$

Theorems [5.15](#) and [5.16](#) together show that [MAXCOV](#) (respectively, [SEQCOV](#)) does not belong to the class of [OWA](#) rules (respectively, sequential [OWA](#) rules).

More importantly, these results show that [MAXCOV](#) and [SEQCOV](#) outperform all [OWA](#) and sequential [OWA](#) rules with respect to the key axiom of [DJR](#). Whether this theoretical advantage holds up under empirical scrutiny is a question we address in the next and final section.

5.7 Numerical Simulations

In this section, we present our empirical findings. We tested our rules on synthetic data and found that, under the metrics considered, the utilitarian and harmonic rules perform well. Moreover, their sequential formulations scale to practically relevant input sizes while retaining a reasonable approximation of their non-sequential counterparts.

5.7.1 Setup

In this section, we detail the setup (implementation, data, metrics) of our simulations.

5.7.1.1 Implementation Sketch

Since some of our results concern the runtime performance of our algorithms, we sketch the implementation of our rules and our hardware. For full details refer to the code, which is publicly available online [[Bernreiter et al., 2024b](#)].

Given an [ABSAF](#), first, we pre-compute all the preferred extensions. To do so, we use [ASPARTIX](#) [[Egly et al., 2008](#)], a powerful argumentation solver based on answer set programming [[Gebser et al., 2022](#)]. For core-representation, we additionally pre-compute

$\mu(i)$ for all $i \in N$ by iterating over all preferred extensions and picking the one that maximally overlaps with $A(i)$.

The procedure is then analogous for both OWA and coverage-maximising rules. For non-sequential rules, we simply brute-force and iterate over all subsets of $\text{prf}(\text{Fram})$ of size k , and pick one maximising the objective function (be it the OWA score or the number of perfectly represented voters). For sequential rules, we proceed in k steps. At each step, we iterate over all preferred extensions, and pick an unselected one that maximises, together with the partial outcome selected thus far, the objective.

We implemented our framework in Python, and ran our simulations on a device equipped with two Intel Xeon E5-2650 v4 (12 cores @ 2.20GHz) CPUs and 256GB of RAM.

5.7.1.2 Data Generation

Since there is no dataset of ABSAFs, we need to generate our data. Note that we cannot take advantage of datasets consisting of plain-text discussions with user approvals,¹¹ as there are still challenges in mining AFs from text [Baumann et al., 2020], though breakthroughs in large language models might make this challenge solvable in the near future [Gorur et al., 2025]. Moreover, we also cannot use existing benchmarks of AFs (and only generate the votes), as these involve very large AFs.¹² This is beyond our scope, as we want to focus on more moderately-sized discussions. Hence, we have to generate both the AFs and the ballots.

Sampling the frameworks. To sample the frameworks, we use AFBENCH-GEN2 [Cerutti et al., 2016], a tool for generating AFs that implements various graph-generation methods. We use the Barabási-Albert sampling algorithm [Barabási and Albert, 1999], which produces scale-free graphs, i.e., graphs whose degree distribution follows a power law. According to Barabási and Albert, several real-world networks exhibit this feature, including, for instance, social networks. This is the result of the “preferential attachment” mechanism, where new nodes are more likely to be attached to existing nodes with many edges. Note that Barabási-Albert graphs are, by default, acyclic; this is unsuitable for AFs, as acyclic AFs are trivial.¹³ Hence, with AFBENCH-GEN2, users can specify the probability of a node being in a cycle, and the tool will add extra attacks accordingly. We will later refer to this probability as p_{cyc} , and, as always, to the number of arguments as m .

In our simulations, we are interested in additionally controlling the amount of preferred extensions in the generated AFs. There are two reasons for this. On the

¹¹See, e.g., the “Computational Democracy Project”, from which we took Example 5.2: <https://compdemocracy.org/> (accessed: 2026-04-17).

¹²See, e.g., <http://argumentationcompetition.org/> (accessed: 2026-04-17).

¹³Observe that acyclic AFs always have a unique preferred extension. Intuitively, this extension can be constructed iteratively by accepting all arguments with no incoming attacks, removing the arguments they attack, and repeating the process. An explicit proof is given in the original paper by Dung [1995, Theorem 30].

one hand, we need enough preferred extensions to make the problem of choosing an outcome nontrivial. On the other hand, we need to limit the number of extensions for performance reasons. As there is no default option to generate **AFs** with a given number of extensions in **AFBENCHGEN2**, we employ rejection sampling. Concretely, we generate a large number of **AFs** (with varying arguments) and then filter out those with a desirable number of extensions.

In particular, we generate **AFs** with $m \in \{10, 11, \dots, 49, 50\}$ and $p_{\text{cyc}} \in \{1/4, 1/2, 3/4\}$. For each combination of these values, we generate 20 **AFs**. From this pre-generated set we then select a certain number of **AFs** (depending on the experiment) with the desired number of preferred extensions (also depending on the experiment).

Sampling the approvals. To sample the approvals, we use a mixed Mallows model [Mallows, 1957, Szufa et al., 2020], a popular method to sample approval elections. Here, there are some central votes (from now on, “ground truths”), of which the ballots are noisy signals. We might think of a ground truth as some rational viewpoint a voter might lean towards: each voter’s preference will be a (possibly irrational) approximation of such viewpoints.

Formally, let

$$d: 2^C \times 2^C \rightarrow \mathbb{R}_{\geq 0}$$

be a distance over sets of arguments. For a ground truth $S \subseteq C$, the probability of sampling ballot $S' \subseteq C$ is proportional to $\Phi^{d(S, S')}$, where $\Phi \in [0, 1]$ is called the dispersion parameter. A polynomial-time algorithm for sampling approval ballots with this model is described by Szufa et al. [2022]. To build intuition, observe that if $\Phi = 0$ then we can only sample ballots that are at zero distance from the ground truth. On the other hand, if $\Phi = 1$ then all approval ballots are equally likely to be sampled (the so-called “impartial statistical culture” [Szufa et al., 2022]).

In our model, we choose as distance measure the function $d(S, S') = |S'| - |S' \cap S|$, i.e., the number of arguments in S' that are not in S . Observe that our distance measure reflects our notion of representability: if a voter approves a subset of the ground truth, the distance is zero (as the ground truth fully “covers” her approved arguments).

Full sampling algorithm. We now describe the full generation algorithm. We select 50 **AFs** from the pre-generated set of frameworks described above, each of which has 10 preferred extensions. For every such **AF**, we choose uniformly at random 5 preferred extensions that will act as the ground truths. Then, for each ground truth, we sample 20 approval ballots centred around it via the Mallows model. This captures a scenario in which voters are divided into different “ideological camps” centred around different rational viewpoints. We repeat this process 5 times per **AF**, creating five **ABSAFs** per **AF**. The above generation is repeated in full for $\Phi \in \{0, 0.1, \dots, 0.9, 1\}$. We obtain, in the end, a total of $50 \cdot 11 \cdot 5 = 2750$ **ABSAFs**.

5.7.1.3 Description of the Simulation Experiments

All simulation experiments are conducted for both the base definition of representation (Definition 5.5) and core-representation (Definition 5.7). Here, we show results for the former; the others can be found in Appendix B. Likewise, for experiments where we set $k = 3$, in Appendix B we include experiments for $k \in \{2, 5\}$.

In our first simulation experiment, we study the effect of varying k when selecting a representative outcome. Given an OWA rule (defined by a certain score function $\omega^{(n)}$), we plot on the x -axis the values of the dispersion value Φ . For each value of Φ and each fixed k , we plot the mean $\omega^{(n)}$ -OWA score of the selected outcome of size k , averaged over all generated ABSAFs. We have one plot per each considered rule, namely, utilitarian, harmonic and egalitarian; one each plot, we have one separate curve for each value of $k \in [7]$. The error bars indicate one sample standard deviation above and below the mean. Results are in Figure 5.5, and will be discussed in the next section.

For the second experiment, we fix $k = 3$. This is a number slightly smaller than the number of true ground truths, which is 5, and this choice, as we will see, is motivated by the previous experiment. We compare the utilitarian, egalitarian, harmonic and MAXCOV rules with a baseline and two oracles. “Oracle”, here, is not to be understood in the complexity-theoretic sense we used thus far, but as the opposite of “baseline”. That is, as a rule that exhibits ideal but inaccessible performance. The baseline returns k preferred extensions of the input ABSAF selected uniformly at random. The first oracle returns all ground truths (from which the ballots were constructed) while the second returns all preferred extensions. To compare the rules we use the following metrics:

- the average and minimal representation score (over all voters) given by the selected outcome;
- the fraction of voters that are perfectly represented by the outcome;
- the “recovery score” of an outcome. This is computed as $r(\Omega)/k$, where $r(\Omega)$ is the number of ground truths included in Ω . Intuitively, this captures how well a rule can recover the ground truths around which the voters’ preferences centre.

As in the previous simulation experiment, we average the metrics over the dispersion parameters, and error bars indicate sample standard deviation. We run this over all the selected ABSAFs, plot the dispersion parameter on the x -axis, and the metric on the y -axis. We report the results in Figure 5.6 (again, we will discuss the results in the next section).

Finally, we turn our attention to sequential rules. We perform two experiments: in the first one, we assess the approximation ratio of the sequential variant of each rule, computed as the OWA score of the sequential outcome divided by the optimal OWA score. The setup is analogous to the previous experiment, and we again set $k = 3$. For this experiment, we aggregate the results over for all values of the dispersion

parameter, and only present the minimum and maximum approximation ratio. Results are in Table 5.1. Next, we compare scalability of the sequential and non-sequential rules. To do so, we select ABSAFs whose frameworks have an increasing number of preferred extensions, and study the effect of this parameter (number of extensions) on the runtime. We start by selecting 30 AFs with 8 to 12 preferred extensions each. In the next iteration, we select 30 AFs with 13 to 17 preferred extensions, and so on, increasing the minimum and maximum number of preferred extensions by 5 at every iteration. To sample the approvals, for each AF *Fram*, we choose uniformly at random the dispersion value to use from the set $\Phi \in \{0, 0.1, \dots, 0.9, 1\}$, and sample approximately 100 ballots centred around $\lfloor |prf(Fram)| \cdot 1/4 \rfloor$ ground truths (chosen uniformly at random from $prf(Fram)$).¹⁴ Note that, since we pre-compute the preferred extensions, this approach only compares the selection process, in line with the discussion after Proposition 5.11. We run the utilitarian and seq-utilitarian¹⁵ rules on each input, looking for outcomes of size $k = \lfloor |prf(Fram)| \cdot 1/4 \rfloor$ and setting a timeout of 45 seconds. We increase the number of preferred extensions until no instance out of the 30 is solved within the time limit. For the instances that terminate before timeout, we report the average runtime. The results are shown in Table 5.2.

As a final, informal note, we also ran the sequential rules on even larger ABSAFs, to assess their limits. We do not show a full-fledged experiment for this, but report some figures to build intuition. On a framework with 256 extensions, using the data-generation technique outlined in the previous paragraph, we could find an outcome of size $k = 64$ in 22 seconds. With 512 extensions and $k = 128$, we exceed the timeout of 45 seconds.

5.7.2 Discussion of the Results

In this section, we discuss the results of our simulation experiments.

5.7.2.1 Varying the Size of the Outcome

We begin by looking at the results of the varying k experiment (Figure 5.5). First, let us sketch what our expectations were with this experiment. Since there are 5 ground truths, our expectation (for low dispersion) was that, for $k = 5$, our method should be able to represent the voters well. By considering $k < 5$ we investigate how robust our approach is (i.e., whether the score “dips”) when selecting outcomes smaller than the number of ground truths. Conversely, by looking at $k \in \{6, 7\}$, we check whether we gain any performance by outcomes larger than the ground truth.

¹⁴Note that the exact number of voters depends on the value of $\lfloor |prf(Fram)| \cdot 1/4 \rfloor$, and in our setup, was never less than 96. Regardless, the number of voters has a limited impact on the runtime.

¹⁵Careful readers might object that these two rules compute the same output; we have already mentioned this fact elsewhere for rules defined by the weight function $\omega^{(n)} = (1, \dots, 1)$. For instance, think of AV in the ABC voting model. However, in this experiment we are not comparing the outcomes, but only the runtime of the sequential algorithm to the runtime of the non-sequential one. We choose the utilitarian rule for simplicity, as we expect the runtime to not depend substantially on $\omega^{(n)}$.




Figure 5.5: Results of the varying k experiment. Error bars indicate sample standard deviations; for readability, they are slightly horizontally offset and alternated across dispersion values.

Considering Figure 5.5, we see the following. First, as Φ grows, performance drops at a similar rate for all k . Secondly, our methods seems quite robust: if we select outcomes slightly smaller than the true number of ground truths (i.e., $k \in \{3, 4\}$), we do not observe substantial performance drops. Next, for $k = 1$, the performance is rather poor, especially for the egalitarian rule, with the large standard deviation suggesting unreliable behaviour. This is reasonable as a single viewpoint, in general, might not represent everyone. Finally, if we select outcomes larger than the number of ground truths ($k > 5$), we do not gain significant advantages. All in all, these observations indicate that $k = 3$ is a balanced and representative setting, which justifies focusing on this value for the subsequent experiments.




