On Proportionality in Complex Domains

Julian Chingoma

On Proportionality in Complex Domains

ILLC Dissertation Series DS-2025-01

INSTITUTE FOR LOGIC, LANGUAGE AND COMPUTATION

For further information about ILLC-publications, please contact

Institute for Logic, Language and Computation
Universiteit van Amsterdam
Science Park 107
1098 XG Amsterdam
phone: +31-20-525 6051
e-mail: illc@uva.nl
homepage: http://www.illc.uva.nl/

These investigations were supported by the Dutch Research Council (NWO) in the context of the Collective Information project funded under the VICI scheme (grant number 639.023.811).

Cover design by Julian Chingoma.
Printed and bound by Ipskamp Printing.

ISBN: 978–94–6473–661–8

On Proportionality in Complex Domains

Academisch Proefschrift

ter verkrijging van de graad van doctor
aan de Universiteit van Amsterdam
op gezag van de Rector Magnificus
prof. dr. ir. P.P.C.C. Verbeek
ten overstaan van een door het College voor Promoties ingestelde commissie,
in het openbaar te verdedigen in de Agnietenkapel
op donderdag 23 januari 2025, te 14.00 uur

door

Julian Zvikomborero Chingoma

geboren te Harare

Promotiecommissie

Promotor:	prof. dr. U. Endriss	Universiteit van Amsterdam
Co-promotor:	dr. R. de Haan	Universiteit van Amsterdam
Overige leden:	prof. dr. P. Faliszewski	AGH University of Science
	dr. rer. nat. R.E.M. Reiffenhäuser	Universiteit van Amsterdam
	prof. dr. G. Schäfer	Universiteit van Amsterdam
	dr. rer. nat. U. Schmidt-Kraepelin	TU Eindhoven
	prof. dr. Y. Venema	Universiteit van Amsterdam

Faculteit der Natuurwetenschappen, Wiskunde en Informatica

Samenvatting

Het enorme aanbod aan collectieve besluitvormingsscenario's die zich in de echte wereld voordoen heeft een schat aan voorbeelden opgeleverd voor onderzoekers in de computationale socialekeuzetheorie om te modelleren en te analyseren. Het vakgebied van de computationele socialekeuzetheorie neemt het perspectief van de computerwetenschap om methoden te bestuderen die worden gebruikt om individuele meningen samen te voegen (te aggregeren) om zo tot een enkele collectieve uitkomst te komen. In het klassieke voorbeeld van de socialekeuzetheorie stemt een electoraat op kandidaten bij een verkiezing. Normaal gesproken is het doel bij deze verkiezingen om één winnaar te kiezen, zoals een president. Dit proefschrift kijkt echter verder dan het geval van verkiezingen met één enkele winnaar en richt zich op gevallen waarin meerdere winnaars worden verkozen. Toepassingen hiervan in de praktijk zijn onder meer het probleem waarbij parlementszetels moeten worden verdeeld over politieke partijen (het probleem van apportionment), het selecteren van een aantal kandidaten voor een sollicitatiegesprek, of een lokale gemeente die een selectie van openbare projecten kiest om uit te voeren. Dit is een gebied dat veel aandacht heeft gekregen van onderzoekers, en een groot deel van hun werk is gewijd aan het bestuderen van *eerlijke* aggregatiemethoden, waarbij eerlijkheid specifiek verwijst naar het begrip *evenredige vertegenwoordiging* (ook wel: *proportionele representatie*). In gevallen waar het doel is om meerdere winnaars te verkiezen vereist het ideaal van proportionaliteit dat we hanteren doorgaans dat een groep kiezers die zeer vergelijkbare voorkeuren heeft en een α-fractie van het electoraat vertegenwoordigt, controle moet hebben over de keuze van een α-fractie van de winnaars. Dit is een idee dat we zeer wenselijk achten en dat veel aandacht heeft gekregen in de literatuur over computationele socialekeuzetheorie. Dit proefschrift zal onderzoeken in hoeverre proportionele representatie kan worden ontwikkeld voor een aantal *complexe domeinen*.

Welke complexe domeinen beschouwen we precies? Neem als basis het reguliere geval van een verkiezing waarin we meerdere kandidaten voor een commissie moeten kiezen. Wat als de zetels in de commissie gekoppeld zijn aan een rol en sommige rollen waardevoller zijn dan andere? Of stel je voor dat er een restrictie is die stelt dat het opnemen van kandidaat A betekent dat kandidaat B niet in de commissie kan zitten. We wijzen een aantal complexe varianten van reguliere geval aan, zoals onder meer de bovenstaande. Binnen deze complexere gevallen achten we het wenselijk dat het begrip evenredige vertegenwoordiging op passende wijze wordt aangepast. Dit is dan ook de primaire focus van dit proefschrift. Het hoofddoel van het proefschrift is, concreter gezegd, om te bepalen hoe we in deze complexe domeinen uitkomsten kunnen produceren die proportioneel representatief zijn voor de deelnemende kiezers. Daarbij passen we proportionaliteitsbegrippen uit het reguliere geval aan, terwijl we rekening houden met de extra complexiteit die met deze domeinen gepaard gaat.

Het eerste deel van het proefschrift bestaat uit twee hoofdstukken. Het eerste hoofdstuk gaat over het probleem van apportionment, maar dan met de kiezers die sommige parlementszetels waardevoller vinden dan andere. Dit idee van zetels met verschillende waarden zien we terug in het tweede hoofdstuk, dat het geval behandelt waarbij meerdere kandidaten toegewezen worden aan zetels in een commissie, maar waarbij de commissiezetels niet gelijk worden behandeld. In beide hoofdstukken is het doel om proportionaliteit te importeren in de betreffende complexe domeinen.

Het volgende deel van het proefschrift bestaat uit één hoofdstuk en vertegenwoordigt een korte afwijking van het hoofddoel van het proefschrift. Hier wordt onderzocht in hoeverre aggregatiemethoden die worden gebruikt om meerdere winnaars te selecteren kunnen worden gesimuleerd in het algemene raamwerk van *judgment aggregation*. Deze analyse geeft ons inzicht in de interne werking van deze aggregatiemethoden.

Het hieropvolgende deel van het proefschrift omvat een terugkeer naar het project om proportionaliteit aan te passen aan complexe domeinen. Hoewel beide hoofdstukken van dit deel een ander specifiek interessegebied behandelen, hebben de hoofdstukken een gemene deler. In beide domeinen is er sprake van een restrictie die de mogelijke uitkomsten beperkt. Wat zijn deze twee domeinen nu precies? Eén domein betreft verkiezingen over een reeks publieke kwesties, waarbij er voor elke kwestie een keuze gemaakt wordt uit meerdere alternatieven. Het andere domein gaat over kiezers die samenwerken om gezamenlijk een shortlist van kandidaten te maken. Voor beide gevallen proberen we ervoor te zorgen dat er proportionele uitkomsten worden geproduceerd, terwijl de restricties worden gerespecteerd.

Ten slotte biedt het proefschrift een verkenning van aggregatiemethoden die erop gericht zijn om een groot aantal winnaars te selecteren. Deze verkenning wordt voornamelijk gedaan door te onderzoeken in hoeverre men proportionaliteitsbegrippen kan verheffen van de reguliere gevallen van verkiezingen met meerdere winnaars naar een aantal van hun complexere tegenhangers.

Summary

The vast array of collective decision-making scenarios that occur in the real world has provided a wealth of examples for *computational social choice* researchers to model and analyse. The field of computational social choice has taken the perspective of computer science in studying the methods used to aggregate group opinions towards producing a single collective outcome. The classical social-choice example sees an electorate voting over candidates in some election. Typically, the goal in these elections is to elect a single winning candidate such as a president. However, this thesis goes beyond the case of single-winner elections and focuses on scenarios where multiple candidates are to be selected as winners instead of just one candidate. Real-world applications of this include the apportionment problem where parliamentary seats are to be distributed to political parties, choosing a number of candidates to form a shortlist to attend a job interview, or a local municipality choosing a selection of public projects to implement. This is an area that has garnered plenty of attention from researchers and a significant portion of their work has been dedicated to studying *fair* aggregation methods, with fairness specifically referring to the notion of *proportional representation.* In these settings where the goal is to select multiple winners, the ideal of proportionality that we adopt typically requires that a group of voters that has very similar preferences and represents an α-fraction of the voting population, should have control over the selection of an α-fraction of the winners. This is a notion that we deem to be highly desirable and is one that has received considerable attention within the literature of computational social choice. The thesis will investigate the extent to which proportional representation can be developed for certain *complex domains.*

What are these complex domains? Take, as a foundation, the standard voting scenario where we are to elect multiple candidates to a committee. What if the seats on the committee are associated with a role and some of the roles are more valuable than others? Or consider that there is a constraint that states that including candidate A means that candidate B cannot be in the committee? We identify scenarios such as these, amongst others, as complex variants of the

standard setting where we are to choose more than one winner. Within these more complex settings, we deem it desirable that the notion of proportional representation be suitably adapted. This is then the primary focus of this thesis. Specifically, the thesis' main objective is to determine, when in these complex domains, how to produce outcomes that are proportionally representative of the participating voters. In doing so, we adapt proportionality notions from the standard settings while taking into account the added complexities that come with these domains.

The first part of the thesis consists of two chapters. The first of these chapters deals with the apportionment problem but with the voters considering some parliamentary seats to be more valuable than others. This notion of seats having varying values is then seen again in this part's second chapter where the problem is assigning multiple candidates to some seats in a committee but with committee seats not being treated equally. Within both chapters, the goal is to import proportionality into the complex domains of interest.