Figure 5.6: Average, minimal, perfect representation, and recovery score experiments for $k = 3$. Error bars indicate sample standard deviations; for readability, they are slightly horizontally offset and alternated across dispersion values.

5.7.2.2 Comparing the Rules

We now consider the second simulation experiment, where we compare our rules. In general, we expected here that, as Φ increases, the performance of the rules decreases, as the votes become more diverse and hence less easily representable in a short summary. Looking at Figure 5.6, we note that this is the case. Moreover, all rules outperform the baseline that returns k random extensions. This is promising, although this baseline achieves quite high performance in the average and perfect representation metrics. Next, although the oracle returning all preferred extensions is, as expected, always best performing, the performance of the oracle returning all ground truths decreases. The latter oracle is even outperformed by some rules in some metrics, although, of course, in the recovery score it achieves perfect performance, fully overlapping with the other oracle. This is reasonable: as Φ increases, votes become less similar to the ground truths.

Furthermore, we can see from the average and minimum representation metrics that

the harmonic rule seems to be a good compromise between the utilitarian and egalitarian rules. In particular, the harmonic rule performs almost as well as the egalitarian one in the minimum representation metric, which, given the poor performance of the latter rule in other metrics, suggests that the harmonic rule should be preferred to it. Furthermore, we can see that MAXCOV performs poorly for minimum representation, and its performance decreases starkly as Φ increases for average representation. Finally, looking at the recovery and perfect representation scores, the utilitarian rule and MAXCOV seem the best performing rules; however, the harmonic rule performs comparably. Still, the large standard deviations for the recovery score suggest substantial variability across instances. Hence, the results for this metric should be interpreted with caution.

In general, despite the axiomatic advantages of MAXCOV (Theorems 5.15 and 5.16), in our setup, the utilitarian rule seems to perform better with respect to most metrics, and comparably even under the angle of perfect representation. Depending on the application scenario, we can hence recommend the utilitarian or harmonic rules, the latter offering a compromise between total representation and fairness.

rule	min ratio	max ratio
utilitarian	99.26%	99.77%
harmonic	96.61%	99.28%
egalitarian	89.41%	97.73%
MAXCOV	98.75%	99.84%

Table 5.1: Approximation ratio experiment for $k = 3$.

# pref. ext.	non-sequential		sequential	
	successes	time	successes	time
$\{8, \dots, 12\}$	30/30	0.03	30/30	0.01
$\{13, \dots, 17\}$	30/30	0.73	30/30	0.02
$\{18, \dots, 22\}$	30/30	5.63	30/30	0.03
$\{23, \dots, 27\}$	1/30	17.24	30/30	0.06
$\{28, \dots, 32\}$	0/30	-	30/30	0.12

Table 5.2: Performance experiment.

5.7.2.3 Sequential Rules

Finally, we notice from table Table 5.1 that the sequential formulations offer a good approximation ratio, always well above 90%, often close to 100%. The egalitarian rule here is an exception, having a minimal approximation ratio of around 89%. Still, our previous discussion already suggests that this rule should not, in general, be adopted in practice. Moreover, from Table 5.2 we can see, as hinted by our theoretical findings

(Proposition 5.12), that sequential rules significantly improve over the runtime performance of the non-sequential ones, and seem to scale quite well. The results in these two tables suggest that sequential rules could be a scalable and well-performing method for applying our framework in real-world scenarios.

5.8 Conclusion

In this chapter, we presented a new framework for obtaining representative summaries of online discussions by combining techniques from multi-winner voting and abstract argumentation.

Summary. Let us present our main contributions.

- We began by noting the increasing interest in platforms for public deliberation, such as Polis. Our objective is to offer a tool for summarising these discussions in a representative way, i.e., one that reflects the various positions expressed in the discussion in proportion to the support they receive.
- In Section 5.3, we introduce our object of study, [Approval-Based Social Argumentation Frameworks \(ABSAFs\)](#). Conceptually, in this model, citizens can post their comments and approve (or not) of comments posted thus far. Technically, an [ABSAF](#) consists of an [Argumentation Framework \(AF\)](#) (capturing the discussion) together with a profile of approval preferences over its arguments. We use abstract argumentation theory to identify the potential viewpoints: those are the preferred extensions, which can be thought of as maximally-consistent sets of comments that defend themselves against counterarguments. We then use the preferences to filter out some of these extensions.
- In Section 5.4, we study how one can perfectly represent *all* positions in the discussion. We give conditions under which this is possible and study the complexity of the related decision problems. We find that, for smaller instances, we can efficiently find an outcome that represents the “defensible cores” of all users’ preferences.
- In Section 5.5, we take a different approach: having fixed the number of viewpoints we want to select, how can we best represent the users? We introduce several methods for doing so, mostly based on the [OWA](#) operators and their sequential approximation.
- In Section 5.6 we introduce axioms that capture guarantees of proportional representation in our setting. These are based on the justified representation axioms from the [ABC](#) voting literature. We show that [OWA](#) rules fail our [Deliberative Justified Representation \(DJR\)](#) axiom, but another rule we introduce, [MAXCOV](#), satisfies it. The same holds for its sequential approximation.

- Finally, in Section 5.7, we perform simulations to assess the performance of our rules. The harmonic and utilitarian rules seem to be the best-performing rules overall. We find that the sequential approximations of these rules still retain good performance and scale to larger instances.

We can give specific suggestions on which rules to employ. Axiomatically, MAXCOV showed the best behaviour and can be recommended when satisfying DJR is required. If this can be negotiated, our simulations suggest that the utilitarian and harmonic rules are most appealing, with the former being more efficient in terms of total representation and the latter more representative.

Future work. This work can be extended along several directions.

- First, we assumed that discussions were already formalised as AFs. In practice, when using real-world data, this formalisation is a crucial step and is the main obstacle to the applicability of our method. In Example 5.2, we manually created the AF representing the Canadian election reform discussion. However, while manual annotation is feasible for small discussions, fair moderation is important to avoid any subjective biases. The formalisation of larger debates, instead, may require argument mining techniques that can convert (often previously annotated) natural language text into AFs. This remains an open challenge, despite decades of research on the topic [Palau and Moens, 2009, Cabrio and Villata, 2012, Baumann et al., 2020, Goffredo et al., 2023, Li et al., 2025]. In particular, the recent revolution in natural language processing, with large language models, offers a promising approach to extract abstract argumentation from text. See, for instance, the work of Gorur et al. [2025]. Given this, we envision that our framework might find real-world applicability in the future.
- ABSAFs, as a model, can be generalised and improved. As an example, Rossie et al. [2024] present a generalisation of our formalism in which users can approve, disapprove, or remain neutral towards arguments. This is a popular format in social media, and hence a useful generalisation (indeed, we listed this as future work in our original paper [Bernreiter et al., 2024a]). Moreover, our general approach is independent of the choice of the underlying formalism for identifying consistent sets of comments. We have chosen a specific formalism for concreteness, and we have chosen abstract argumentation in particular, as it is attractive for our purposes. Indeed, as already mentioned, it requires only minimal semantic annotation, and at the same time is well-understood. However, future work should consider richer formalisms as well, ranging from generalisations of AFs (like the aforementioned bipolar variant [Cayrol and Lagasquie-Schiex, 2005] or claim-augmented AFs [Dvořák and Woltran, 2020]), to competing models of argument like structured argumentation, where arguments have a complex inner structure [Besnard et al., 2014].

- A distinct limitation concerns the status of the attack relation. In this chapter, we assumed that the attack relation is unequivocal and accepted by all participants. This assumption might be reasonable in settings where a team of moderators or an automated approach is capable of (and has the authority to) annotating the discussion with attack relations. However, if we want to crowdsource these relations from participants, different users might have different opinions about what constitutes an attack. Future work should investigate how to capture such diverging opinions about argument attacks within the model. One approach would be to replace AFs with related models that natively support voting over attack relations, such as weighted AFs [Dunne et al., 2011, Bistarelli and Santini, 2021]. Our approach could then be applied with minimal modification to select a representative subset of the extensions of such weighted frameworks. Ideally, however, one might want to account for preferences over both the arguments and the attack relations *simultaneously* when selecting the outcomes, which would require adapting our approach significantly. It is also worth noting that approaches in which participants vote on both arguments and relations already exist in the literature, such as the relational reasoning model [Ganzer et al., 2023]. However, the latter focuses on deciding whether a particular proposition is collectively accepted and does not naturally support the production of representative summaries, which is our goal.
- Finally, Fichte et al. [2026] recently studied a notion of *diversity* for extensions of AFs. It would be interesting to compare and possibly combine our approaches to obtain a set of viewpoints that is both diverse and representative.

Epilogue

This thesis dealt with multi-winner elections with approval ballots. We studied a range of different settings—electing parliaments and committees, making several decisions sequentially or simultaneously, and selecting a set of viewpoints from a larger discussion. We addressed fundamental questions about established formalisms, such as [Approval-Based Committee \(ABC\)](#) voting, and proposed novel models with concrete applications.

In Chapter 3 (“Representation”), we introduced two axioms for the [ABC](#) voting model. The first, *justified upper quota*, aims to prevent the overrepresentation of voter groups; the second, *justified near quota*, aims to balance underrepresentation and overrepresentation. Using these axioms, we characterised two rules in the class of composite Thiele rules (a generalisation of the well-known class proposed by [Thiele \[1895\]](#)). These are, respectively, Adams and Sainte-Laguë approval voting, each extending the corresponding apportionment method. We further studied the relationship between our axioms and established notions of proportionality. In Chapter 4 (“Manipulation”), we considered a setting in which several decisions must be made and studied a specific form of strategic manipulation: *free-riding*. We showed that this phenomenon is inevitable if we wish to guarantee representation in such elections, especially when decisions are made over time. On the positive side, free-riding might not be as risk-free as one would initially expect. We argued this by showing that free-riding can reduce the final satisfaction of the manipulating voter and by assessing its associated risk through simulations. Finally, in Chapter 5 (“Deliberation”), we applied the machinery of [ABC](#) voting to the problem of obtaining “representative summaries” of online discussions. In particular, we developed a model consisting of a set of interrelated comments (the debate) together with the approvals of the users over these comments. We used the theory of abstract argumentation [[Dung, 1995](#)] to identify sets of comments that are jointly acceptable, and ideas from multi-winner voting to select acceptable sets of comments that represent the users. We proposed several methods and studied them axiomatically, computationally, and via simulations.

Each of Chapters 3 to 5 has a concluding section detailing research avenues and open questions for that specific setting; we will not repeat them here. Looking at approval-based multi-winner voting more broadly, beyond the specific technical issues that remain open [Lackner and Skowron, 2023, Chapter 7], an important direction of work is to bring this rich theory to fruition.

It has been observed by Brams and Fishburn [2010] that, despite the substantial attention that approval voting has attracted in research, it has seen limited political application, even if it has been implemented in the internal elections of several scientific institutions. While Brams and Fishburn refer mostly to the single-winner case, similar conclusions seem plausible for the multi-winner case, as large-scale adoption of approval voting appears to be rare. While advocating for specific electoral processes goes beyond research itself, touching on activism and outreach (Brams and Fishburn even go so far as to explain why they think academics are not the best “sales people” for this job), there are potential research directions to explore. Approval voting is a relatively novel format, of which relatively few people have experience; hence, it may encounter resistance. This raises a broad question: what analysis would be needed in order to move these methods from theory to practice? What evidence could help policymakers and laypeople assess the applicability of approval voting? For example, daring to think about a high-stakes context, what would be the effects of adopting approval voting in parliamentary elections, e.g., by allowing voters to approve multiple parties, a setting studied in social choice theory under the name of *party-approval elections* [Brill et al., 2018, Delemazure et al., 2023]? Beyond the usual social choice questions, how would such a change affect party positioning and campaigning? Another question would be: in which contexts are the elections studied in Chapter 4 applicable? That is, are there political settings in which proportional representation *over time* is desirable, and if so, what kinds of institutions would they suit?

A related but orthogonal question would be to identify scenarios in which the mathematical models developed in this context—most notably ABC voting—would be useful *beyond* the prototypical electoral case. One such example, deliberation platforms, is presented in Chapter 5, and the literature already contains others, such as the use of proportional methods for clustering data points [Kellerhals and Peters, 2024] and the use of ABC voting rules in blockchain technologies [Boehmer et al., 2024a]. What other settings might benefit from the theory of multi-winner elections despite not *prima facie* fitting that descriptor? Answering this question could help expand the practical reach of computational social choice, in line with recent calls to apply its methods to real-world problems [Baumeister et al., 2026].