The thesis' next part consists of a single chapter and represents a brief deviation from the main objective of the thesis. Here, there is an investigation into the extent that aggregation methods that are used to select multiple winners, can be simulated in the general framework of *judgment aggregation*. This analysis provides us with insights into the inner workings of these aggregation methods.

The part of the thesis that follows sees a return to the task of adapting proportionality to complex domains. While each of this part's two chapters deals with a particular domain of interest, there is a common thread throughout this part. Specifically, for each domain, there is the presence of a constraint that restricts the outcomes that are feasible. Now, what exactly are these two domains? One scenario concerns a vote over a set of public issues with there being multiple alternatives that can be selected for each issue. The other scenario sees voters working to collectively create a shortlist of candidates. For both of these, we look to ensure proportional outcomes are returned while respecting a constraint.

In the end, the thesis provides an exploration of aggregation methods that aim to choose many winners. And for the most part, this exploration is done through investigating to what extent one can lift proportionality notions from standard multiwinner voting settings to some of their more complex counterparts.

Contents

Samenvatting **v**

Summary **vii**

Acknowledgments **xiii**

Setting the Stage

1 Introduction **3**
1.1 Formal Foundations . . . 6
1.2 Objectives and Thesis Outline . . . 9

2 Background **13**
2.1 Preliminaries . . . 13
2.2 Apportionment . . . 15
2.2.1 The Model . . . 15
2.2.2 Apportionment Axioms for Proportionality . . . 15
2.2.3 Apportionment Methods . . . 16
2.3 Approval-based Multiwinner Voting . . . 17
2.3.1 The Model . . . 17
2.3.2 Proportionality in Multiwinner Voting . . . 18
2.3.3 Proportional Rules . . . 21

3 Justified Representation in Complex Domains: A Glimpse at Existing Approaches **27**
3.1 Apportionment Extensions . . . 27
3.2 Multiwinner Voting Extensions . . . 29
3.3 Takeaways . . . 36

Part One: Some Seats Have More Value Than Others

4 Apportionment with Weighted Seats **39**
4.1 The Model 41
4.2 Weighted Lower Quota 45
4.3 Weighted Upper Quota 56
4.4 House Monotonicity for Weights 66
4.5 Experiments 69
4.6 Chapter Summary 75

5 Approval-based Multiwinner Voting with Weighted Seats **79**
5.1 The Model 82
5.2 Weighted Justified Representation 84
5.2.1 Up-to-one Weakening of WEJR 86
5.2.2 An Alternative Relaxation of WEJR 93
5.3 Priceability with Weights 95
5.4 Chapter Summary 102

Detour

6 Simulating Multiwinner Voting Rules (Beyond Approval Ballots) in Judgment Aggregation **107**
6.1 Multiwinner Voting Beyond Approvals 108
6.1.1 The Model 109
6.1.2 Ranking-based Multiwinner Rules 109
6.2 Judgment Aggregation 110
6.2.1 The Model 111
6.2.2 Majoritarian Rules 111
6.3 Simulating Multiwinner Voting Rules 112
6.3.1 Preference Agenda and Constraints 112
6.3.2 Extracting Election Winners 114
6.3.3 Simulation Results 115
6.4 Computational Considerations 126
6.5 Chapter Summary 130

Part Two: Constraining the Feasible Committees

7 Constrained Public Decisions **133**
7.1 The Model 136
7.2 Proportionality via Justified Representation 138
7.3 Justified Representation with Cohesiveness 141
7.4 Justified Representation with Agreement 150

7.5 Proportionality via Priceability 160
7.6 Chapter Summary 163

8 Constrained Shortlisting Using Approvals 165
8.1 The Model 166
8.2 Defining Proportionality Axioms 167
8.3 Rules to Provide Proportionality 169
8.4 Chapter Summary 176

Coming to a Close

9 Conclusion 181
9.1 Taking It All In 181
9.2 Taking It Further 183

Bibliography 187

Index 201

List of Symbols 203

Acknowledgments

Firstly, I'd like to thank Piotr Faliszewski, Rebecca Reiffenhäuser, Guido Schäfer, Ulrike Schmidt-Kraepelin and Yde Venema for being part of my thesis committee. I look forward to chatting about the thesis with you.

Now, I wish to thank my supervisors, Ronald and Ulle. I feel extremely fortunate to have had your support. Not only are you both highly impressive from an scientific standpoint but you are both absolute pleasures to be around. Thank you for all of the laughs and for putting up with me these last few years. To Ronald, I really appreciate your consistent availability especially during this recent period with the addition to your family. Your words of encouragement and sharp insights have been a massive reason as to why this thesis was completed. To Ulle, thank you for the constant guidance and for allowing me to be part of your exceptional research group in Amsterdam. It has shown me how things in academia should be done and is certainly something for me to aspire towards. It is a testament to your own character that through your group, I was introduced to some of the very best people that I've ever had the privilege of meeting.

Next, I have to send the warmest thanks to the many faces of the Amsterdam COMSOC group over the years: Zoi, Sirin, Arianna, Arthur, Simon, Adrian, Jan, Oliviero, Markus and Federico. When asked about my PhD experience, I always mention that it is the people around me that have been the best part. I arrived to Amsterdam to a group that made every effort to make me feel welcomed and I will forever be grateful for that. I do not imagine I would've made it through the early parts of this journey without you all. And in the years that followed, to all the people that came and went, I will always remember our time spent together (however short/long it may have been) and how you have contributed to my development not only as an academic, but as a person. The trips to the office were always easier knowing I'd run into you.

The COMSOC community as a whole is full of great individuals and I've been fortunate that I can mention a few as collaborators. To Jan and Adrian, our times working together in Amsterdam were some of the most enjoyable moments I've had in my so-far brief stint in academia. Your insights and enthusiasm inspired me more than you could ever imagine. Then beyond Amsterdam, to Arianna and Umberto, I say thank you for going through the trouble of hosting me for my research visits and thank you for our collaborative journey thus far. The opportunity to work with people like yourselves is the biggest reason that I wish to continue down the academic career path.

I must also mention my appreciation of the ILLC as an institute. The ILLC has provided me with an excellent environment to take on this PhD challenge that is even more daunting for those like me moving from across the world. I'm thankful for all of the support (administrative and otherwise) and for all of the people within the institute that I had the chance to spend time with.

Speaking of office life, I would be remiss not to highlight my two constant Lab42 office mates: Daira and Simon. I often looked forward to the conclusion of my PhD but most times, I dreaded the day where I would turn up to work and your faces would not be there to greet me. I could not have achieved this without your presence and seeing you both complete your own PhD journeys has filled me with so much joy. I'll always cherish your friendship and I'll be watching your future steps with much excitement. Thank you so very much.

As tough as doing a PhD has been, I cannot overstate the amount of fun that I've had while staying in Amsterdam. I put all of the credit on the shoulders of my friends. Thinking about each and every one of you, and all of the experiences we shared, makes it clear how lucky I've been. The board game nights, all of the chats about food, the murder mysteries, the concert and festival fun, the laughs on those terrace lunches, having a dance at the club, the (far too few) trips to the cinema, the lectures about food, the attempts at making art, the nights spent bouldering, coercing me onto boat rides, the raclette-filled Christmas dinners, the games of football (in Utrecht too!), the braais, the chaotic squash battles, all of the King's Day shenanigans and the (way too many) interrogations about my food taste. Because of you all, I end these four years with experiences more valuable than any diploma and I hope that I will always get to call all of you my friends.

I must also send love to all of my friends and family whether in South Africa, Zimbabwe or elsewhere in the world. It is you that shaped the person that arrived in Amsterdam all those years ago and without your support throughout these years, I could not have achieved this.

Finally, and most importantly, I have to say thank you to my Mother and to my Aunt Hesphina, for everything.

Setting the Stage

Chapter 1

Introduction

Decision-making is an integral feature of everyday life. People are constantly making choices amongst an (often substantial) offer of options. Consider a mundane morning routine where a person selects the outfit they will wear and the route they will take on their commute to work. On the individual level, *how* one makes such (seemingly inconsequential) choices may be of little interest. However, when we zoom out of the individual bubble, we see that *collective* decision-making is a cornerstone of society at large. On the one hand, the decisions made by a collective may be at a small scale such as a friend group choosing a restaurant for their weekly dinner, or a travel destination for their joint summer vacation. On the other hand, the stakes can be much greater and have an impact at a larger scale. The following is a notable example that immediately springs to mind: voters casting their ballots to choose a country's next president.

Collective choice often manifests itself as the aggregation of individual views over a set of alternatives with the goal of selecting a *single* winning alternative. However, the thesis focuses on scenarios where *multiple* winning alternatives are to be selected. In the real world, we are not short of such cases. From short-listing some candidates that will be up for some award, to the distribution of parliamentary seats to various political parties, or even to a friend group deciding the activities they will partake in during the weekend. It is clear that this too is a collective choice problem that is ubiquitous in society and can have wide-ranging effects. Observe that a feature of many such scenarios is that there is a limit on the number of alternatives that can 'win', i.e., be part of the final collective outcome. For example, scheduling conflicts and limited time available may restrict the friend group to selecting very few of the available activities; or when distributing parliamentary seats, the number of available seats means only a small fraction of the running political candidates can get a seat. So, tough choices must be made and thus, we require some way of guiding our choices.

As individuals, we do not (often) make decisions arbitrarily. Usually, whether we are cognisant of this or not, we apply some criteria that inform our choices. Ultimately, the goal is to make *good* decisions. For example, taking a cycle does indeed provide a shorter route to work but using a public transport system (that is efficient) may be a more relaxing experience. These are factors that the decision-maker may consider. Now of course, this presence of desiderata is not limited to decision-making at the individual level. On an intuitive level, there are many ideals of a collective decision that seem uncontroversial. Examples include the desire that every individual in the group be treated equally; or that every alternative has a chance to be a winner; or even requiring that, whatever the final choice may be, it must be a choice that is fair to the group. We highlight that last point. *Fairness*. This thesis takes the stance that fairness should be at the forefront of one's thinking when considering what makes a collective outcome a good one. But what is it that makes a collective decision fair? Consider a single-winner election where voters indicate a single candidate that they support. One might jump straight to the idea that a fair result has the most supported candidate as the winner. In the case of choosing multiple winners however, this *tyranny of the majority (or largest minority)* is often seen as highly undesirable from the context of fairness.

> Imagine a group of five friends are deciding on the ten movies that they will collectively watch. What if three of them are in complete agreement with each other while being in complete disagreement with the remaining two? If we are concerned with what is fair then it seems unreasonable that this majority get to choose all ten of the movies with the two dissenting friends having zero say on the outcome.

Given the above example, we must move away from the majoritarian notion of fairness when aiming to choose a set of multiple winners. Of course, there is still the issue of there being varying notions of fairness. In a voting scenario, one may desire that a diverse range of opinions are represented in the selection of winners. Or when dividing resources amongst individuals, one may wish that the chosen division does not cater to the tastes of only a select few and thus, lead to envy from others. So, what exactly is meant by "fairness" within this thesis? Moving forward, what seems most appealing is the notion of *proportional representation*, or *proportionality*. This notion is a common, real-world requirement. The most prominent example is that of parliamentary elections. Here, some parliamentary seats are to be distributed to political parties vying for power within the parliament and in many of the world's countries, the seat distributions must adhere to some notion of proportionality.

> Consider there being three parties—let's say Party A, Party B and Party C—with there being, let's say, 125 available seats. Suppose that voters in the electorate must submit a vote in support of a single party. Now, say that parties A, B and C receive 60%, 30% and 10% of the votes, respectively. First, it would seem that the most sensible solution would have 75 seats, which is exactly 60% of the seats, assigned to Party A. Now, the other parties' vote shares do not yield a seamless cut of the remaining 50 seats. So, to produce a proportional solution, we aim to assign to Party B, and to Party C, a number of seats that is as close as possible to a 30% share, and a 10% share, of the seats, respectively.

This example shows that proportionality is a reasonable requirement for this setting. However, it is a notion that finds itself useful in scenarios beyond that of parliamentary elections. Let us turn our attention to *committee elections* where, given some candidates, the goal is select a subset of the candidates to form a committee. Instances of committee elections are widely seen in practice: from shortlisting tasks where panel members vote to whittle down a larger list of candidates (for some award or to attend job interviews), to the selection of movies to be showcased on a television network's evening broadcast where the network users' viewing history serves as the voting data, and even to choosing the locations where certain public facilities are to be placed within a neighbourhood. In many such examples that the real world provides, it is key that the decisions that are produced are ones that are proportionally representative of the voting population. However, what should proportionality (roughly) look like in the context of committee elections? Despite not being as cut and dry (at a conceptual level) as ensuring proportionality within parliamentary elections, there does seem to be a solid intuition that presents itself.

> Alice, Bob, Carol and Dan comprise a panel that is tasked with shortlisting ten candidates. Suppose that Alice and Bob, who represent 50% of the voting population, agree on five candidates that should be part of the ten-person shortlist. In this scenario, it would be reasonable for Alice and Bob to demand control of 50% of the shortlist and include these five candidates.

So we can (somewhat) seamlessly take the notion of proportionality from the domain of parliamentary elections to the richer domain of committee elections. Now, what if we make the domains even more *complex*?

> The executives of a television network are to vote to decide on the three shows (out of eight options) that will be aired on their evening programme. Moreover, they must also decide on the timeslots that the three shows will be placed. Suppose the available timeslots are 17:00, 20:00 and 23:00 and from the network's viewing data, the executives know that the 20:00 primetime slot draws the most viewers (so is most valuable) while the late-night 23:00 slot is the least popular. Thus, some timeslots are more valuable than others and cannot be treated equally when trying to make a timeslot assignment that is proportional. If a large group of executives have very similar preferences, do we assign a show that they support to the most appealing timeslot at 20:00, or instead, do we assign multiple shows that they support to the two less valuable slots? What of smaller voting groups with more obscure preferences? What choice of overall assignment ensures that they are sufficiently represented?

It is simply not enough that the selection of movies represents the networks viewers in a proportional manner, but the assignment to timeslots must also respect the ideal of proportionality. Let us consider another scenario. Suppose that while conducting a shortlisting task, there are constraints placed on the candidates that may be shortlisted. An example constraint may state that shortlisting candidate A means you cannot shortlist candidate B. Or, the constraint could state that at least one candidate from amongst candidates A, B and C must be in the final shortlist. The presence of such constraints makes the shortlisting task much more complex. However, despite there being more moving parts to consider, it would still be sensible to desire a proportionally representative shortlist to be made. This sets us up to state the thesis' goal, at least informally. With this thesis, we tackle the following dilemma:

When dealing with complex scenarios, how does a decision-maker (i) define what proportionality should be, and (ii) design decision-making methods and/or tools that produce proportional outcomes?

Formal work done on proportionality, both in parliamentary elections and committee elections, form the foundational pieces of the thesis. Next, we discuss some of this work and specify how the thesis intends to build upon these foundations.

1.1 Formal Foundations

In order to outline the main goal of the thesis, it is necessary to establish the field of study that we work in.

How groups of individuals can make *good* decisions has been subject of ongoing study in the field of *social choice theory*. Generally, this field studies how to come to a collective choice via the *aggregation* of the individual opinions of the participants in the decision-making process (Arrow, 1951; Arrow et al., 2002, 2011; Brandt et al., 2016). This work has resulted in a slew of proposed *aggregation methods* and a similarly large number of *axioms* that are used to assess the quality of the choices made by the aggregation methods. Ideally, we wish to use aggregation methods that produce outcomes that meet all of these desirable notions, but unfortunately, as negatively shown by the social choice literature's classical results, we cannot have it all (Arrow et al., 2002). So, whichever aggregation method is used will represent a compromise of some sort. Using the axiomatic approach provides an established way to make such a compromise and we will employ it in this thesis. Within this broad area of social choice theory, there have been contributions made by researchers from a variety of disciplines. This includes the fields of economics, political science and mathematics. The subarea that this thesis' content lies in, however, primarily takes a computer-science perspective of social choice.

Computational Social Choice. The burgeoning field of *computational social choice* (Brandt et al., 2016) has seen significant strides made in research over the last two decades. We are in a world where there is a push for digital platforms being used in democratic processes, the increasing presence of artificial agents plus more and more societal interactions occurring online. With this mind, it is vital to ensure that various aspects meet certain standards such as transparency, legitimacy and most relevant to us, fairness. We believe that computational social choice offers a rich pool of methods and tools to tackle the collective decision-making components within this increasingly digitised world.

A wide variety of formal frameworks have been proposed in the literature to analyse scenarios where several voters report an individual view and a reasonable compromise must be made to produce a single collective view. Examples of collective decision-making that have been formally studied include: voting (Zwicker, 2016); the division of some resources between the agents (Thomson, 2016); the matching of students to schools, or of kidney donors to transplant patients (Klaus et al., 2016); a local municipality funding certain public projects based on the preferences of their jurisdiction's residents (Rey and Maly, 2023); and judgment aggregation in a court of law (Kornhauser and Sager, 1983). But also many problems long studied in AI—such as belief merging (Doyle and Wellman, 1991), collective argumentation (Bodanza et al., 2017), and consensus clustering (Goder and Filkov, 2008)—can be seen in this vein (Endriss, 2020).

Thus, with this rich bed of literature, we have the foundations to design models that capture the real-world scenarios that we are interested in. Moreover, as proportionality is our desired aim within said scenarios, we are even more fortunate that there has already been extensive research on providing proportional representation within the computational social choice literature.

Apportionment. The *apportionment* task can be described as the allocation of resources in a proportional manner to entities with different entitlements. This task is one of the core problems of social choice (Balinski, 2005): in federal systems (e.g., the US), states receive seats in parliament according to their populations, while in proportional representation systems (e.g., the Netherlands), political parties receive seats according to their share in the popular vote. Outside of parliamentary elections, the need for apportionment arises in the context of fair allocation (Chakraborty et al., 2021), the presentation of statistics (Balinski and Rachev, 1993) and the handling of bankruptcies (Csóka and Herings, 2018), just to mention a few other applications. But, in line with the paradigmatic example of apportionment, we stick to the terminology of seats being assigned to parties for the rest of the thesis. Within the literature, providing proportional representation has been heavily studied in the case of parliamentary elections (Balinski, 1982; Pukelsheim, 2014).

Approval-based Multiwinner Voting. *Multiwinner voting* deals with the task of selecting some subset of candidates, typically of a fixed size k, from a set of m total candidates based on the preferences of the voters (Faliszewski et al., 2017; Lackner and Skowron, 2023). Much of the extensive research dedicated to this model has dealt with the case where voters submit preferences via *approval ballots* that indicate which subset of the candidates they approve of (Lackner and Skowron, 2023). This is called *approval-based multiwinner voting* but moving forward, we often simply refer to this as 'multiwinner voting'. The attention that the approval-based multiwinner voting model has received is unsurprising given the wide real-world applicability of the model (Lackner and Skowron, 2023). From electing a representative committee of political candidates to producing the shortlists of candidates for an award. From finding group recommendations (Lu and Boutilier, 2011, 2015; Skowron et al., 2016) to returning search results (Skowron et al., 2017). The practical applications that are present for this model are plentiful. In fact, this model can also be seen as a generalisation of the apportionment model (Brill et al., 2017) and this even further justifies our focus on it. This committee election scenario can be seen as a generalisation of the apportionment task where voters voice their support for individual political candidates from different parties instead of putting their full voting weight behind a single party. Much like the apportionment model, approval-based multiwinner voting has also seen significant research conducted on proportionality (Lackner and Skowron, 2023).