This brings us to a close. It has been a rather technical dissertation, chiefly aimed at the theoretical foundations of collective decision-making, and the connection to democratic practice might not always have been apparent. Hopefully, we at least illustrated the intricacies hidden within a seemingly innocuous aspect of public life, voting; reminding us that democracy is not a motionless ideal but a system of concrete and amendable institutions.

Additional Results on Representation

In this appendix, we present additional results from Chapter 3.

A.1 Justified Upper Quota Swaps and Might Break EJR+

In Chapter 3, we have shown that JUQ swaps might break EJR+ (Proposition 3.23). Here, we show that this holds even when starting from a particularly well chosen EJR+ committee, such as the output of the Method of Equal Shares (MES) [Peters and Skowron, 2020]. We first define the rule.

Definition A.1. MES works sequentially in rounds. We begin with the empty committee, $W_0 = \emptyset$, and at each round ℓ we add a candidate to $W_{\ell-1}$ to obtain W_ℓ . Initially, all voters are assigned a budget of k/n . Let $b_\ell(i)$ be the budget of voter i before round ℓ (hence, $b_1(i) = k/n$ for all $i \in N$). In round ℓ , we consider the set $C_\ell \subseteq C$ of all candidates in $C \setminus W_\ell$ such that $\sum_{i \in N(c)} b_\ell(i) \geq 1$. If C_ℓ is empty, the rule terminates and returns W_ℓ . Otherwise, we compute $\alpha(c)$ for each $c \in C_\ell$ as

$$\alpha(c) = \min \left\{ \alpha \in \mathbb{R} : \sum_{i \in N(c)} \min(\alpha, b_\ell(i)) = 1 \right\}$$

and, according to some tiebreaking, select some candidate c^* with

$$c^* \in \arg \min_{c \in C_\ell} \alpha(c).$$

We then set $b_{\ell+1}(i) = \max(0, b_\ell(i) - \alpha(c^*))$ if $i \in N(c^*)$ and $b_{\ell+1}(i) = b_\ell(i)$ otherwise. Finally, we set $W_{\ell+1} = W_\ell \cup \{c^*\}$.

MES returns all committees that can be obtained by the above procedure plus some way of breaking ties.

Note that **MES** might terminate in less than k rounds, and thus return a committee with fewer than k candidates. To enforce output committees to have size exactly k , one needs to extend **MES** with a so-called “completion method” [Lackner and Skowron, 2023, Chapter 2]. We remark that Brill and Peters [2023] proved that **MES** satisfies **EJR+**.

Proposition A.1. A sequence of **JUQ** swaps inducing a sequence of committees $\{W_1, W_2, \dots\}$ that starts from a committee satisfying **EJR+** might terminate with a committee that violates **EJR+** (while satisfying **JUQ**).

This holds even if the initial committee W_1 was selected by **MES** (with any completion method) and even if swap decisions are based solely on the satisfaction vectors of the supporters of the candidates being considered for replacement. That is, when deciding which swap to perform on committee W_ℓ , we consider two swaps (c_1, d_1) and (c_2, d_2) to be indistinguishable whenever

$$(|A(i) \cap W_\ell|)_{i \in N(c_1)} = (|A(i) \cap W_\ell|)_{i \in N(c_2)}$$

holds, where the voters in each vector are ordered according to their natural ordering in $N = \{1, \dots, n\}$.

Proof. We structure our counterexample in three parts. First, we construct an **ABC** election, then derive the outcome of **MES** (with arbitrary completion), and finally, consider sequences of **JUQ** swaps starting from that outcome.

Constructing the instance. We construct a counterexample instance as follows. For notational convenience, let $s = 31$. Moreover, let $k = 140$. Set the number of voters to be $n = 6sk$, which notably means that the quota of a group of $6s$ voters will be exactly 1.

To assist with our argument, consider two **ABC** voting sub-instances. First, what we call the “single gadget”, an election with $n = 36s$ voters and with $k = 6$:

Candidate/Approvals	6s	6s	6s	6s	6s	6s
c_1, c_2, c_3	✓	✓	✓			
c_4	✓			✓		
c_5		✓			✓	
c_6			✓			✓

Here, we have that every submitted ballot is submitted by exactly $6s$ voters. For example, $6s$ voters submit $\{c_1, \dots, c_4\}$. Second, what we call the “double gadget”, with $n = 72s$ voters and $k = 12$, where each group of voters $R_i \subseteq N$ and $S_i \subseteq N$ contains exactly $3s$ voters:

Cand./Group	R_1	R_2	R_3	R_4	R_5	R_6	R_7	R_8	R_9	R_{10}	R_{11}	R_{12}	S_1	S_2	S_3	S_4	S_5	S_6	S_7	S_8	S_9	S_{10}	S_{11}	S_{12}
c'_1, c'_2, c'_3	✓	✓	✓	✓	✓	✓																		
c'_4, c'_5, c'_6							✓	✓	✓	✓	✓	✓												
c'_7	✓						✓						✓						✓					
c'_8		✓						✓						✓						✓				
c'_9			✓						✓						✓						✓			
c'_{10}				✓						✓						✓						✓		
c'_{11}					✓						✓						✓						✓	
c'_{12}						✓						✓						✓						✓

Observe that in the single gadget **MES** selects all 6 candidates, which is an exhaustive outcome. There, candidates c_1, c_2 and c_3 all violate **UQ**. Similarly, **MES** selects all candidates in the double gadget, and $\{c'_1, \dots, c'_6\}$ all witness **UQ** violations.

To build our final **ABC** election, we add 15 copies of the single gadget to our instance (giving $15 \cdot 6 \cdot 6s$ voters), as well as one copy of the double gadget. We assume that all gadgets are disjoint, i.e., a voter from any gadget approves only candidates from that gadget. We will refer to candidates $\{c'_1, \dots, c'_6\}$ of the double gadget as candidates of “type A” and $\{c'_7, \dots, c'_{18}\}$ of “type B”. We additionally add the following voter sets to the instance:

- sets $\{T_1, \dots, T_{12}\}$, each containing $9s + 1$ voters;
- sets $\{U_1, \dots, U_{12}\}$, each containing $9s + 1$ voters;
- The set Z , containing $12s - 24$ additional voters to ensure that $n = 6sk$.

This gives us a total of $140 \cdot 6s = 26040$ voters.

We subdivide the $216s + 24$ voters from $\{T_1, \dots, T_{12}\}$ and $\{U_1, \dots, U_{12}\}$ into 35 arbitrary disjoint groups of equal size (each containing 192 voters), and call these groups $\{Y_1, \dots, Y_{35}\}$. Note that the choice of $s = 31$ ensures that $216s + 24$ was divisible by 35.

We then add the following candidates to the instance, including them in the approval sets of some voters from the double gadget, and some of the additional voters above:

- “type C” candidates:
 - candidates $\hat{c}_i$ with $i \in [12]$, approved by voters in $S_i \cup T_i$;
 - candidates $\hat{c}_{i+12}$ with $i \in [12]$, approved by voters in $S_i \cup U_i$;
- “type D” candidates:
 - candidates $\bar{c}_i$ for $i \in [34]$, approved by voters in Y_i and Y_{i+1} ;
 - candidate $\bar{c}_{35}$, approved by voters in Y_{35} and Y_1 ;
- candidate c_Z , approved by the voters in Z .

We will refer to the voters and candidates belonging to the second gadget, together with these additional voters and candidates, as the “enhanced double gadget”.

This concludes the construction.

Computing the MES outcome. Let n_A, n_B, n_C , and n_D be the number of supporters for candidate types A, B, C and D , respectively. Note that, within each type, this number is the same, so there is no ambiguity here. Similarly, let q_A, q_B, q_C , and q_D be the quotas of each candidate type. Observe the following:

- type A candidates are approved by $n_A = 18s$ voters and thus have $q_A = 3$;
- type B candidates are approved by $n_B = 12s$ voters and thus have $q_B = 2$;
- type C candidates are each approved by $12s + 1$ voters, which means $2 < q_C < 3$;
- type D candidates are each approved by $18s > 384 > 12s + 1$ voters, and thus $2 < q_D < 3$ and $q_D > q_C$.

While executing **MES**, we can treat each single gadget separately, as they are disjoint, and also disjoint from the (enhanced) double gadget. We can also treat the voters in Z separately. **MES** selects every candidate from every single gadget ($15 \cdot 6 = 90$ in total), as well as c_Z . From the enhanced double gadget, **MES** selects, in order:

- all 6 type A candidates $\{c'_1, \dots, c'_6\}$, exhausting the budgets of voters in groups $\{R_1, \dots, R_{12}\}$;
- all 35 type D candidates $\{\bar{c}_1, \dots, \bar{c}_{35}\}$. At this point, the voters in groups $\{T_1, \dots, T_{12}\}$ and $\{U_1, \dots, U_{12}\}$ have purchased 35 candidates, from their combined **MES** budgets of $(216s + 24)/6s \approx 36.129$. Thus, by symmetry, each of the 24 voter groups has less than 0.048 budget remaining among its $9s + 1$ members. Meanwhile, each voter group in $\{S_1, \dots, S_{12}\}$ has an **MES** budget of 0.5 among its $3s$ members. We can conclude that no type C candidate is affordable by its supporters.
- all 6 type B candidates $\{c'_7, \dots, c'_{12}\}$.

Thus, **MES** selects a total of 138 candidates. To complete this committee to size $k = 140$, we choose any two type C candidates, and call the resulting set W .

Checking the JUQ swaps. First, observe that no type B candidate currently constitutes a **UQ** violation for W , even if any one additional candidate is added to it. For instance, a committee in which c'_7 violates **UQ** must also contain $\hat{c}_1, \hat{c}_7, \hat{c}_{13}$ and $\hat{c}_{19}$, as well as two of $\{c'_1, c'_2, c'_3\}$, and two of $\{c'_4, c'_5, c'_6\}$ in order to give all voters in groups $\{R_1, R_7, S_1, S_7\}$ a satisfaction of 3. In other words, it is necessary (but not sufficient) for at least 4 type C candidates to be in a committee in order to involve any type B candidate in a **UQ** violation.

In W , the only candidates that witness a violation of **UQ** are $\{c_1, c_2, c_3\}$ in each of the 15 single gadgets, as well as the type A candidates from $\{c'_1, c'_2, c'_3\}$ and $\{c'_4, c'_5, c'_6\}$ in the double gadget. Observe that each of these candidates has exactly the same amount of supporters, namely 18s, with each of those supporters obtaining a satisfaction of 4 from committee W . Removing one candidate from any set of three candidates listed above leads to the other two no longer violating **UQ**.

As per the statement we are trying to prove, we assume that our method of selecting **JUQ** swaps uses the satisfaction of the supporters of a candidate to select which swap to execute. Thus, it cannot distinguish between these violations, so we will assume that the first candidate removed in a **JUQ** swap is c'_1 . Each of the 22 remaining type C candidates not selected by our completion method can be added to the outcome without violating **UQ**. Observe that adding the third type C candidate cannot bring any additional candidate over quota: type C candidates can never constitute a **UQ** violation and a type B candidate can only constitute a **UQ** violation once at least four type C candidates have been added. Thus, we will assume that the second candidate removed in a **JUQ** swap is c'_4 .

As we continue to perform **JUQ** swaps, we will remove 15 **UQ**-violating candidates of the single gadgets and 5 of the 6 type B candidates, in some order, while adding in the remaining 20 type C candidates. Notably, only the last of these **JUQ** swaps will make the remaining type B candidate violate **UQ**. Assume this is c'_{12} , without loss of generality. In addition to c'_{12} , the committee currently contains 5 candidates from each of the 15 single gadgets, c_Z , 4 type A candidates $\{c'_2, c'_3, c'_5, c'_6\}$, all 24 type C candidates, and all 25 type D candidates. Observe that c'_{12} witnesses a **JUQ** violation with either c'_1 or c'_4 , the type A candidates we initially removed. Notably, we did not want to add a type A candidate back into the committee as part of a **JUQ** swap until 5 of the 6 type B candidates were removed. Performing either **JUQ** swap leads to a committee for which the other type A candidate witnesses an **EJR+** violation. Furthermore, no candidate in that committee witnesses a violation of **UQ**, and thus the sequence of **JUQ** swaps terminates in a committee violating **EJR+**. $\square$

A.2 Counterexamples for Other Voting Rules

In this section, we present several counterexamples for familiar **ABC** voting rules for **JUQ** and **JNQ**. We tested all the rules present in the **ABCVOTING** Python package [Lackner et al., 2023b] at the time of writing, as well as some additional rules of interest. For the sake of simplicity, for the rules implemented in **ABCVOTING**, we refer to them by their name in the package.¹ For the other rules, we provide references.

A.2.1 Counterexamples for Justified Upper Quota

We now show counterexamples for **JUQ**.

¹See this link for more information: <https://abcvoting.readthedocs.io/en/latest/intro-abcrules.html> (accessed: 2026-02-05.)