Complex Domains in Social Choice. In the social choice literature, it is common practice to take some well-established model and consider the richer, more complex variants of it. A notable example is the *fair division* problem (Procaccia, 2016; Thomson, 2016) where the task is to divide some items—that may be *divisible* (Procaccia, 2016) or *indivisible* (Bouveret et al., 2016)—amongst a group of individuals who each hold preferences on the items. This task has seen many variants being studied. Some examples include: fair division under constraints; dividing not only items with positive utility for their recipients, but also items that yield negative utility; or where items are presented, and must be allocated, in an online manner (Amanatidis et al., 2022).

With this line of thinking, the standard fair division setting was brought closer to real-world cases and allowed for the modelling of a larger number of interesting practical scenarios. The same can be done for the aforementioned models of apportionment and multiwinner voting. Let us hone in on the latter case. This multiwinner voting model can be generalised to become an instance of *participatory budgeting* (Rey and Maly, 2023). This can be considered a committee election instance where candidates, now called projects, come with an associated cost and instead of electing a fixed number of projects as an outcome, the task is to select a bundle of projects subject to a budget limit. In recent years, the study of the participatory budgeting setting has trended upwards and notably, much of this study has been dedicated towards how to produce proportionally representative outcomes for participatory budgeting (discussed in greater detail in Chapter 3). Now, although the multiwinner voting model can be used to conduct a participatory budgeting process, the richer model is a clearly better fit. It captures more information and is much closer to reality at a conceptual level. Also, just as importantly, it allows for the consideration of proportionality properties that are more natural to the setting itself. For example, it seems natural to require (in a loose sense) that a group of like-minded voters that are an α fraction of the voter population should control an α fraction of the municipality's budget. Using the richer, more general model makes this possible. So when moving up levels of generality—from parliamentary elections to committee elections, and then from committee elections to participatory budgeting—there is evidence that investigating proportional representation is a worthwhile venture.

1.2 Objectives and Thesis Outline

Of course the study of proportionality has not been limited to the above domains as we will observe later on in the thesis. Indeed, if one looks at the vast pool of (complicated) real-world scenarios, there are many more models that deserve, and have yet, to be investigated with the view of ensuring proportional representation. Consider the following two examples (that hint at the type of domains that we will study later in the thesis). Suppose there is a committee election where the seats in the committee are associated with certain roles. For example, one candidate must chair the committee while another is the committee's treasurer. In such a case, the role of the chair may be considered more valuable than that of the treasurer and thus, when looking to return a proportional committee, the seats cannot be treated equally as some seats yield more value for the voters than others. More broadly, consider the example of a friend group selecting their weekend activities.

In its most plain variant, this is a simple shortlisting task where the group is free to select whichever activities that wish to do. However, it is natural that there will be time constraints, conflicting schedules of activities, a limited budget to pay for the selected activities' associated costs, etc. This then takes the simple shortlisting task and constrains the combinations of activities that are feasible. Despite such a scenario's convoluted nature, the goal of returning a proportionally fair outcome remains at the forefront of our thinking. In this thesis, we look to identify a handful of these domains and conduct this proportionality study.

Goal of the Thesis. Now, we can state our aim for the chapters to come. Besides identifying complex domains that are relevant and natural, the primary aim of thesis can be summed up as providing an answer to the following question:

Can we lift notions of proportional representation from our chosen foundational settings (of apportionment and approval-based multiwinner voting) to more complex domains that are enriched versions of these foundations?

This goal has been stated as a 'yes or no' question, but in reality the answer will be 'yes, to some extent' with the extent to which we can do so depending on the proportionality notion/s and complex domain being considered. And looking ahead, it is clear that there are some complex domains that are much more difficult to transfer *certain* proportionality notions to than other domains.

We continue this section by providing a breakdown of the thesis' structure as well as the contributions within each of the thesis' chapters.

Setting the Stage

If you are reading this then you are well within the portion of the thesis that is meant to (*i*) prepare the reader (both in a technical and conceptual sense) for the thesis' later parts where the main objective is tackled, and (*ii*) situate the thesis' contributions amongst ongoing research holding similar objectives. Much of the motivation thus far is also stated in (Chingoma, 2023). The remainder of this part sees two chapters follow this introduction. In Chapter 2, the most pertinent background knowledge required for the thesis is presented. This is done in two parts with the background regarding proportionality given for both apportionment and approval-based multiwinner voting. This beginning part of the thesis is then closed with Chapter 3 which is titled *Justified Representation in Complex Domains: A Glimpse at Existing Approaches*. This chapter can be treated as a sort of 'mini-survey' that highlights some cases within the literature where proportionality is studied in a complex domain. Thus, Chapter 3 gives a look at work that is related to ours and gives a sense of the type of content to follow in the thesis' subsequent parts.

Part One: Some Seats Have More Value Than Others

Consider the tasks of assigning seats within a parliament and a committee. Most work done on proportionality has considered these tasks with the seats being of equal value to the voters. This part of the thesis takes a different approach and investigates what can be achieved in terms of proportional representation when some seats are considered to be more valuable than others. This is captured by having each seat associated with some weight that represents the seat's intrinsic value, e.g, a seat with a weight of 2 is considered more valuable than a seat that has weight 1. This is applied to two models and each application represents a chapter. Chapter 4 is the first of these chapters and it is titled *Apportionment with Weighted Seats*. This chapter's content is primarily based on (Chingoma et al., 2024a). Here, we take the apportionment model, enrich it by supposing that seats have weights, and then for most of the chapter, we adapt, from the standard apportionment literature, the most prominent proportionality axioms and apportionment methods to this weighted model. This part's second chapter, Chapter 5, is titled *Approval-based Multiwinner Voting with Weighted Seats*. This chapter extends the work done in Chapter 4 to the standard approval-based multiwinner voting model. In doing so, we lift proportionality notions from multiwinner voting to the case where the committee seats have varying weights.

Detour

This part of the thesis is comprised of a single chapter. The sole chapter of this part is Chapter 6 and it is titled *Simulating Multiwinner Voting Rules (Beyond Approval Ballots) in Judgment Aggregation*. Its content is based on some of the work from (Chingoma et al., 2022). Specifically, it is based on the paper's work that uses a general social-choice framework, called *judgment aggregation*. In this chapter, we use this judgment aggregation framework to simulate rules for a multiwinner voting model that is not limited to approval ballots. This means that with this chapter, we are deviating from the proposed task regarding proportionality in complex domains, hence referring to this part of the thesis as a detour. We believe that the inclusion of this chapter, and this change in thematic direction, is beneficial as the chapter touches on the structure of different multiwinner voting rules and ultimately, it provides a better understanding of how rules with certain properties differ from each other.

Part Two: Constraining the Feasible Committees

This part puts the thesis back on track with respect to the outlined main objective: proportionality in complex domains. It consists of two chapters which each following a similar route of adapting standard proportionality notions to the respective complex settings. In this part, the settings are made more complex with the introduction of constraints that restrict the collective outcomes that are feasible. We first encounter Chapter 7 that is titled *Constrained Public Decisions* and is based on (Chingoma et al., 2024b). This chapter looks at proportionality in a model of public decisions that is modified with the addition of a constraint on the possible outcomes. An example of the public decisions model (without constraints) would see voters deciding over a set of issues such as "do you agree with a particular political policy?" and the voters respond with one of the following options: "strongly agree", "agree", "indifferent", "disagree", or "strongly disagree". The part's second chapter, Chapter 8, is titled *Constrained Shortlisting Using Approvals* and it is mostly based on the proportionality parts from (Chingoma et al., 2022). Here, proportional representation is considered in a model of approval-based shortlisting where the shortlists that are feasible are determined by some constraint. In this case, the standard shortlisting model is one where voters submit approval ballots over a list of candidates with the goal of creating a shortlist of the candidates with this shortlist not having a predetermined size.

Coming to a Close

With Chapter 9, the thesis closes with a brief summary of the work that was covered as well as a brief discussion of promising routes for future research that expand on some of the thesis' content.

Chapter 2

Background

The two models of *apportionment* and *approval-based multiwinner voting* form the foundation for the work conducted in the thesis. This chapter serves to present some necessary background material related to proportionality in the contexts of these two settings. Specifically, this chapter details various notions, axioms and voting rules that will pop up frequently throughout the remainder of the thesis.

Take note that the presented work is only a selection of the most relevant information. For those interested in more detailed explorations on these topics, the following books are recommended readings:

- *Apportionment:* see the books by Balinski (1982) and Pukelsheim (2014).
- *Approval-based multiwinner voting:* see the book by Lackner and Skowron (2023) that also serves as the main source of this chapter's content for both apportionment and approval-based multiwinner voting.

2.1 Preliminaries

Before diving into the details, we first define some basic notation that will be seen often in this thesis. For any given set U, we use $\mathcal{P}(U)$ to represent its powerset, $\mathcal{P}_+(U)$ the set of all of its nonempty subsets, and for some positive integer k, we write $\mathcal{P}_k(U)$ to represent the set of all of its subsets of size k. Additionally, we sometimes write $[k]$ for the set $\{1, \ldots, k\}$, where k is a positive integer.

Basics of Computational Complexity. As mentioned in the introduction, a factor that we give serious consideration is the applicability of our results in practice. Consequently, we must consider whether our decision-making methods can produce collective outcomes in an efficient manner from a computational standpoint. So we give a brief word on some notions of *computational complexity* that we assume basic familiarity with. See (Arora and Barak, 2009) for an extensive coverage of the subject.

We begin by touching on our use of the *Big O notation*. For some functions f, g, we say that $f(n) \in O(g(n))$ if there exists some real numbers c and d such that $f(n) \leqslant c \cdot g(n)$ for all $n \geqslant d$. Also, we assume that *polynomial-time* computation is a well understood notion.

Now, here are the main complexity classes of decision problems that shall be mentioned: P, NP and coNP. Both *membership* within, and *hardness* for, these classes will be mentioned. Regarding the latter, we then assume at least basic knowledge of *polynomial-time reductions* from some complexity class. It is then assumed that the reader is also familiar with the notion of a decision problem being *complete* for some class in $\{\mathsf{NP}, \mathsf{coNP}\}$.

Regarding the class NP, we (at times) distinguish between those decision problems that are *weakly* NP-hard and those are *strongly* NP-hard. The former are the decision problems that are polynomial-time solvable when all numbers in the input are represented in *unary* while the latter are those that do not admit a polynomial-time algorithm to solve said decision problems (assuming $\mathsf{P} \neq \mathsf{NP}$) even when all the numbers in the input have a unary representation.

Environments. To further aid the reader, we highlight some non-standard features in the environments that appear throughout the thesis.

- *Examples:* these environments have a light orange background with an orange leftbar of a darker shade.

Example 2.1. Here is an example.

- *Remarks:* these environments have italicised text and the end of a remark is marked with a ◁ symbol.

Remark 2.1. *Here is a remark.* ◁

- *Proofs:* these environments have a light grey background with a grey leftbar of a darker shade. The end of the proofs are marked with a □ symbol (as is customary).

Proof. Here is a proof. □

2.2 Apportionment

As stated in the introduction, apportionment is concerned with the distribution of parliamentary seats to political parties in a manner that is proportionally representative of the electorate. Moving forward, we refer to this setting/model that we present as the *standard* apportionment setting/model.

In this section, we detail this standard apportionment model along with its prominent proportionality axioms and some of the well-known methods used to assign seats to parties.

2.2.1 The Model

Take a finite set of *voters* $N = \{1, \ldots, n\}$ who shall each vote for one of the m *parties* $\{1, \ldots, m\}$. A *vote vector* $\boldsymbol{v} = (v_1, \ldots, v_m) \in [n]^m$, with $v_1 + \cdots + v_m = n$, specifies how many votes (out of the total number n) that party $p \in [m]$ garnered. An *election instance* is a pair $(\boldsymbol{v}, k)$ consisting of a vote vector $\boldsymbol{v}$ and the number of seats k.

In line with the work in the literature, we say there is *full supply* if each party has at least k members, and is thus able to fill all of the available seats by itself (we borrow the full supply term from Lang and Skowron (2018)). Unless stated otherwise, we assume full supply to hold throughout the thesis.

Recall that the apportionment goal is to divide k *seats* over some parties. So, an outcome in this setting is a *seat assignment*, which is a vector $\boldsymbol{s} = (s_1, \ldots, s_k) \in [m]^k$, where $s_t = p$ means that party $p \in [m]$ is assigned seat $t \in [k]$. Given a seat assignment $\boldsymbol{s} = (s_1, \ldots, s_k)$, we can talk of the *representation* of a party p, denoted by $r(p, \boldsymbol{s}) = |\{s_t \in \boldsymbol{s} \mid s_t = p\}|$, which indicates how many seats were assigned to party p within the seat assignment $\boldsymbol{s}$.

An *apportionment method* M takes an election instance $(\boldsymbol{v}, k)$ as input and maps it to a winning seat assignment $M(\boldsymbol{v}, k)$ (so we assume that these apportionment methods are *resolute*).[1]

2.2.2 Apportionment Axioms for Proportionality

A primary concern of apportionment is how to assign the available seats to parties in a *proportional* manner. The notion of the *quota of a party* then becomes central to the apportionment task. The quota of a party p is $q(p) = k \cdot {}^{v_p}/_{n}$. Intuitively, this quota for a party is the share of the k seats that this party deserves. Ideally, a perfectly proportional assignment would assign to each party p, a number of the k seats that corresponds *precisely* to this share defined by the quota. However, this quota may not be integral. The immediate fallback is to 'approximate' the quota

[1] When we say that a method (or later on, a voting rule) is resolute, this means that this method (or voting rule) returns only a single winning outcome as opposed to *irresolute* methods (or voting rules) that may return a set of winning outcomes.

for a party. In the apportionment literature, this is formalised in the following two ways:

Definition 2.1 (Lower Quota, LQ)**.** *An apportionment method M satisfies* lower quota (LQ)*, if for every election instance $(\boldsymbol{v}, k)$, it holds for the resulting seat assignment $\boldsymbol{s} = M(\boldsymbol{v}, k)$ that there exists no party p such that $r(p, \boldsymbol{s}) < \lfloor k \cdot {}^{v_p}/_{n} \rfloor$.*

Definition 2.2 (Upper Quota, UQ)**.** *An apportionment method M satisfies* upper quota (UQ)*, if for every election instance $(\boldsymbol{v}, k)$, it holds for the resulting seat assignment $\boldsymbol{s} = M(\boldsymbol{v}, k)$ that there exists no party p such that $r(p, \boldsymbol{s}) > \lceil k \cdot {}^{v_p}/_{n} \rceil$.*

LQ and UQ are two of the most extensively studied axioms in the apportionment literature (Balinski, 1982; Pukelsheim, 2014). This prompts us to choose these two axioms as our focus points when studying apportionment.

2.2.3 Apportionment Methods

We now detail some of the more important apportionment methods that have been defined, starting with the class of *divisor methods*.

Definition 2.3 (Divisor method)**.** *Given an election instance $(\boldsymbol{v}, k)$ and a function $f : \mathbb{R} \to \mathbb{R}$, the divisor method for f works in k rounds, as follows. In each round $t \in [k]$, seat t is assigned to the party p with the highest value for:*

$$ratio_p = \begin{cases} \frac{v_p}{f(g_p(t))} & \text{if } f(g_p(t)) \neq 0 \\ \infty & \text{if } f(g_p(t)) = 0, \end{cases}$$

where $g_p(t)$ is the number of the seats assigned to party p in earlier rounds. If required, a tie-breaking rule is used to choose between parties with equal ratio.

Within this class of divisor methods, two stand out as particularly appealing.

Definition 2.4 (Adams method)**.** *The Adams method is the divisor method defined by $f(g_p(t)) = g_p(t)$.*

Definition 2.5 (D'Hondt method)**.** *The D'Hondt method is the divisor method defined by $f(g_p(t)) = g_p(t) + 1$.*[2]

There exist other, notable divisor methods such as the *Saint-Laguë* method which is the divisor method with $g_p(t) + 0.5$. However, amongst the divisor methods, we narrow our focus on the Adams and D'Hondt methods since Adams is the unique divisor method to satisfy UQ while D'Hondt is the unique divisor method to satisfy LQ. So note that Adams fails LQ and D'Hondt fails UQ.

We now detail an apportionment method that has drawn substantial attention from researchers and sits outside the class of divisor methods.

[2]In the literature, the D'Hondt method is also known as the *Jefferson* method (Balinski, 1982; Pukelsheim, 2014).

Definition 2.6 (Largest Remainder Method, LRM)**.** *The LRM method first assigns each party p with exactly* $\lfloor k \cdot {}^{v_p}\!/\!_n \rfloor$ *seats. Then to assign the remaining r seats, the method assigns exactly one seat to each of the r parties with the largest value for* $k \cdot {}^{v_p}\!/\!_n - \lfloor k \cdot {}^{v_p}\!/\!_n \rfloor$.[3]

To conclude this section, we highlight that LRM satisfies both LQ and UQ, and so it outperforms the two chosen divisor methods in this respect. However, it is trumped by the divisor methods on other properties (such as those related to *monotonicity* (Balinski, 1982; Pukelsheim, 2014)) that are not detailed here, which motivates the widespread study and use of the divisor methods alongside LRM. Again, see (Balinski, 1982) and (Pukelsheim, 2014) for more extensive discussions and analyses of these apportionment methods along with many others.

2.3 Approval-based Multiwinner Voting

In this section, we present the background related to the approval-based multiwinner voting (MWV) model using approval ballots, and in particular, we detail some axioms and voting rules related to proportional representation.[4] Throughout the thesis, we will refer to this as the *standard* multiwinner voting setting/model.

2.3.1 The Model

First, we provide the model that again contains a set of voters $N = \{1, \ldots, n\}$. We then have a set of m *candidates* $C = \{a, b, c, \ldots\}$ and each voter $i \in N$ submits an *approval ballot* $A_i \subseteq C$ that indicates the set of candidates that voter i approves of. This is also referred to as a *dichotomous* preference order. For a candidate $c \in C$, we denote the set of voters that approve of c as $N(c) = \{i \in N \mid c \in A_i\}$. An *approval profile* is then a vector $\boldsymbol{A} = (A_1, \ldots, A_n)$ of approval ballots, one for each voter. An *election instance* is then a pair $(\boldsymbol{A}, k)$.

The goal is to select, as an outcome, a set of k candidates $W \in \mathcal{P}_k(C)$ that we refer to as a *committee*. We define a voter i's *satisfaction* with a committee W to be $|A_i \cap W|$, i.e., the number of committee members approved by voter i.

Note that we can model the apportionment setting in this approval-based multiwinner voting model. This can be done using *party-list profiles* where for all $i, j \in N$, it holds that either $A_i = A_j$, or $A_i \cap A_j = \emptyset$. Additionally, the property of full supply can be ensured by only considering profiles where $|A_i| \geqslant k$ for every voter $i \in N$.