Example A.1. Let $m = 6$, $n = 10$, and $k = 2$. Consider profile

$$\begin{aligned} A(1) &= \{a\}, & A(6) &= \{b, c\}, \\ A(2) &= \{b, c, d\}, & A(7) &= \{f\}, \\ A(3) &= \{b, c\}, & A(8) &= \{d\}, \\ A(4) &= \{b\}, & A(9) &= \{b, c\}, \\ A(5) &= \{e\}, & A(10) &= \{b\}. \end{aligned}$$

In this example, $\{b, c\}$ does not satisfy [JUQ](#), as we can swap c with a . The rules `av`, `sav`, `pav`, `slav` (our [SL](#)), `geom2`, `seqpav`, `revseqpav`, `seqslav`, `seqphragmen`, `minimaxphragmen`, `maximin-support`, `minimaxav`, `equal-shares` (in all its variants), and `eph` select $\{b, c\}$ in this instance.

Example A.2. Let $m = 4$, $n = 4$, and $k = 3$. Consider profile

$$\begin{aligned} A(1) &= \{b, c, d\}, \\ A(2) &= \{a, d\}, \\ A(3) &= \{b, c, d\}, \\ A(4) &= \{b, c, d\}. \end{aligned}$$

In this example, $\{a, b, d\}$ does not satisfy [JUQ](#), as we can swap a with c . The rules `cc`, `seqcc`, `lexcc`, `monroe`, `consensus-rule`, `phragmen-enestroem`, and `leximinmaxav` can select $\{a, b, d\}$ in this instance.

Example A.3. Let $m = 4$, $n = 4$, and $k = 2$. Consider profile

$$\begin{aligned} A(1) &= \{a, b, c\}, \\ A(2) &= \{a, d\}, \\ A(3) &= \{a, c, d\}, \\ A(4) &= \{a, d\}. \end{aligned}$$

In this example, $\{a, c\}$ does not satisfy [JUQ](#), as we can swap c with d . The `leximaxphragmen` and `rsd` rules can select $\{a, c\}$ in this instance.

Example A.4. Let $m = 5$, $n = 4$, and $k = 2$. Consider profile

$$\begin{aligned} A(1) &= \{e\}, \\ A(2) &= \{a, c, d\}, \\ A(3) &= \{c\}, \\ A(4) &= \{a, b, c\}. \end{aligned}$$

In this example, $\{a, c\}$ does not satisfy [JUQ](#), as we can swap a with e . The rule `greedymonroe` can select $\{a, c\}$ in this instance.

Example A.5. Let $m = 5$, $n = 6$, and $k = 3$. Consider profile

$$\begin{aligned} A(1) &= \{a\}, \\ A(2) &= \{e\}, \\ A(3) &= \{c\}, \\ A(4) &= \{b, e\}, \\ A(5) &= \{b, c, d\}, \\ A(6) &= \{d\}. \end{aligned}$$

In this example, $\{b, c, e\}$ does not satisfy [JUQ](#), as we can swap b with a . The sequential formulation of Adams (defined analogously as sequential simple Thiele rules, see the definition by [Lackner and Skowron \[2023, Section 2.3\]](#)) can select $\{b, c, e\}$ in this instance.

Example A.6. Let $m = 9$, $n = 15$, and $k = 8$. Consider profile

$$\begin{array}{lll} A(1) = \{a, b, c\}, & A(6) = \{c, i\}, & A(11) = \{e\}, \\ A(2) = \{a, d, e\}, & A(7) = \{c, i\}, & A(12) = \{f\}, \\ A(3) = \{a, f, g\}, & A(8) = \{d\}, & A(13) = \{f\}, \\ A(4) = \{b, h\}, & A(9) = \{d\}, & A(14) = \{g\}, \\ A(5) = \{b, h\}, & A(10) = \{e\}, & A(15) = \{g\}. \end{array}$$

In this example, $\{a, b, c, d, e, f, g\}$ does not satisfy [JUQ](#), as we can swap a with h . [GJCR](#) (Definition 3.11) can select $\{a, b, c, d, e, f, g\}$ in this instance. Note that the violation persists even if we consider the exhaustive committee $\{a, b, c, d, e, f, g, h\}$.

Example A.7. Let $m = 3$, $n = 4$, and $k = 2$. Consider profile

$$\begin{aligned} A(1) &= \{a\}, \\ A(2) &= \{a, b\}, \\ A(3) &= \{a, b, c\}, \\ A(4) &= \{a, b, c\}. \end{aligned}$$

In this example, $\{a, c\}$ does not satisfy **JUQ**, as we can swap c with b . The Greedy Cohesive Rule introduced by Peters et al. [2021] (in the formulation given by Brill et al. [2023]) can select $\{a, c\}$ in this instance.

A.2.2 Counterexamples for Justified Near Quota

We now show counterexamples for **JNQ**.

Example A.8. Let $m = 7$, $n = 5$, and $k = 3$. Consider profile

$$\begin{aligned} A(1) &= \{b, e, f\}, \\ A(2) &= \{f\}, \\ A(3) &= \{a\}, \\ A(4) &= \{a, b, c, g\}, \\ A(5) &= \{g\}. \end{aligned}$$

In this example, $\{a, b, f\}$ does not satisfy **JNQ**. We can swap b with g . The quota of $N(b)$ is $6/5$, while the quota of voter 5 is $3/5$. The rules `av`, `sav`, `cc`, `lexcc`, `geom2`, `seqpav`, `revseqpav`, `seqcc`, `minimaxphragmen`, `maximin-support`, `monroe`, `greedy-monroe`, `minimaxav`, `leximinmaxav`, `equal-shares` (in all its variants), `phragmen-enestrom`, can select $\{a, b, f\}$ in this instance.

Example A.9. Let $m = 5$, $n = 5$, and $k = 3$. Consider profile

$$\begin{aligned} A(1) &= \{b, d\}, \\ A(2) &= \{a, d\}, \\ A(3) &= \{a, b\}, \\ A(4) &= \{c\}, \\ A(5) &= \{e\}. \end{aligned}$$

In this example, $\{a, b, d\}$ does not satisfy **JNQ**, as we can swap b with e . The rules `pav`, `seqslav`, `seqphragmen`, `rsd` and `consensus-rule` can select $\{a, b, d\}$ in this instance.

Example A.10. Let $m = 7$, $n = 7$, and $k = 3$. Consider profile

$$\begin{aligned} A(1) &= \{a, c\}, \\ A(2) &= \{a, b, e\}, \\ A(3) &= \{b, c, e\}, \\ A(4) &= \{c, e, f, g\}, \\ A(5) &= \{d\}, \\ A(6) &= \{b, c, e, f, g\}, \\ A(7) &= \{a\}. \end{aligned}$$

In this example, $\{a, b, c\}$ does not satisfy [JNQ](#), as we can swap c with e . [Adams-AV](#) (Definition 3.7) and eph can select $\{a, b, c\}$ in this instance.

Example A.11. Let $m = 5$, $n = 5$, and $k = 2$. Consider profile

$$\begin{aligned} A(1) &= \{a, b, d\}, \\ A(2) &= \{c, d\}, \\ A(3) &= \{b, d\}, \\ A(4) &= \{e\}, \\ A(5) &= \{a, c, d\}. \end{aligned}$$

In this example, $\{b, c\}$ does not satisfy [JNQ](#), as we can swap b with d . The rule [leximaxphragmen](#) can select $\{b, c\}$ in this instance.

Additional Results on Deliberation

In this appendix, we present additional results from Chapter 3. For all plots shown in this appendix, error bars indicate sample standard deviations; for readability, they are slightly horizontally offset and alternated across dispersion values.

B.1 Approximation Ratio

In Tables B.1 to B.3, we present the additional results on the approximation ratio of sequential rules.

Regular representation			Core-representation		
rule	min ratio	max ratio	rule	min ratio	max ratio
utilitarian	99.20%	99.78%	utilitarian	99.20%	99.80%
harmonic	95.31%	98.72%	harmonic	95.31%	98.51%
egalitarian	88.42%	96.06%	egalitarian	88.42%	96.12%
MAXCOV	99.09%	99.93%	MAXCOV	99.09%	99.50%

Table B.1: Sequential approximation ratio experiment ($k = 2$).

rule	min ratio	max ratio
utilitarian	99.26%	99.79%
harmonic	96.61%	98.73%
egalitarian	89.41%	95.10%
MAXCOV	98.73%	99.31%

Table B.2: Sequential approximation ratio experiment ($k = 3$, core-representation).

B. ADDITIONAL RESULTS ON DELIBERATION

Regular representation			Core-representation		
rule	min ratio	max ratio	rule	min ratio	max ratio
utilitarian	99.74%	99.90%	utilitarian	99.74%	99.89%
harmonic	94.20%	99.79%	harmonic	94.20%	98.99%
egalitarian	73.96%	97.89%	egalitarian	73.96%	92.26%
MAXCOV	99.43%	99.97%	MAXCOV	99.33%	99.63%

Table B.3: Sequential approximation ratio experiment ($k = 5$).

B.2 Varying the Number of Selected Extensions

In Figure B.1, we present the simulations where we vary k , the number of selected viewpoints, for core-representation.




Figure B.1: Results of the experiments with variable k for core-representation.

B.3 Representation Metrics

In Figures B.2 to B.6, we present the additional results on average, minimal, perfect representation and recovery score.






Figure B.2: Average, minimal, perfect representation, and recovery score experiments for $k = 2$ (regular representation).


B. ADDITIONAL RESULTS ON DELIBERATION




(a) Average representation.



(b) Minimal representation.



(c) Prefect representation.



(d) Recovery score.

Figure B.3: Average, minimal, perfect representation, and recovery score experiments for $k = 2$ (core-representation).




Figure B.4: Average, minimal, perfect representation, and recovery score experiments for $k = 3$ (core-representation).

B. ADDITIONAL RESULTS ON DELIBERATION







Figure B.5: Average, minimal, perfect representation, and recovery score experiments for $k = 5$ (regular representation).




(a) Average representation.



(b) Minimal representation.



(c) Perfect representation.



(d) Recovery score.

Figure B.6: Average, minimal, perfect representation, and recovery score experiments for $k = 5$ (core-representation).

Overview of Generative AI Tools Used

No results presented in this thesis, including proofs, theorems, examples, or simulation code, were generated using AI tools. Likewise, no substantial portions of the written content were produced by AI; all writing, conceptual, and structural work was carried out either by the author or derived from pre-existing publications authored together with other human researchers, which are appropriately credited in the thesis. AI tools were used solely for limited editorial support, including grammar correction, wording suggestions, feedback on clarity, and assistance with certain $\text{\LaTeX}$ issues. German text included in this thesis was initially machine-translated and subsequently carefully reviewed by native speakers.

Acronyms

- ABC** Approval-Based Committee. 3–6, 9, 18–20, 23, 25, 26, 28, 33–38, 40, 46, 47, 49, 55, 58, 63, 67–69, 72, 75–77, 79–85, 87, 88, 134, 151–153, 155, 158, 162, 168–170, 173, 175–177, 184, 188, 191, 192, 194, 195, 197
- ABSAF** Approval-Based Social Argumentation Framework. 151, 154, 158, 160, 162–165, 168–170, 172–177, 179–184, 188, 189
- Adams-AV** Adams Approval Voting. 37, 44–46, 48–51, 55, 58, 59, 64, 67, 68, 75, 76, 201
- AF** Argumentation Framework. 154–160, 164, 167, 168, 181, 182, 184, 188–190
- ALQ** Apportionment Lower Quota. 40, 41, 49, 67, 68
- ANQ** Apportionment Near Quota. 40, 41, 62–66
- AUQ** Apportionment Upper Quota. 40, 41, 49, 51–54, 67, 68, 76
- AV** Approval Voting. 26–29, 31, 32, 45, 85, 184
- CC** Chamberlin-Courant Approval Voting. 27, 29–32, 44, 45, 51, 59, 67, 173, 174
- CNF** Conjunctive Normal Form. 21, 102, 105, 109, 112, 166
- DJR** Deliberative Justified Representation. 177–180, 188, 189
- EJR** Extended Justified Representation. 18, 19, 31–33, 35–38, 40, 41, 43, 46, 47, 49, 59, 60, 62, 66–73, 75, 76, 193, 194, 197
- FPT** Fixed-Parameter Tractability. 22, 168, 172, 173
- GJCR** Greedy Justified Candidate Rule. 60, 69–71, 199
- JNQ** Justified Near Quota. 5, 37, 63–67, 69, 70, 75, 76, 197, 200, 201
- JR** Justified Representation. 18, 19, 31, 67, 153, 175–177
- JUQ** Justified Upper Quota. 5, 37, 47–50, 55–64, 66–70, 72, 74–76, 193, 194, 196–200

MES Method of Equal Shares. 193–196

OWA Ordered Weighted Averaging. 5, 6, 27–29, 80, 84–87, 90, 95, 96, 98–102, 108, 112, 117, 119–123, 126, 128–130, 133, 134, 137–140, 142–145, 147, 152, 169–173, 178–181, 183, 185, 188, 204

PAV Proportional Approval Voting. 18, 27, 28, 31–33, 35–37, 39, 40, 43–45, 59, 67, 68, 88, 102, 120, 122, 124, 134, 170