[3]As is the case with D'Hondt, LRM is also known by a different name within the literature, namely as the *Hamilton method* (Balinski, 1982).

[4]For more of approval-based multiwinner voting, Lackner and Skowron (2023) provide an overview of recent advances on this topic.

2.3.2 Proportionality in Multiwinner Voting

This section details the proportionality axioms from the multiwinner voting literature that shall be referred to often in the thesis.

The Core and Justified Representation

For us to label an outcome as proportionally representative, we deem that it must meet the following ideal: if some group of voters represents an α-fraction of the voter population, then this voter group should be able to control, roughly speaking, an α-fraction of the outcome. The outcome in the approval-based multiwinner voting context is the selection of candidates that constitutes the committee. The following axiom—that is an adaptation by Aziz et al. (2017a) of the well-known, game-theoretic core notion—captures this very ideal in this multiwinner setting.

Definition 2.7 (Core)**.** *Given an election instance* $(\boldsymbol{A}, k)$*, we say that a committee* W *is in the core if for every group of voters* $N' \subseteq N$ *and every set of candidates* $C' \subseteq C$ *such that*

$$|N'| \geqslant |C'| \cdot n/k,$$

there exists a voter $i \in N'$ *such that* $|A_i \cap C'| \leqslant |A_i \cap W|$*.*

It is an open problem whether the core is non-empty for all election instances. This represents a drawback for the core as an option to lift to richer settings. So, moving forward, we focus on properties that are known to always be satisfiable. To define these axioms, we must define the notion of a *cohesive group*.

Definition 2.8 (ℓ-cohesiveness)**.** *For an integer* $\ell \in \{1, \ldots, k\}$*, we say that group of voters* $N' \subseteq N$ *is* ℓ*-cohesive if both of the following conditions hold:*

- $|N'| \geqslant \ell \cdot n/k$.
- $|\bigcap_{i \in N'} A_i| \geqslant \ell$.

Using this cohesiveness notion, we can define the most central axiom of this thesis called *extended justified representation (EJR)* (Aziz et al., 2017a).

Definition 2.9 (Extended Justified Representation, EJR)**.** *Given an election instance* $(\boldsymbol{A}, k)$*, we say a committee* W *provides* extended justified representation (EJR) *if for every* ℓ*-cohesive group of voters* $N' \subseteq N$*, there exists a voter* $i \in N'$ *such that* $|A_i \cap W| \geqslant \ell$*.*

EJR can be seen as a restriction of the core by imposing a cohesiveness requirement on those voter groups that would deviate (Peters and Skowron, 2020). Thus, identifying those groups that are cohesive becomes of practical importance. Unfortunately, the problem of deciding whether there exists an ℓ-cohesive group of voters is NP-complete (Skowron et al., 2017). Consequently, checking whether a given committee W provides EJR is coNP-complete (Aziz et al., 2017a).

We move on to the following weakening of EJR that was defined by Sánchez-Fernández et al. (2017). This axiom deems a cohesive group N' to be represented if the group's total satisfaction, summing over the voters in N', is at least ℓ, instead requiring some voter from N' to reach this satisfaction threshold of ℓ.

Definition 2.10 (Proportional Justified Representation, PJR)**.** *Given some election instance* $(\boldsymbol{A}, k)$*, we say a committee* W *provides* proportional justified representation (PJR) *if for every* ℓ*-cohesive group of voters* $N' \subseteq N$*, it holds that* $|(\bigcup_{i \in N'} A_i) \cap W| \geqslant \ell$*.*

From a computational perspective, it is also coNP-complete to check if a committee W provides PJR (Aziz et al., 2018).

Next is the weakest of the justified-representation axioms that we consider and we see that it coincides with EJR when one only considers ℓ-cohesive groups for $\ell = 1$ (Aziz et al., 2017a).

Definition 2.11 (Justified Representation, JR)**.** *For an election instance* $(\boldsymbol{A}, k)$*, we say a committee* W *provides* justified representation (JR) *if for every* ℓ*-cohesive group of voters* $N' \subseteq N$ *for* $\ell = 1$*, there exists a voter* $i \in N'$ *such that* $|A_i \cap W| \geqslant 1$*.*

Now with this axiom, we find more positive news with regards to computational complexity as it is known that checking if a committee W provides JR can be done in polynomial time (Lackner and Skowron, 2023).

Of the known notions of justified representation, those mentioned thus far are those that will feature significantly we in the remainder of the thesis. However, we would be remiss not to mention (at least briefly) some of the following justified-representation-like notions in the literature. The axioms of *fully justified representation* (Peters et al., 2021b) and *EJR*$^{+}$ (Brill and Peters, 2023) are recently introduced strengthenings of EJR that have been shown to always be satisfiable. Another axiom that implies EJR was defined by Brill et al. (2022) as they employed the justified-representation notion towards the task of representing *individual* voters instead of voter groups. And also, the notion of *proportionality degree* (or *average satisfaction of cohesive groups* provides a quantitative measure of the satisfaction afforded to cohesive groups (Sánchez-Fernández et al., 2017; Skowron et al., 2017; Skowron, 2021).

Priceable Committees

We now pivot from justified-representation to a market-based notion of proportionality called *priceability* (Peters and Skowron, 2020; Lackner and Skowron, 2023). Loosely speaking, priceability deems a committee W to be proportional only if, in a simulated market where voters have access to the same amount of virtual funds, the voters that support the committee members are able to fund these candidates' presence in the committee W (subject to further conditions as outlined in the following definition).

Definition 2.12 (Priceability)**.** *Suppose that voters have a personal budget* b*. A* price system $\boldsymbol{ps} = (b, \{p_i\}_{i\in[n]})$ *is a pair where* b *is the individual* budget *per voter and each voter* $i \in N$ *has a* payment function $p_i : C \to [0, b]$ *such that:*

C1. *If* $p_i(c) > 0$*, then* $c \in A_i$ *(a voter only pays for candidates she approves of).*

C2. $\sum_{c\in C} p_i(c) \leqslant b$ *(a voter does not spend more than her budget).*

A price system supports a committee W *if all of the following hold:*

C3. *For every* $c \in W$, $\sum_{i\in N} p_i(c) = 1$ *(payments for this candidate equal to the price of* 1*).*

C4. $\sum_{i\in N} p_i(c) = 0$ *for every* $c \notin W$ *(candidates outside the committee are not paid for).*

C5. *For every* $c \notin W$*, it holds that:*

$$\sum_{i\in N(c)} \left(b - \sum_{c'\in W} p_i(c') \right) \leqslant 1$$

(supporters of any unelected candidate do not collectively hold, in terms of their unspent budget, strictly more funds than the price of a candidate).

If there exists a price system $\mathbf{ps} = (b, \{p_i\}_{i\in[n]})$ that supports a committee W then we say a committee W is *priceable*. And if a rule F_α always returns committees that are priceable, then we say that F_α is priceable.

Now, we know that priceable committees always exist (Peters and Skowron, 2020; Peters et al., 2021a). Also, for a given committee W, we have that checking whether W is priceable can be done in polynomial time by making use of an *Integer Linear Program (ILP)* (Peters et al., 2021a).

As priceability is not the strongest requirement, we note that there exist other, stronger notions of priceability such as *Stable Priceability* and *Balanced Stable Priceability* (Lackner and Skowron, 2023; Peters et al., 2021a). We do not focus on these as it is known that there are election instances where no committees exist that satisfy stable priceability, or balanced stable priceability (Peters et al., 2021a). However, regarding stable priceable committees, we note that it is easy to check if such a committee exists for a given election instance, and whenever a stable priceable committee exists, it is in the core (Peters et al., 2021a).

Now, we note a connection between priceability and an axiom of justified representation. Specifically, a priceable committee W with size $|W| = k$ also provides PJR (Peters and Skowron, 2020). However, this does not hold for EJR as it turns out to be incompatible with priceability (Peters and Skowron, 2020).

2.3.3 Proportional Rules

In this section, we outline the multiwinner voting rules that play a prominent in the rest of this thesis (see (Lackner and Skowron, 2023) for a more comprehensive look at these rules as well as others that we do not mention). Formally, an *approval-based multiwinner voting rule* is a function $F_\alpha : (\mathcal{P}_+(C))^n \to \mathcal{P}_+(\mathcal{P}_k(C))$ mapping an approval profile to a set of winning committees that are each of size k. Note that frequently, for purposes of readability, we simply refer to these as multiwinner voting rules (unless the context requires clarification on the use of ballots that are not approval-based). Ideally, there will be a single winning committee, but in general F_α may be irresolute. So in practice, a method to break ties may have to be used post-election to produce a single, winning committee.

To start with, we define the *Greedy Cohesive Rule (GCR)* (Peters et al., 2021b) that was designed for the purpose of providing EJR. It is for this reason that it may not be considered the most natural rule but it is useful showing that one can always achieve the provision of EJR.

Definition 2.13 (Greedy Cohesive Rule, GCR)**.** *The rule starts with an initially empty, current committee W and all voters being active. In iterations, the rule looks for the largest ℓ-cohesive group of active voters $N' \subseteq N$ and selects ℓ candidates that are approved by all voters in N' to be part of the current committee. The rule then makes the voters in N' inactive and continues to the next iteration, unless there are either (i) no more active voters, or (ii) there are no more available seats on the current committee W, i.e., $|W| = k$. In case either (i) or (ii) occurs, the rule terminates.*

Now, regarding the proportionality axioms, GCR satisfies EJR, but it does not always return priceable outcomes. Also note that a committee W returned by GCR may be *partial*, i.e., $|W| < k$. Thus, in practice, one must consider a method of completing a GCR committee to be of size k (Peters and Skowron, 2020).

We continue by presenting a class of rules introduced by Thiele (1895) that is home to two rules of particular interest to us.

Definition 2.14 (Thiele Methods)**.** *Given an election instance $(\boldsymbol{A}, k)$, the* Thiele method F^u_α *that is induced by some scoring vector* $\boldsymbol{u}^{(k)} = (u_1, \dots, u_k)$ *is the following:*

$$F^u_\alpha(\boldsymbol{A}) = \operatorname*{argmax}_{W \in \mathcal{P}_k(C)} \sum_{i \in N} \sum_{t=1}^{|W \cap A_i|} u_t.$$

This class captures the standard *Approval Voting (AV)* rule which is the Thiele method with scoring vector $\boldsymbol{u} = (1, 1, \dots, 1)$. However, AV is not priceable and it does not even satisfy JR (Lackner and Skowron, 2023). So, AV is not one that is considered to be 'fair' or 'proportional' in any sense, unlike the following two Thiele methods that we emphasise.

Definition 2.15 (Proportional Approval Voting, PAV)**.** *The rule called* Proportional Approval Voting (PAV) *rule is the Thiele method with the* harmonic scoring vector $\boldsymbol{u}^{(k)} = (1, 1/2, 1/3, \dots, 1/k)$.

It is known that PAV satisfies EJR (Aziz et al., 2017a) and this places it as a rule to consider when dealing with proportional representation. However, it is NP-hard to compute the outcomes of PAV (Aziz et al., 2018). This leads us to consider the following local-search variant of PAV that represents one of the select few multiwinner voting rules that satisfies EJR while also being polynomial-time computable (Aziz et al., 2018). To define this rule we must define the following:

Definition 2.16 (PAV score of a committee)**.** *Given an election instance* $(\boldsymbol{A}, k)$*, the* PAV-score *of a committee* $W \subseteq C$ *is defined as:*

$$\textit{PAV-score}(W) = \sum_{i \in N} \sum_{t=1}^{|W \cap A_i|} \frac{1}{t}$$

With this notion, we can define the PAV variant called *Local-Search PAV (LS-PAV)* (Aziz et al., 2018). Intuitively, LS-PAV operates as follows: starting with some arbitrarily chosen committee, the rule repeatedly searches for 'large enough' improvements in the PAV score (via candidate swaps between current committee members and unelected candidates) until no such improvements are possible.

Definition 2.17 (Local-Search PAV, LS-PAV)**.** *The rule begins with an arbitrary set of candidates* $W \subseteq C$ *that is considered the current committee and then it operates in rounds until termination. In each round, the rule looks for a pair of candidates* $c \in W, d \notin W$ *such that:*

$$\textit{PAV-score}(W \setminus \{c\} \cup \{d\}) \geqslant \textit{PAV-score}(W) + \frac{n}{k^2}.$$

If no such candidate pair exists then the rule terminates and returns the current committee W *as the final outcome. Otherwise, for candidate pair* $c \in W, d \notin W$*, it sets* $W = W \setminus \{c\} \cup \{d\}$*, i.e., the current committee* W *is updated by swapping candidates* c *and* d*.*

Having presented PAV (and its local-search variant), we now present a Thiele method whose study is also frequently motivated through its fairness qualities. It is the following rule that is named after Chamberlin and Courant (1983) who also defined this rule independently from Thiele (Lackner and Skowron, 2023).

Definition 2.18 (Approval-based Chamberlin Courant, α-CC)**.** *The rule called* Approval-based Chamberlin Courant, (α-CC) *rule is the Thiele method with scoring vector* $\boldsymbol{u}^{(k)} = (1, 0, \ldots, 0)$.

Unlike PAV, α-CC does not satisfy EJR, nor does it satisfy PJR. This rule is only guaranteed to produce committees that provide JR while also having the downside of being NP-hard to compute. However, we give it some attention due it being suited to tasks for diversity (or degressive proportionality) (Faliszewski et al., 2017). This concludes our dive into the class of Thiele methods as we move to multiwinner voting rules that are *sequential* in nature.[5]

Below, we define of an often studied rule in the multiwinner voting literature. It is a member of the rules proposed by Phragmén (Janson, 2016; Lackner and Skowron, 2023) and its sequential formulation is as follows.

Definition 2.19 (Phragmén's Sequential Rule, seq-Phragmén)**.** *During the run of this method, voters continuously earn funds and may pay to include a candidate into the committee. Each candidate* $c \in C$ *has a price of* 1*. Each voter begins with a personal budget of* 0 *and this budget continuously grows. If at some moment, the voters that approve of an unelected candidate* $c \in C \setminus W$*, collectively hold enough funds to pay the price of* 1*, then this candidate* c *is put into the committee* W*, the funds of all voters approving of candidate* c *are set to* 0*, and then the process continues until* k *candidates have been put into the committee* W*. Ties are broken arbitrarily if necessary.*

Now, we know that seq-Phragmén is polynomial-time computable. This is certainly an advantage that it has over our aforementioned Thiele methods of interest (namely, PAV and α-CC). However, seq-Phragmén only satisfies PJR while failing EJR. In this respect, it is outperformed by PAV. Now, another aspect where seq-Phragmén exhibits better qualities is when one looks at priceability. Specifically, seq-Phragmén always returns priceable outcomes whereas Peters and Skowron (2020) showed that no rule from the class of so-called *welfarist* rules—which includes the Thiele methods (and so includes PAV and α-CC)—is priceable.[6] Despite this contrast, a connection between PAV and seq-Phragmén presents itself

[5]Note that, to circumvent the issue of PAV and α-CC being NP-hard to compute, one may use these variants of Thiele methods that work in a sequential fashion (see (Lackner and Skowron, 2023) for details). We do not look to these sequential variants however, as they provide weaker proportionality guarantees than their non-sequential counterparts.

[6]The welfarist rules are those that maximise some function of the satisfaction that voters obtain from a committee (Peters and Skowron, 2020).

when one models apportionment (via restricting one's view to party-list profiles) as quite interestingly, both PAV and seq-Phragmén correspond to the D'Hondt method in the apportionment setting (Brill et al., 2017).

To close this presentation of rules, we detail another sequential rule that has drawn significant attention from those studying proportional representation within social choice. This rule is called the *Method of Equal Shares (MES)* (Peters and Skowron, 2020; Peters et al., 2021b). To be precise, what follows is a definition of the *first phase* of MES.

Definition 2.20 (Method of Equal Shares, MES)**.** *The rule will run over some rounds $r \in \{1, \dots, k\}$. Over rounds, voters may pay to assign candidates to seats in the committee. In every round r, to assign some candidate to a seat, a price of* 1 *must be paid. Let $b_i(r)$ be the budget that voter $i \in N$ has available to start round r. We set $b_i(1) = {}^k/_n$ as the initial budget of every voter $i \in N$. The rule starts at round* 1 *with an empty committee $W^0 = \emptyset$, and will add candidates to the committee over rounds. In round r, we say a candidate c is ρ-*affordable *for some $\rho \in \mathbb{R}_{\geqslant 0}$, with $c \in C \setminus W$, if:*

$$\sum_{i \in N(c)} \min(\rho, b_i(r)) \geqslant 1.$$

If there exists no candidate that is ρ-affordable, then the rule terminates and returns W^{r-1}, otherwise, for a ρ-affordable candidate c for a minimum ρ (use some arbitrary tie-breaking if there are multiple ρ-affordable candidates), the rule will fix $W^r = W^{r-1} \cup \{c\}$, i.e., add candidate c to the current committee from previous rounds. The budget of each voter $i \in N(c)$ is then set to $b_i(r+1) = b_i(r) - \min(\rho, b_i(r))$ while for voters $i \notin N(c)$, we set $b_i(r+1) = b_i(r)$.

It is known that MES's first phase satisfies EJR (Peters and Skowron, 2020). We have previously seen that it is computationally hard to check if a committee W provides EJR and PJR, but with MES, we have a polynomial-time computable rule that returns a committee W providing EJR (so, also PJR).

As the committees returned by this first phase of MES may only be partial (much like the aforementioned GCR) and we are generally interested in committees that are of size k, we must consider how to complete these MES committees. Mostly, we assume that arbitrary unelected candidates are chosen to fill any partial committee. In such cases, we write MES[arbitrary] to denote the combination of MES's first phase and this arbitrary completion. We can also consider completing a partial, MES committee using seq-Phragmén by having the following: voters begin the seq-Phragmén process with their leftover budgets from the end of MES's first phase (instead of kickstarting seq-Phragmén with each voter holding 0 in funds), and run seq-Phragmén as usual. We shall write MES[seq-Phragmén] to refer to the resultant rule.

While all MES committees provide EJR, regardless of the completion method used, the way of completion becomes relevant for the priceability property. Specifically, the partial committees of MES are always priceable, but the MES coupled with some completion method only returns priceable committees if the completion method preserves priceability. So, for instance, MES[seq-Phragmén] is priceable while MES[arbitrary] is not.[7] Interestingly, in the apportionment setting, a committee W with $|W| = k$ is priceable if and only if it corresponds to the seat assignment returned by the D'Hondt method and from this, we see that MES[seq-Phragmén] is also equivalent to D'Hondt in the apportionment setting (as we previously noted for PAV and seq-Phragmén) (Peters and Skowron, 2020). All detailed rules and their properties are summarised in Table 2.1.