PER Proportional Elimination Rule. 69–71

SDJR Strong Deliberative Justified Representation. 176, 177

SL Sainte-Laguë Approval Voting. 37, 43, 45, 64–68, 75, 198

STV Single Transferable Vote. 15, 16, 79

UQ Upper Quota. 46–49, 60–62, 67–69, 71, 72, 74, 195–197

UQER Upper Quota Elimination Rule. 49, 60–62, 69–71, 75

Bibliography

- ACE Electoral Knowledge Network. Comparative Data: Electoral Systems (ES005), 2026. URL <https://aceproject.org/epic-en/CDTable?view=country&question=ES005>. Accessed: 2026-04-22.
- J. Adams. *Thoughts on Government: Applicable to the Present State of the American Colonies*. John Dunlap, 1776.
- E. Adar and B. A. Huberman. Free Riding on Gnutella. *First Monday*, 5(10), 2000.
- D. S. Ahn and S. Oliveros. Combinatorial Voting. *Econometrica: Journal of the Econometric Society*, 80(1):89–141, 2012.
- F. Aleskerov. The History of Social Choice in Russia and the Soviet Union. *Social Choice and Welfare*, 25(2):419–431, 2005.
- G. Alfano, S. Greco, F. Parisi, and I. Trubitsyna. On Preferences and Priority Rules in Abstract Argumentation. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI-2022)*, pages 2517–2524, 2022.
- G. Alfano, S. Greco, F. Parisi, and I. Trubitsyna. Abstract Argumentation Framework with Conditional Preferences. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI-2023)*, pages 6218–6227. AAAI Press, 2023.
- P. Alimonti and V. Kann. Some APX-Completeness Results for Cubic Graphs. *Theoretical Computer Science*, 237(1):123–134, 2000.
- G. Amanatidis, N. Barrot, J. Lang, E. Markakis, and B. Ries. Multiple Referenda and Multiwinner Elections Using Hamming Distances: Complexity and Manipulability. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2015)*, pages 715–723. IFAAMAS, 2015.
- S. Arora and B. Barak. *Computational Complexity: A Modern Approach*. Cambridge University Press, 2009.
- K. J. Arrow. *Social Choice and Individual Values*. Wiley, 1951.
- K. J. Arrow, A. Sen, and K. Suzumura, editors. *Handbook of Social Choice and Welfare*. Elsevier, 2002.

- K. Atkinson and T. J. M. Bench-Capon. Value-Based Argumentation. In D. Gabbay, M. Giacomin, G. Simari, and M. Thimm, editors, *Handbook of Formal Argumentation*, volume 2, pages 299–354. College Publications, 2021.
- E. Awad, R. Booth, F. Tohmé, and I. Rahwan. Judgement Aggregation in Multi-Agent Argumentation. *Journal of Logic and Computation*, 27(1):227–259, 2017.
- H. Aziz, S. Gaspers, J. Gudmundsson, S. Mackenzie, N. Mattei, and T. Walsh. Computational Aspects of Multi-Winner Approval Voting. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2015)*, pages 107–115. IFAAMAS, 2015.
- H. Aziz, M. Brill, V. Conitzer, E. Elkind, R. Freeman, and T. Walsh. Justified Representation in Approval-Based Committee Voting. *Social Choice and Welfare*, 48(2):461–485, 2017.
- H. Aziz, A. Bogomolnaia, and H. Moulin. Fair Mixing: The Case of Dichotomous Preferences. In *Proceedings of the 2019 ACM Conference on Economics and Computation (EC-2019)*, pages 753–781. ACM, 2019.
- M. Balinski and H. P. Young. *Fair Representation: Meeting the Ideal of One Man, One Vote*. Brookings Institution Press, 2nd edition, 2001. Originally published in 1982.
- A.-L. Barabási and R. Albert. Emergence of Scaling in Random Networks. *Science*, 286(5439):509–512, 1999.
- S. Barberà, W. Bossert, and P. K. Pattanaik. Ranking Sets of Objects. In S. Barberà, P. Hammond, and C. Seidl, editors, *Handbook of Utility Theory*, volume 2, pages 893–977. Springer, 2004.
- P. Baroni, M. Caminada, and M. Giacomin. Abstract Argumentation Frameworks and Their Semantics. In P. Baroni, D. Gabbay, M. Giacomin, and L. van der Torre, editors, *Handbook of Formal Argumentation*, pages 159–236. College Publications, 2018.
- N. Barrot, J. Lang, and M. Yokoo. Manipulation of Hamming-Based Approval Voting for Multiple Referenda and Committee Elections. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2017)*, pages 597–605. IFAAMAS, 2017.
- J. J. Bartholdi III, C. A. Tovey, and M. A. Trick. The Computational Difficulty of Manipulating an Election. *Social Choice and Welfare*, 6(3):227–241, 1989a.
- J. J. Bartholdi III, 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, 1989b.
- R. Baumann and H. Strass. On the Maximal and Average Numbers of Stable Extensions. In *International Workshop on Theory and Applications of Formal Argumentation (TAFIA-2013)*, pages 111–126. Springer, 2014.

- R. Baumann, G. Wiedemann, M. Heinrich, A. D. Hakimi, and G. Heyer. The Road Map to FAME: A Framework for Mining and Formal Evaluation of Arguments. *Datenbank-Spektrum*, 20(2):107–113, 2020.
- D. Baumeister, D. Neugebauer, and J. Rothe. Collective Acceptability in Abstract Argumentation. *Journal of Applied Logics*, 8(6):1503–1542, 2021.
- D. Baumeister, R. E. Berker, N. Boehmer, S. Bouveret, A. Darmann, P. Faliszewski, M. Lackner, J. Lang, N. Mattei, and A. Novaro. Computational Social Choice: Research & Development. In *Proceedings of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2026)*, pages 3940–3945. IFAAMAS, 2026. “Blue Sky Ideas” track.
- A. Baychkov, M. Lackner, J. Maly, O. Nardi, and J. Peters. Beyond Lower Quota: Avoiding Overrepresentation in Multi-Winner Voting. In *Proceedings of the 27th ACM Conference on Economics and Computation (EC-2026)*. ACM, 2026.
- X. Bei, X. Lu, and W. Suksompong. Truthful Cake Sharing. *Social Choice and Welfare*, 64(1):309–343, 2025.
- J. Bentham. *An Introduction to the Principles of Morals and Legislation*. T. Payne and Son, 1789.
- A. Bergson. A Reformulation of Certain Aspects of Welfare Economics. *The Quarterly Journal of Economics*, 52(2):310–334, 1938.
- M. Bernreiter, W. Dvořák, and S. Woltran. Abstract Argumentation with Conditional Preferences. In *Computational Models of Argument: Proceedings of COMMA 2022*, pages 92–103. IOS Press, 2022.
- M. Bernreiter, W. Dvořák, A. Rapberger, and S. Woltran. The Effect of Preferences in Abstract Argumentation under a Claim-Centric View. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI-2023)*, pages 6253–6261. AAAI Press, 2023.
- M. Bernreiter, J. Maly, O. Nardi, and S. Woltran. Combining Voting and Abstract Argumentation to Understand Online Discussions. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2024)*, pages 170–179. IFAAMAS, 2024a.
- M. Bernreiter, J. Maly, O. Nardi, and S. Woltran. Supplementary Material for “Combining Voting and Abstract Argumentation to Understand Online Discussions”, 2024b. URL <https://doi.org/10.5281/zenodo.10571624>. Repository on Zenodo.
- P. Besnard, A. Garcia, A. Hunter, S. Modgil, H. Prakken, G. Simari, and F. Toni. Introduction to Structured Argumentation. *Argument & Computation*, 5(1):1–4, 2014.
- S. Bistarelli and F. Santini. Weighted Argumentation. *Journal of Applied Logics*, 8(6):1589–1622, 2021.

- D. Black. On the Rationale of Group Decision-Making. *Journal of Political Economy*, 56(1): 23–34, 1948.
- D. Black. *The Theory of Committees and Elections*. Springer, 1958.
- G. A. Bodanza, F. Tohmé, and M. Auday. Collective Argumentation: A Survey of Aggregation Issues Around Argumentation Frameworks. *Argument & Computation*, 8(1):1–34, 2017.
- N. Boehmer, M. Brill, A. Cevallos, J. Gehrlein, L. Sánchez-Fernández, and U. Schmidt-Kraepelin. Approval-Based Committee Voting in Practice: A Case Study of (Over-)Representation in the Polkadot Blockchain. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI-2024)*, pages 9519–9527. AAAI Press, 2024a.
- N. Boehmer, P. Faliszewski, Ł. Janeczko, A. Kaczmarczyk, G. Lisowski, G. Pierczyński, S. Rey, D. Stoliczki, S. Szufa, and T. Wąs. Guide to Numerical Experiments on Elections in Computational Social Choice. In *Proceedings of the 33rd International Joint Conference on Artificial Intelligence (IJCAI-2024)*, pages 7962–7970, 2024b. Survey Track.
- N. Boehmer, M. Brill, and U. Schmidt-Kraepelin. Proportional Representation in Matching Markets: Selecting Multiple Matchings under Dichotomous Preferences. *Social Choice and Welfare*, 64:179–220, 2025a.
- N. Boehmer, S. Fish, and A. D. Procaccia. Generative Social Choice: The Next Generation. In *Proceedings of the 42nd International Conference on Machine Learning (ICML-2025)*, pages 4627–4659. PMLR, 2025b.
- A. Bonner. *The Art and Logic of Ramon Llull: A User’s Guide*, volume 95. Brill, 2007.
- A. Bonner, P. Planas Mulet, and L. Badia. Graphs and Lattices in the Early Versions of Ramon Llull’s Art. *Historia Mathematica*, 70:31–52, 2025.
- J.-C. d. Borda. Mémoire sur les élections au scrutin. *Histoire de l’Académie Royale des Sciences*, 1781.
- S. Botan. Manipulability of Thiele Methods on Party-List Profiles. In *Proceedings of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2021)*, pages 223–231. IFAAMAS, 2021.
- S. J. Brams. Approval Voting in Multicandidate Elections. *Policy Studies Journal*, 9(1), 1980.
- S. J. Brams and P. C. Fishburn. Approval Voting. *American Political Science Review*, 72(3): 831–847, 1978.
- S. J. Brams and P. C. Fishburn. *Approval Voting*. Birkhäuser, 1983.

- S. J. Brams and P. C. Fishburn. Going from Theory to Practice: The Mixed Success of Approval Voting. In J.-F. Laslier and M. R. Sanver, editors, *Handbook on Approval Voting*, pages 19–37. Springer, 2010.
- S. J. Brams and D. M. Kilgour. Satisfaction Approval Voting. In R. Melnik, editor, *Mathematical and Computational Modeling: With Applications in Natural and Social Sciences, Engineering, and the Arts*, pages 273–298. John Wiley & Sons, 2015.
- S. J. Brams, D. M. Kilgour, and W. S. Zwicker. Voting on Referenda: The Separability Problem and Possible Solutions. *Electoral Studies*, 16(3):359–377, 1997.
- S. J. Brams, D. M. Kilgour, and W. S. Zwicker. The Paradox of Multiple Elections. *Social Choice and Welfare*, 15(2):211–236, 1998.
- S. J. Brams, D. M. Kilgour, and M. R. Sanver. A Minimax Procedure for Electing Committees. *Public Choice*, 132(3):401–420, 2007.
- F. Brandl, F. Brandt, D. Peters, and C. Stricker. Distribution Rules under Dichotomous Preferences: Two Out of Three Ain’t Bad. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC-2021)*, pages 158–179. ACM, 2021.
- F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
- R. Bredereck, P. Faliszewski, A. Kaczmarczyk, and R. Niedermeier. An Experimental View on Committees Providing Justified Representation. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI-2019)*, pages 109–115, 2019.
- G. Brewka, M. Truszczynski, and S. Woltran. Representing Preferences Among Sets. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI-2010)*, pages 273–278. AAAI Press, 2010.
- M. Brill and J. Peters. Robust and Verifiable Proportionality Axioms for Multiwinner Voting. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC-2023)*, page 301. ACM, 2023. Full version arXiv:2302.01989 [cs.GT].
- M. Brill, J.-F. Laslier, and P. Skowron. Multiwinner Approval Rules as Apportionment Methods. *Journal of Theoretical Politics*, 30(3):358–382, 2018.
- M. Brill, P. Gözl, D. Peters, U. Schmidt-Kraepelin, and K. Wilker. Approval-Based Apportionment. *Mathematical Programming*, 203(1):77–105, 2022.
- M. Brill, S. Forster, M. Lackner, J. Maly, and J. Peters. Proportionality in Approval-Based Participatory Budgeting. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI-2023)*, pages 5524–5531. AAAI Press, 2023.
- M. Brill, R. Freeman, S. Janson, and M. Lackner. Phragmén’s Voting Methods and Justified Representation. *Mathematical Programming*, 203(1):47–76, 2024.

- L. Bulteau, N. Hazon, R. Page, A. Rosenfeld, and N. Talmon. Justified Representation for Perpetual Voting. *IEEE Access*, 9:96598–96612, 2021.
- J. Burdges, A. Cevallos, P. Czaban, R. Habermeier, S. Hosseini, F. Lama, H. K. Alper, X. Luo, F. Shirazi, A. Stewart, and G. Wood. Overview of Polkadot and its Design Considerations, 2020. URL <https://arxiv.org/abs/2005.13456>. E-print, arXiv:2005.13456 [cs.CR].
- S. R. Buss and L. Hay. On Truth-Table Reducibility to SAT. *Information and Computation*, 91(1):86–102, 1991.
- E. Cabrio and S. Villata. Generating Abstract Arguments: A Natural Language Approach. In *Computational Models of Argument: Proceedings of COMMA 2012*, pages 454–461. IOS Press, 2012.
- M. Caminada and G. Pigozzi. On Judgment Aggregation in Abstract Argumentation. *Autonomous Agents and Multi-Agent Systems*, 22:64–102, 2011.
- C. D. Campbell and G. Tullock. A Measure of the Importance of Cyclical Majorities. *The Economic Journal*, 75(300):853–857, 1965.
- R. Camps, X. Mora, and L. Saumell. The Method of Eneström and Phragmén for Parliamentary Elections by Means of Approval Voting, 2024. URL <https://arxiv.org/abs/1907.10590>. E-print, arXiv:1907.10590 [econ.TH].
- I. Caragiannis, E. Hemaspaandra, and L. A. Hemaspaandra. Dodgson’s Rule and Young’s Rule. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
- L. Carroll. *The Principles of Parliamentary Representation*. Harrison and Sons, 1884.
- A. M. Carstairs. *A Short History of Electoral Systems in Western Europe*. Routledge, 2013.
- A. Casella. Storable Votes. *Games and Economic Behavior*, 51(2):391–419, 2005.
- M. M. Casey and E. Elkind. Justified Representation: From Hare to Droop, 2025. URL <https://arxiv.org/abs/2508.00811>. E-print, arXiv:2508.00811 [cs.GT].
- C. Cayrol and M.-C. Lagasquie-Schiex. On the Acceptability of Arguments in Bipolar Argumentation Frameworks. In *European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty (ECSQARU-2005)*, pages 378–389. Springer, 2005.
- J. Cembrano, J. Correa, U. Schmidt-Kraepelin, A. Tsigonias-Dimitriadis, and V. Verdugo. New Combinatorial Insights for Monotone Apportionment. In *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA-2025)*, pages 1308–1328. SIAM, 2025.

- F. Cerutti, M. Giacomin, and M. Vallati. Generating Structured Argumentation Frameworks: AFBenchGen2. In *Computational Models of Argument: Proceedings of COMMA 2016*, pages 467–468. IOS Press, 2016.
- A. Cevallos and A. Stewart. A Verifiably Secure and Proportional Committee Election Rule. In *Proceedings of the 3rd ACM Conference on Advances in Financial Technologies (AFT-2021)*, pages 29–42. ACM, 2021.
- J. R. Chamberlin and P. N. Courant. Representative Deliberations and Representative Decisions: Proportional Representation and the Borda Rule. *American Political Science Review*, 77(3):718–733, 1983.
- N. Chandak, S. Goel, and D. Peters. Proportional Aggregation of Preferences for Sequential Decision Making. *Journal of Artificial Intelligence Research*, 85, 2026.
- W. Chen and U. Endriss. Preservation of Semantic Properties in Collective Argumentation: The Case of Aggregating Abstract Argumentation Frameworks. *Artificial Intelligence*, 269:27–48, 2019.
- J. Cohen. Deliberative Democracy. In S. W. Rosenberg, editor, *Deliberation, Participation and Democracy: Can the People Govern?*, pages 219–236. Palgrave Macmillan UK, 2007.
- J. M. Colomer. Ramon Llull: From ‘Ars Electionis’ to Social Choice Theory. *Social Choice and Welfare*, 40(2):317–328, 2013.
- J. M. Colomer and I. McLean. Electing Popes: Approval Balloting and Qualified-Majority Rule. *The Journal of Interdisciplinary History*, 29(1):1–22, 1998.
- M. d. Condorcet. *Essai sur l’application de l’analyse à la probabilité des décisions rendues à la pluralité des voix*. Imprimerie Royale, 1785.
- V. Conitzer and T. Walsh. Barriers to Manipulation in Voting. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
- V. Conitzer and L. Xia. Paradoxes of Multiple Elections: An Approximation Approach. In *Proceedings of the 13th International Conference on Principles of Knowledge Representation and Reasoning (KR-2012)*, 2012.
- V. Conitzer, M. Rognlie, and L. Xia. Preference Functions that Score Rankings and Maximum Likelihood Estimation. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI-2009)*, pages 109–115, 2009.
- V. Conitzer, R. Freeman, and N. Shah. Fair Public Decision Making. In *Proceedings of the 2017 ACM Conference on Economics and Computation (EC-2017)*, pages 629–646. ACM, 2017.

- S. A. Cook. The Complexity of Theorem-Proving Procedures. In *Proceedings of the 3rd Annual ACM Symposium on Theory of Computing (STOC-1971)*, pages 151–158. ACM, 1971.
- A. H. Copeland. A ‘Reasonable’ Social Welfare Function. Seminar on Mathematics in Social Sciences, University of Michigan. Unpublished, 1951.
- S. Coste-Marquis, C. Devred, and P. Marquis. Symmetric Argumentation Frameworks. In *European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty (ECSQARU-2005)*, pages 317–328. Springer, 2005.
- S. Coste-Marquis, C. Devred, S. Konieczny, M.-C. Lagasquie-Schiex, and P. Marquis. On the Merging of Dung’s Argumentation Systems. *Artificial Intelligence*, 171(10):730–753, 2007.
- N. Creignou, R. Pichler, and S. Woltran. Do Hard SAT-Related Reasoning Tasks Become Easier in the Krom Fragment? *Logical Methods in Computer Science*, 14(4), 2018.
- G. Cullity. The Free Rider Problem. In E. N. Zalta and U. Nodelman, editors, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Fall 2025 edition, 2025. URL <https://plato.stanford.edu/archives/fall2025/entries/free-rider/>.
- M. Cygan, F. V. Fomin, Ł. Kowalik, D. Lokshtanov, D. Marx, M. Pilipczuk, M. Pilipczuk, and S. Saurabh. *Parameterized Algorithms*. Springer, 2015.
- G. Dalaqua. Representative Democracy in Mill: A Response to Rousseau’s Critique of Representation. *Revue D’Études Benthamiennes*, 27, 2025.
- S. De, L. Gelauff, A. Goel, S. Milli, A. D. Procaccia, and A. Siu. Question the Questions: Auditing Representation in Online Deliberative Processes. In *Proceedings of the ACM Web Conference 2026 (WWW-2026)*, pages 1640–1650. ACM, 2026.
- T. Delemazure, T. Demeulemeester, M. Eberl, J. Israel, and P. Lederer. Strategyproofness and Proportionality in Party-Approval Multiwinner Elections. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI-2023)*, pages 5591–5599. AAAI Press, 2023.
- J. Delobelle, A. Haret, S. Konieczny, J.-G. Mailly, J. Rossit, and S. Woltran. Merging of Abstract Argumentation Frameworks. In *Proceedings of the 15th International Conference on Principles of Knowledge Representation and Reasoning (KR-2016)*, 2016.
- V. D’Hondt. *Question électorale: la représentation proportionnelle des partis*. Bruylant-Christophe & Cie, 1878.
- C. L. Dodgson. A Method of Taking Votes on More than Two Issues, 1876. Reprinted in McLean and Urken (1995).

- C. L. Dodgson. *Principles of Parliamentary Representation*. Harrison and Sons, 1884.
- S. Dovi. Political Representation. In E. N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Fall 2018 edition, 2018. URL <https://plato.stanford.edu/archives/fall2018/entries/political-representation/>.
- A. Downs. *An Economic Theory of Democracy*. Harper, 1957.
- H. R. Droop. On Methods of Electing Representatives. *Journal of the Statistical Society of London*, 44(2):141–202, 1881.
- P. M. Dung. On the Acceptability of Arguments and its Fundamental Role in Nonmonotonic Reasoning, Logic Programming and n -Person Games. *Artificial Intelligence*, 77(2): 321–358, 1995.
- P. E. Dunne. Computational Properties of Argument Systems Satisfying Graph-Theoretic Constraints. *Artificial Intelligence*, 171(10-15):701–729, 2007.
- P. E. Dunne, A. Hunter, P. McBurney, S. Parsons, and M. Wooldridge. Weighted Argument Systems: Basic Definitions, Algorithms, and Complexity Results. *Artificial Intelligence*, 175(2):457–486, 2011.
- P. E. Dunne, P. Marquis, and M. Wooldridge. Argument Aggregation: Basic Axioms and Complexity Results. In *Computational Models of Argument: Proceedings of COMMA 2012*, pages 129–140. IOS Press, 2012.
- P. E. Dunne, W. Dvořák, T. Linsbichler, and S. Woltran. Characteristics of Multiple Viewpoints in Abstract Argumentation. *Artificial Intelligence*, 228:153–178, 2015.
- W. Dvořák and S. Woltran. Complexity of Abstract Argumentation under a Claim-Centric View. *Artificial Intelligence*, 285:103290, 2020.
- W. Dvořák and P. E. Dunne. Computational Problems in Formal Argumentation and their Complexity. In P. Baroni, D. Gabbay, M. Giacomin, and L. van der Torre, editors, *Handbook of Formal Argumentation*, pages 631–687. College Publications, 2018.
- M. Döring, M. Brill, and J. Heitzig. The River Voting Method. In *Proceedings of the 40th AAAI Conference on Artificial Intelligence (AAAI-2026)*, pages 16846–16854. AAAI Press, 2026.
- U. Egly, S. A. Gaggl, and S. Woltran. Aspartix: Implementing Argumentation Frameworks Using Answer-Set Programming. In *Proceedings of the 24th International Conference on Logic Programming (ICLP-2008)*, pages 734–738. Springer, 2008.
- E. Elkind and A. Ismaili. OWA-Based Extensions of the Chamberlin–Courant Rule. In *Proceedings of the 4th International Conference on Algorithmic Decision Theory (ADT-2015)*, pages 486–502. Springer, 2015.

- E. Elkind, P. Faliszewski, J.-F. Laslier, P. Skowron, A. Slinko, and N. Talmon. What Do Multiwinner Voting Rules Do? An Experiment Over the Two-Dimensional Euclidean Domain. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI-2017)*, pages 494–501. AAAI Press, 2017.
- E. Elkind, T. Y. Neoh, and N. Teh. Temporal Elections: Welfare, Strategyproofness, and Proportionality. In *Proceedings of the 27th European Conference on Artificial Intelligence (ECAI-2024)*, pages 3292–3299. IOS Press, 2024a.
- E. Elkind, S. Obraztsova, and N. Teh. Temporal Fairness in Multiwinner Voting. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI-2024)*, pages 22633–22640. AAAI Press, 2024b.
- E. Elkind, S. Obraztsova, J. Peters, and N. Teh. Verifying Proportionality in Temporal Voting. In *Proceedings of the 39th AAAI Conference on Artificial Intelligence (AAAI-2025)*, pages 13805–13813. AAAI Press, 2025.
- J. Elster. Introduction. In J. Elster, editor, *Deliberative Democracy*, volume 1. Cambridge University Press, 1998.
- J. M. Enelow and M. J. Hinich. *Advances in the Spatial Theory of Voting*. Cambridge University Press, 1990.
- G. H. Eneström. Om aritmetiska och statistiska metoder för proportionella val. *Öfversigt af Kongliga Vetenskaps-Akademiens förhandlingar*, 53:543–570, 1896.
- P. Faliszewski, P. Skowron, A. Slinko, and N. Talmon. Multiwinner Voting: A New Challenge for Social Choice Theory. In U. Endriss, editor, *Trends in Computational Social Choice*, pages 27–47. AI Access, 2017.
- J. K. Fichte, M. Hecher, Y. Mahmood, and Z. Wang. Diversity of Extensions in Abstract Argumentation. In *Proceedings of the 35th International Joint Conference on Artificial Intelligence (IJCAI-2026)*, 2026.
- S. Fish, P. Gözl, D. C. Parkes, A. D. Procaccia, G. Rusak, I. Shapira, and M. Wüthrich. Generative Social Choice. *Journal of the ACM*, 73(2):1–52, 2026.
- P. C. Fishburn. *The theory of social choice*. Princeton University Press, 1973.
- P. C. Fishburn. An Analysis of Simple Voting Systems for Electing Committees. *SIAM Journal on Applied Mathematics*, 41(3):499–502, 1981a.
- P. C. Fishburn. Majority Committees. *Journal of Economic Theory*, 25(2):255–268, 1981b.
- R. Freeman, M. Brill, and V. Conitzer. General Tiebreaking Schemes for Computational Social Choice. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2015)*, pages 1401–1409. IFAAMAS, 2015.

- R. Freeman, S. M. Zahedi, and V. Conitzer. Fair and Efficient Social Choice in Dynamic Settings. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI-2017)*, pages 4580–4587, 2017.
- E. J. Friedman, V. Gkatzelis, C.-A. Psomas, and S. Shenker. Fair and Efficient Memory Sharing: Confronting Free Riders. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI-2019)*, volume 33, pages 1965–1972. AAAI Press, 2019.
- R. Frigg and S. Hartmann. Models in Science. In E. N. Zalta and U. Nodelman, editors, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Winter 2025 edition, 2025. URL <https://plato.stanford.edu/archives/win2025/entries/models-science/>.
- R. Fuller. *SDI Digital Democracy Report*. Solonian Democracy Institute, 2022.
- G. Gabriel and U. Dathe, editors. *Gottlob Frege - Werk und Wirkung. Mit unveröffentlichten Vorschlägen für ein Wahlgesetz von Gottlob Frege*. Brill, 2000.
- M. Gallagher and P. Mitchell, editors. *The Politics of Electoral Systems*. OUP Oxford, 2005.
- J. Ganzer, N. Criado, M. Lopez-Sanchez, S. Parsons, and J. A. Rodriguez-Aguilar. A Model to Support Collective Reasoning: Formalization, Analysis and Computational Assessment. *Journal of Artificial Intelligence Research*, 77:1021–1086, 2023.
- M. Gardner. Quotation from an interview, 1998. In Kendrick Frazier, *A Mind at Play: An Interview with Martin Gardner*, *Skeptical Inquirer*, 22(2), 34–39.
- M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. Freeman, 1979.
- S. Gaspers and R. Li. Enumeration of Preferred Extensions in Almost Oriented Digraphs. In *Proceedings of the 44th International Symposium on Mathematical Foundations of Computer Science (MFCS-2019)*, pages 74:1–74:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.
- M. Gebser, R. Kaminski, B. Kaufmann, and T. Schaub. *Answer Set Solving in Practice*. Springer Nature, 2022.
- A. Gibbard. Manipulation of Voting Schemes: A General Result. *Econometrica: Journal of the Econometric Society*, 41(4):587–601, 1973.
- C. Girard. Acclamation Voting in Sparta: An Early Use of Approval Voting. In J.-F. Laslier and M. R. Sanver, editors, *Handbook on Approval Voting*, pages 15–17. Springer, 2010.
- M. T. Godziszewski, P. Batko, P. Skowron, and P. Faliszewski. An Analysis of Approval-Based Committee Rules for 2D-Euclidean Elections. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI-2021)*, pages 5448–5455. AAAI Press, 2021.

- P. Goffredo, E. Cabrio, S. Villata, S. Haddadan, and J. T. Sanchez. DISPUTool 2.0: A Modular Architecture for Multi-Layer Argumentative Analysis of Political Debates. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI-2023)*, pages 16431–16433. AAAI Press, 2023.
- O. Goldreich. On Promise Problems: A Survey. In O. Goldreich, A. L. Rosenberg, and A. L. Selman, editors, *Theoretical Computer Science: Essays in Memory of Shimon Even*, pages 254–290. Springer, 2006.
- I. J. Good and T. N. Tideman. A Principle for Selecting a Representative Committee. Unpublished manuscript, Virginia Polytechnic Institute and State University, 1976.
- D. Gorur, A. Rago, and F. Toni. Can Large Language Models Perform Relation-Based Argument Mining? In *Proceedings of the 31st International Conference on Computational Linguistics (COLING-2025)*, pages 8518–8534. Association for Computational Linguistics, 2025.
- U. Grandi, R. Bredereck, T. Delemazure, U. Endriss, J. Maly, N. Mattei, N. Maudet, O. Nardi, and S. Szufa. Social Choice with Text: Collective Decision Making in the LLM Era. In *Proceedings of the ACM AI Leadership Summit*, 2026. E-print available at HAL Open Science Archive (hal-05548117): <https://hal.science/hal-05548117>.
- D. Grossi and G. Pigozzi. *Judgment Aggregation: A Primer*. Springer Nature, 2022.
- T. Groves and J. Ledyard. Optimal Allocation of Public Goods: A Solution to the “Free Rider” Problem. *Econometrica: Journal of the Econometric Society*, pages 783–809, 1977.
- D. Halpern, G. Kehne, A. D. Procaccia, J. Tucker-Foltz, and M. Wüthrich. Representation with Incomplete Votes. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI-2023)*. AAAI Press, 2023.
- T. Hare. *The Machinery of Representation*. W. Maxwell, 1857.
- P. Harrenstein, M.-L. Lackner, and M. Lackner. A Mathematical Analysis of an Election System Proposed by Gottlob Frege. *Erkenntnis*, 87(6):2609–2644, 2020.
- E. Hemaspaandra, L. A. Hemaspaandra, and J. Rothe. Exact Analysis of Dodgson Elections: Lewis Carroll’s 1876 Voting System is Complete for Parallel Access to NP. *Journal of the ACM (JACM)*, 44(6):806–825, 1997.
- E. S. Herron, R. J. Pekkanen, and M. S. Shugart, editors. *The Oxford Handbook of Electoral Systems*. Oxford University Press, 2018.
- C. G. Hoag and G. H. Hallett. *Proportional Representation*. The Macmillan Company, 1926.
- D. S. Hochbaum. Approximation Algorithms for NP-Hard Problems. *SIGACT News*, 28(2):40–52, 1997.

- W. H. Holliday and E. Pacuit. Split Cycle: A New Condorcet-Consistent Voting Method Independent of Clones and Immune to Spoilers. *Public Choice*, 197(1):1–62, 2023.
- A. Hylland. Proportional Representation without Party Lists. In R. Malnes and A. Underdal, editors, *Rationality and Institutions: Essays in Honour of Knut Midgaard on the Occasion of his 60th Birthday*, pages 126–153. Universitetsforlaget, 1992.
- A. Igarashi, M. Lackner, O. Nardi, and A. Novaro. Repeated Fair Allocation of Indivisible Items. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI-2024)*, pages 9781–9789. AAAI Press, 2024. An earlier version was presented at COMSOC 2023.
- S. Janson. Phragmén’s and Thiele’s Election Methods, 2018. URL <https://arxiv.org/abs/1611.08826>. E-print, arXiv:1611.08826 [cs.GT].
- C. Jerrett and E. Anshelevich. Low Cost, Fair, and Representative Committees in a Metric Space, 2025. URL <https://arxiv.org/abs/2510.21039>. E-print, arXiv:2510.21039 [cs.GT].
- S. Kaci, L. W. N. van der Torre, S. Vesic, and S. Villata. Preference in Abstract Argumentation. In D. Gabbay, M. Giacomin, G. Simari, and M. Thimm, editors, *Handbook of Formal Argumentation*, volume 2, pages 211–248. College Publications, 2021.
- I. Kahana and N. Hazon. The Leximin Approach for a Sequence of Collective Decisions. In *Proceedings of the 26th European Conference on Artificial Intelligence (ECAI-2023)*, pages 1198–1206. IOS Press, 2023.
- R. M. Karp. Reducibility Among Combinatorial Problems. In *Proceedings of a Symposium on the Complexity of Computer Computations*, The IBM Research Symposia Series, pages 85–103. Plenum Press, 1972.
- L. Kellerhals and J. Peters. Proportional Fairness in Clustering: A Social Choice Perspective. In *Advances in Neural Information Processing Systems 37 (NeurIPS-2024)*, pages 111299–111317. Curran Associates Inc., 2024.
- D. M. Kilgour. Approval Balloting for Multi-Winner Elections. In J.-F. Laslier and M. R. Sanver, editors, *Handbook on Approval Voting*, pages 105–124. Springer, 2010.
- B. Kluiving, A. de Vries, P. Vrijbergen, A. Boixel, and U. Endriss. Analysing Irresolute Multiwinner Voting Rules with Approval Ballots via SAT Solving. In *Proceedings of the 24th European Conference on Artificial Intelligence (ECAI-2020)*, pages 131–138. IOS Press, 2020.
- A. Kozachinskiy, A. Shen, and T. Steifer. Optimal Bounds for Dissatisfaction in Perpetual Voting. In *Proceedings of the 39th AAAI Conference on Artificial Intelligence (AAAI-2025)*, pages 13977–13984. AAAI Press, 2025.

- S. Kraiczý and E. Elkind. A Lower Bound for Local Search Proportional Approval Voting. In *Proceedings of the 32nd European Symposium on Algorithms (ESA-2024)*, pages 82:1–82:14, 2024.
- M. Lackner. Perpetual Voting: Fairness in Long-Term Decision Making. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI-2020)*, pages 2103–2110. AAAI Press, 2020.
- M. Lackner and J. Maly. Proportional Decisions in Perpetual Voting. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI-2023)*. AAAI Press, 2023.
- M. Lackner and O. Nardi. Avoiding Overrepresentation: Upper Quota Axioms for Committee Voting, 2025. Presented at M-Pref 2025.
- M. Lackner and P. Skowron. Approval-Based Multi-Winner Rules and Strategic Voting. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI-2018)*, pages 340–346, 2018.
- M. Lackner and P. Skowron. Consistent Approval-Based Multi-Winner Rules. *Journal of Economic Theory*, 192:105173, 2021.
- M. Lackner and P. Skowron. *Multi-Winner Voting with Approval Preferences*. Springer, 2023.
- M. Lackner, J. Maly, and O. Nardi. Free-Riding in Multi-Issue Decisions. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2023)*, pages 2040–2048. IFAAMAS, 2023a.
- M. Lackner, P. Regner, and B. Krenn. abcvoting: A Python Package for Approval-Based Multi-Winner Voting Rules. *Journal of Open Source Software*, 8(81):4880, 2023b.
- M. Lackner, J. Maly, and O. Nardi. martinlackner/free-riding: Journal Submission, 2025. URL <https://doi.org/10.5281/zenodo.17963528>. Repository on Zenodo.
- M. Lackner, J. Maly, and O. Nardi. Free-Riding in Multi-Issue Decisions. *Journal of Artificial Intelligence Research*, 86, 2026.
- G. Laffond and J. Lainé. Does Choosing Committees from Approval Balloting fulfill the Electorate’s Will? In J.-F. Laslier and M. R. Sanver, editors, *Handbook on Approval Voting*, pages 125–150. Springer, 2010.
- S. P. Lalley and E. G. Weyl. Quadratic Voting: How Mechanism Design can Radicalize Democracy. *AEA Papers and Proceedings*, 108:33–37, 2018.
- J. Lang and L. Xia. Sequential Composition of Voting Rules in Multi-Issue Domains. *Mathematical Social Sciences*, 57(3):304–324, 2009.

- J. Lang and L. Xia. Voting in Combinatorial Domains. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
- J.-F. Laslier and M. R. Sanver, editors. *Handbook on Approval Voting*. Springer, 2010a.
- J.-F. Laslier and M. R. Sanver. Introduction to the Handbook on Approval Voting. In J.-F. Laslier and M. R. Sanver, editors, *Handbook on Approval Voting*, pages 1–12. Springer, 2010b.
- B. E. Lee. Representing the Insincere: Strategically Robust Proportional Representation, 2018. URL <https://arxiv.org/abs/1801.09346>. E-print, arXiv:1801.09346 [cs.GT].
- R. LeGrand. Analysis of the Minimax Procedure. Technical report, Department of Computer Science and Engineering, Washington University, St. Louis, 2004. WUCSE-2004-67.
- R. LeGrand, E. Markakis, and A. Mehta. Some Results on Approximating the Minimax Solution in Approval Voting. In *Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2007)*, pages 1185–1187. IFAAMAS, 2007.
- J. Leite and J. G. Martins. Social Abstract Argumentation. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI-2011)*, pages 2287–2292, 2011.
- L. A. Levin. Universal Search Problems. *Problems of Information Transmission*, 9(3):265–266, 1973. Originally in Russian.
- H. Li, V. Schlegel, Y. Sun, R. Batista-Navarro, and G. Nenadic. Large Language Models in Argument Mining: A Survey, 2025. URL <https://arxiv.org/abs/2506.16383>. E-print, arXiv:2506.16383 [cs.CL].
- A. Lijphart. The Field of Electoral Systems Research: A Critical Survey. *Electoral Studies*, 4(1):3–14, 1985.
- A. Lijphart. *Patterns of Democracy: Government Forms and Performance in Thirty-Six Countries*. Yale University Press, 1999.
- M. Lines. Approval Voting and Strategy Analysis: A Venetian Example. *Theory and Decision*, 20(2):155–172, 1986.
- C. L. Mallows. Non-Null Ranking Models. I. *Biometrika*, 44(1):114–130, 1957.
- J. Maly, O. Nardi, and Z. Terzopoulou. Sociotropic Behavior in Voting, 2025. Presented as a poster at COMSOC 2025.

- J. Maly, O. Nardi, and S. Rey. Proportionality Press: Illustrating Proportionality through Comment Selection. In *Proceedings of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2026)*, pages 4104–4106. IFAAMAS, 2026. Demo track.
- T. Masařík, G. Pierczyński, and P. Skowron. A Generalised Theory of Proportionality in Collective Decision Making. In *Proceedings of the 25th ACM Conference on Economics and Computation (EC-2024)*, pages 734–754. ACM, 2024.
- J. P. Mayberry. Quota Methods for Congressional Apportionment are Still Non-Unique. *Proceedings of the National Academy of Sciences*, 75(8):3537–3539, 1978.
- I. McLean. Voting. In R. Wilson and A. Moktefi, editors, *The Mathematical World of Charles L. Dodgson (Lewis Carroll)*. Oxford University Press, 2019.
- I. McLean, A. B. Urken, and F. Hewitt. *Classics of Social Choice*. University of Michigan Press, 1995.
- I. S. McLean, A. McMillan, and B. L. Monroe. *A Mathematical Approach to Proportional Representation: Duncan Black on Lewis Carroll*. Springer, 2012.
- R. Meir. Iterative Voting. In U. Endriss, editor, *Trends in Computational Social Choice*, pages 69–86. AI Access, 2017.
- R. Meir. *Strategic Voting*. Morgan & Claypool Publishers, 2018.
- R. Meir, A. D. Procaccia, and J. S. Rosenschein. A Broader Picture of the Complexity of Strategic Behavior in Multi-Winner Elections. In *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2008)*, volume 2, pages 991–998. IFAAMAS, 2008a.
- R. Meir, A. D. Procaccia, J. S. Rosenschein, and A. Zohar. Complexity of Strategic Behavior in Multi-Winner Elections. *Journal of Artificial Intelligence Research*, 33:149–178, 2008b.
- J. S. Mill. *Considerations on Representative Government*. Longman, Green, Longman, Roberts, and Green, 1865.
- S. Modgil. Reasoning about Preferences in Argumentation Frameworks. *Artificial Intelligence*, 173(9-10):901–934, 2009.
- D. Monderer and L. S. Shapley. Potential Games. *Games and Economic Behavior*, 14(1): 124–143, 1996.
- B. L. Monroe. Fully Proportional Representation. *American Political Science Review*, 89(4): 925–940, 1995.
- H. Moulin. Condorcet’s Principle Implies the No Show Paradox. *Journal of Economic Theory*, 45(1):53–64, 1988.

- M. A. Müller, B. I. Urh, T. Zotescu, and U. Endriss. Breaking the Cycle. Preference-Based Aggregation for Cyclic Argumentation Frameworks. In *Computational Models of Argument: Proceedings of COMMA 2024*, pages 157–168. IOS Press, 2024.
- R. G. Niemi. The Problem of Strategic Behavior under Approval Voting. *American Political Science Review*, 78(4):952–958, 1984.
- A. Niskanen and M. Järvisalo. τ -toksia: An Efficient Abstract Argumentation Reasoner. In *Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning (KR-2020)*, pages 800–804, 2020.
- P. Norris. Choosing Electoral Systems: Proportional, Majoritarian and Mixed Systems. *International Political Science Review*, 18(3):297–312, 1997.
- OECD. *Innovative Citizen Participation and new Democratic Institutions: Catching the Deliberative Wave*. OECD Publishing, 2020.
- O. Olabisi and A. Agrawal. Understanding Position Bias Effects on Fairness in Social Multi-Document Summarization. In *Proceedings of the Eleventh Workshop on NLP for Similar Languages, Varieties, and Dialects (VarDial-2024)*, pages 117–129. Association for Computational Linguistics, 2024.
- M. J. Olson. *The Logic of Collective Action*. Harvard University Press, 1965.
- R. M. Palau and M. Moens. Argumentation Mining: The Detection, Classification and Structure of Arguments in Text. In *Proceedings of the 12th International Conference on Artificial Intelligence and Law (ICAIL-2009)*, pages 98–107. ACM, 2009.
- D. Peters. Proportionality and Strategyproofness in Multiwinner Elections. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2018)*, pages 1549–1557. IFAAMAS, 2018.
- D. Peters and P. Skowron. Proportionality and the Limits of Welfarism. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC-2020)*, pages 793–794. ACM, 2020.
- D. Peters, G. Pierczyński, and P. Skowron. Proportional Participatory Budgeting with Additive Utilities. In *Advances in Neural Information Processing Systems 34 (NeurIPS-2021)*, pages 12726–12737. Curran Associates Inc., 2021.
- J. Peters. *Facets of Proportionality*. PhD Thesis, Technische Universität Berlin, 2025. Available at <https://depositonce.tu-berlin.de/items/0870641f-819e-4d57-b6e9-cef23a6858cb>.
- L. E. Phragmén. Sur une méthode nouvelle pour réaliser, dans les élections, la représentation proportionnelle des partis. *Öfversigt af Kongliga Vetenskaps-Akademiens Förhandlingar*, 51(3):133–137, 1894.

- L. E. Phragmén. *Proportionella val. En valteknisk studie*. Svenska spörsmål 25. Lars Hökersbergs förlag, 1895.
- R. F. Potthoff and S. J. Brams. Proportional Representation: Broadening the Options. *Journal of Theoretical Politics*, 10(2):147–178, 1998.
- G. B. Powell. *Elections as Instruments of Democracy: Majoritarian and Proportional Visions*. Yale University Press, 2000.
- A. D. Procaccia, J. S. Rosenschein, and A. Zohar. Multi-Winner Elections: Complexity of Manipulation, Control and Winner-Determination. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI-2007)*, pages 1476–1481, 2007.
- A. D. Procaccia, J. S. Rosenschein, and A. Zohar. On the Complexity of Achieving Proportional Representation. *Social Choice and Welfare*, 30:353–362, 2008.
- F. Pukelsheim. *Proportional Representation*. Springer, 2017.
- J. Rawls. *A Theory of Justice*. Belknap Press, 1971.
- M. Revel, S. Milli, T. Lu, J. Watson-Daniels, and M. Nickel. Representative Ranking for Deliberation in the Public Sphere. In *Proceedings of the 42nd International Conference on Machine Learning (ICML-2025)*. PMLR, 2025.
- S. Rey, F. Schmidt, and J. Maly. The (Computational) Social Choice Take on Indivisible Participatory Budgeting, 2025. URL <https://arxiv.org/abs/2303.00621>. E-print, arXiv:2303.00621 [cs.GT].
- R. Rogowski. Representation in Political Theory and in Law. *Ethics*, 91(3):395–430, 1981.
- J. Rossie, J. Delobelle, S. Konieczny, C. Lens, and S. Vesic. Collective Satisfaction Semantics for Opinion Based Argumentation. In *21st International Conference on Principles of Knowledge Representation and Reasoning (KR-2024)*, pages 631–641, 2024.
- J.-J. Rousseau. *On the Social Contract; or, Principles of Political Right*. Marc-Michel Rey, 1762. Translated by G. D. H. Cole.
- A. Sainte-Laguë. La représentation proportionnelle et la méthode des moindres carrés. *Annales scientifiques de l’école Normale Supérieure*, 27:529–542, 1910.
- P. A. Samuelson. *Foundations of Economic Analysis*. Harvard University Press, 1947.
- P. A. Samuelson. The Pure Theory of Public Expenditure. *The Review of Economics and Statistics*, pages 387–389, 1954.
- L. Sánchez-Fernández, N. Fernández, J. A. Fisteus, and M. Brill. The Maximin Support Method: An Extension of the D’Hondt Method to Approval-Based Multiwinner Elections. *Mathematical Programming*, 203(1–2):107–134, 2024.

- M. A. Satterthwaite. Strategy-Proofness and Arrow's Conditions: Existence and Correspondence Theorems for Voting Procedures and Social Welfare Functions. *Journal of Economic Theory*, 10(2):187–217, 1975.
- M. Schulze. Free Riding. *Voting Matters*, (18):2–8, 2004.
- M. Schulze. A New Monotonic, Clone-Independent, Reversal Symmetric, and Condorcet-Consistent Single-Winner Election Method. *Social Choice and Welfare*, 36(2):267–303, 2011. Preliminary version in *Voting Matters*, 17:9-19, 2003.
- M. Schulze. The Schulze Method of Voting, 2025. URL <https://arxiv.org/abs/1804.02973>. E-print, arXiv:1804.02973 [cs.GT].
- A. Sen. *Collective Choice and Social Welfare: An Expanded Edition*. Harvard University Press, 2017.
- P. Skowron and A. Górecki. Proportional Public Decisions. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI-2022)*, pages 5191–5198. AAAI Press, 2022.
- P. Skowron, P. Faliszewski, and J. Lang. Finding a Collective Set of Items: From Proportional Multirepresentation to Group Recommendation. *Artificial Intelligence*, 241: 191–216, 2016.
- C. T. Small, I. Vetrov, E. Durmus, H. Homaei, E. Barry, J. Cornebise, T. Suzman, D. Ganguli, and C. Megill. Opportunities and Risks of LLMs for Scalable Deliberation with Polis, 2023. URL <https://arxiv.org/abs/2306.11932>. E-print, arXiv:2306.11932 [cs.SI].
- R. M. Starr. Arrow, Kenneth Joseph (born 1921). In S. N. Durlauf and L. Blume, editors, *The New Palgrave Dictionary of Economics*, pages 433–446. Springer, 2018.
- K. Suzumura. Introduction. In K. J. Arrow, A. Sen, and K. Suzumura, editors, *Handbook of Social Choice and Welfare*. Elsevier, 2010.
- S. Szufa, P. Faliszewski, P. Skowron, A. Slinko, and N. Talmon. Drawing a Map of Elections in the Space of Statistical Cultures. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2020)*. IFAAMAS, 2020.
- S. Szufa, P. Faliszewski, Ł. Janeczko, M. Lackner, A. Slinko, K. Sornat, and N. Talmon. How to Sample Approval Elections? In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI-2022)*, pages 496–502, 2022.
- L. Sánchez-Fernández, E. Elkind, M. Lackner, N. Fernández García, J. A. Fisteus, P. Basanta Val, and P. Skowron. Proportional Justified Representation. *Artificial Intelligence*, 353:104503, 2026.

- M. H. Tessler, M. A. Bakker, D. Jarrett, H. Sheahan, M. J. Chadwick, R. Koster, G. Evans, L. Campbell-Gillingham, T. Collins, D. C. Parkes, M. Botvinick, and C. Summerfield. AI can Help Humans Find Common Ground in Democratic Deliberation. *Science*, 386 (6719), 2024.
- T. N. Thiele. Om Flerfoldsvvalg. *Oversigt over det Kongelige Danske Videnskabernes Selskabs Forhandlinger*, pages 415–441, 1895.
- F. A. Tohmé, G. A. Bodanza, and G. R. Simari. Aggregation of Attack Relations: A Social-Choice Theoretical Analysis of Defeasibility Criteria. In *Proceedings of the 5th International Symposium on Foundations of Information and Knowledge Systems (FoIKS-2008)*, pages 8–23. Springer, 2008.
- R. Tuck. *Free Riding*. Harvard University Press, 2008.
- UNESCO. The Decreta of León of 1188: The Oldest Documentary Manifestation of the European Parliamentary System, 2012. URL https://media.unesco.org/sites/default/files/webform/mow001/spain_decreta_of_leon.pdf. Accessed: 2026-06-08.
- D. van Dalen. *Logic and Structure*. Springer, 2004.
- K. W. Wagner. Bounded Query Classes. *SIAM Journal on Computing*, 19(5):833–846, 1990.
- B. Wampler, S. McNulty, and M. Touchton. *Participatory Budgeting in Global Perspective*. Oxford University Press, 2021.
- R. J. Weber. Comparison of Public Choice Systems. Cowles Foundation Discussion Paper Discussion Paper No. 732, Cowles Foundation for Research in Economics, Yale University, 1978.
- R. J. Weber. Approval Voting. *Journal of Economic Perspectives*, 9(1):39–49, 1995.
- W. Whitman. *Leaves of Grass*. Rees Welsh & Co., 1882.
- L. A. Wolsey. *Integer Programming*. John Wiley & Sons, 2020.
- R. R. Yager. On Ordered Weighted Averaging Aggregation Operators in Multicriteria Decisionmaking. *IEEE Transactions on Systems, Man, and Cybernetics*, 18(1):183–190, 1988.
- H. P. Young. Social Choice Scoring Functions. *SIAM Journal on Applied Mathematics*, 28 (4):824–838, 1975.
- Y. Zhang, N. Zhang, Y. Liu, A. R. Fabbri, J. Liu, R. Kamo, X. Lu, C. Xiong, J. Zhao, D. Radev, et al. Fair Abstractive Summarization of Diverse Perspectives. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-2024)*, pages 3404–3426. Association for Computational Linguistics, 2024.