	EJR	PJR	JR	Priceability	poly-time
PAV	✓	✓	✓	✗	✗
LS-PAV	✓	✓	✓	✗	✓
α-CC	✗	✗	✓	✗	✗
seq-Phragmén	✗	✓	✓	✓	✓
MES	✓	✓	✓	✓	✓
MES[arbitrary]	✓	✓	✓	✗	✓
MES[seq-Phragmén]	✓	✓	✓	✓	✓
GCR	✓	✓	✓	✗	✗
Committee Check	coNP-comp.	coNP-comp.	P	P	

Table 2.1: The first four columns indicate whether an approval-based multiwinner voting rule satisfies a proportionality axiom (✓), or fails it (✗). Only those axioms that are *known* to always be satisfiable are included (so not the core, stable priceability or balanced stable priceability). Note that the result highlighted in yellow ( ) is due to MES's first phase possibly returning partial committees and priceability for size-k committees relying on the completion method that is used with MES. The final column indicates whether a rule is polynomial-time computable (✓) or not (✗). The bottom row then details, for each axiom, the computational complexity associated with checking whether a committee provides the axiom in question.

[7] For an extensive study of the effect of the chosen completion method when looking to complete priceable committees, see (Brill and Peters, 2024).

Chapter 3

Justified Representation in Complex Domains: A Glimpse at Existing Approaches

So far, we have introduced the thesis' goal and laid down the technical foundation that readers require to move through the remainder of the thesis. With this chapter, we offer another way for readers to get prepared for the content to come. Specifically, in this chapter, we wish to highlight existing work from the literature that holds similar goal as we do, at least broadly speaking. That goal is to ensure proportional representation when one increases the complexity (or richness) of a domain. We focus on efforts in guaranteeing proportional representation in the vein of some of the axioms defined in Chapter 2. Specifically, the quota axioms of apportionment, and the justified-representation axioms of multiwinner voting.

Now, before continuing with our examination of the literature, we acknowledge that there are many proportionality investigations within many of these (upcoming) domains that we do not detail. For example, many touch on adaptations of the priceability notion to their settings. We emphasise that this chapter serves only to provide a glance at what has been done, both in terms of scope and detail. For those interested in more detailed look at the various notions (as well as many others), we recommend consulting the various references to the work.

3.1 Apportionment Extensions

This brief section is dedicated to a strand of the computational social choice literature that extends the standard model of apportionment.

Multiple Attributes. We begin by highlighting the work done by Lang and Skowron (2018) on providing *multi-attribute proportional representation* where candidates have multiple attributes and the chosen committee must adhere to certain proportionality constraints for each given attribute. The apportionment model can be seen as a special case of their model where there is only a single attribute per candidate, their associated party, and the chosen committee seats must be distributed in such a way that respects the proportion constraints imposed by the votes of that the parties received. They consider a *respect for the quota (RQ)* axiom from apportionment that requires both their notions of LQ and UQ to be satisfied simultaneously. RQ is then adapted to the multi-attribute setting by requiring (roughly speaking) that LQ and UQ are met for every attribute. In terms of apportionment methods, they assess how their multi-attribute adaptations of LRM and the divisor methods fare when measured against RQ. For example, they find that their multi-attribute LRM satisfies their multi-attribute RQ (under the assumption that full supply holds). Note that there has also been much study of the related model of *biapportionment* (or *biproportional apportionment*) (Ricca et al., 2017; Balinski and Demange, 1989a,b).

Ranking the Parties. Here we consider the work of Airiau et al. (2023) who study a model of *portioning*. Here, a divisible public resource, such as money, is to be distributed to various projects. In the model of Airiau et al. (2023), voters have ordinal preferences over the projects. A ready application of this is in the apportionment setting but with voters providing a ranking of the parties instead of just indicating the party that they support.[1] Notably, they show that an axiom that is similar in spirit to our proportionality axioms of interest, called *SD-core*, is satisfied by members of a class of portioning rules based on *Nash welfare*.

Approving Multiple Parties. The setting of *party-approval elections* studied by Brill et al. (2020, 2024) can be seen as an instance of apportionment where voters are able to cast approval ballots in support of more than one party. Party-approval elections lies between the apportionment setting and that of multiwinner voting as it generalises the former while being a special case of the latter. Due to the latter observation, their approach is mainly through the lens of multiwinner voting where they show numerous positive results. They most notably showed the existence of multiwinner voting rules that satisfy the core in this model (namely PAV and its local-search variant) as well as showed the compatibility of EJR and committee monotonicity by the design of a party-approval rule that combines notions of portioning and apportionment (Brill et al., 2020, 2024).

Apportionment over Time. Harrenstein et al. (2022) take a long-term perspective on apportionment where the apportionment process is no longer a one-shot instance, but occurs over rounds and where unelected candidates carry over to the future rounds. In this work, they find that taking into account the outcomes of past elections results in outcomes satisfying strong proportionality properties.

[1] Note that other models of portioning can be used to directly model the apportionment task but these are not discussed as either they do not extend beyond apportionment, or, they are mentioned in Section 3.2 with respect to multiwinner voting.

More specifically, they formally investigate the ideas of Gottlob Frege and find that in the long run, Frege's proposed apportionment method for this model provides proportional outcomes in the limit of the number of rounds.

3.2 Multiwinner Voting Extensions

Now, we look at work done in building on the multiwinner voting model. We focus on approval ballots but there is also work done with strict rankings (Elkind et al., 2017) and weak order preferences (Aziz and Lee, 2020).

Participatory Budgeting

The introduction made special mention of the participatory budgeting problem and we discuss research on this model here. In the participatory budgeting model, the candidates are now *projects* to be implemented and each project has an associated *cost*. Then the committee target size of k becomes a budget limit b on the sum of costs of the implemented projects.[2] So here, a viable outcome is one where this budget limit b is not exceeded.

The cost associated to each project leads to greater uncertainty on what the voters' utilities for each project are and this has led to more options when it comes to the interpretation of voters' satisfaction with an outcome. Examples include *cardinality-based* satisfaction which mirrors that of approval-based multiwinner voting as voters care about the number of implemented projects that they approve of, *cost-based* satisfaction where voters care about the sum of costs of the implemented projects that they approve of, and many more. In general, the type of voter satisfaction considered has an effect on the results and we make clear if this is the case with any results we mention. Note that there are other justified-representation-like works for participatory budgeting, as well as various extensions of this model, which we will not touch here. We refer interested readers to the extensive survey by Rey and Maly (2023).

The notion of ℓ-cohesiveness is adapted to participatory budgeting as P-cohesive groups. Here, instead of groups being cohesiveness over some number ℓ of committee seats that is an appropriate portion of the k seats (with respect to the group's size), they are instead considered cohesive over some set of projects P whose total cost is an appropriate portion of the budget (with respect to the group's size) (Peters et al., 2021b). This P-cohesiveness can then be used to build axioms of proportionality which are to be parameterised by a satisfaction function under consideration.

[2]Note that for this brief discussion, we specifically consider the *indivisible* model of participatory budgeting where projects cannot be partially funded (so they must either be fully funded, or not be implemented at all).

This gives rise to a version of EJR which requires that some voter within a P-cohesive group has enough satisfaction from the participatory budgeting outcome. Positively, this EJR axiom for participatory budgeting can always be satisfied for any type of satisfaction by use of a participatory budgeting version of GCR. However, this rule is not polynomial-time computable and Peters et al. (2021b) showed that this EJR axiom cannot be satisfied by a polynomial-time computable rule for any type of satisfaction. It is also known that a PB version of MES satisfies this EJR axiom for the cardinality-based satisfaction. The next step moves to weakenings of this participatory budgeting version of EJR. First, use the "up to any" and "up to one" relaxations frequently seen in the fair division literature (Amanatidis et al., 2022). *EJR-X* then states that for those P-cohesive groups that are witness to a violation of EJR, there exists a voter in said group such that any unselected project that they approve of was added to the participatory budgeting outcome, then this voter exceeds the required satisfaction to satisfy EJR. This EJR-X axiom yields more positive computational results as it can be satisfied by the participatory budgeting version of MES for a class of satisfaction functions that includes not only cardinality-based satisfaction but also cost-based satisfaction, among others (Peters et al., 2021b; Brill et al., 2023b). The "up to" weakening can then be taken a step further to reach *EJR-1* which Peters et al. (2021b) showed can be satisfied by the participatory budgeting version of MES for any type of satisfaction that is *additive* (Peters et al., 2021b). Also, there has been significant attention paid to weaker PJR-like axioms for participatory budgeting as well as adaptations of priceability (Peters et al., 2021b; Los at al., 2022; Brill et al., 2023b).

Budget Division and Probabilistic Social Choice

We briefly discussed the portioning model in this chapter's previous section on apportionment extensions. Recall that in this model, a finite public resource is to be divided between a set of projects. This model, introduced to social choice by (Bogomolnaia et al., 2005), is also known as *budget division* and is one that is closely related to that of (indivisible) participatory budgeting. Indeed, using the portioning model, it is possible to model *divisible* participatory budgeting instances where projects can only be partially funded using the portioning model. It is this connection that prompts us to mention it in this chapter. Note that the aforementioned work by Airiau et al. (2023) uses ranking-based ballots so we do not consider it in this section that focuses on approval-based settings. Fain et al. (2016) look at the core adapted to this portioning model. The work by Aziz et al. (2020) then ventures into territory of justified representation by assessing the viability of notions such as PJR, JR and proportionality degree within their setting. Brandl et al. (2022) investigate the Nash product rule that is known to provide proportional outcomes in the form of satisfying an axiom that provides groups of voters with a fair share similar to that guaranteed by justified representation. Other proportionality-related work on these models includes the study of the *price of fairness* where Michorzewski et al. (2020) highlight the utilitarian social welfare tradeoffs that are necessary in demanding proportional representation.

A model that is similar to portioning but whose line of work carries a different interpretation and differing approaches is the model of *probabilistic social choice* (Brandt, 2017). Here, the resource division is equivalent to creating a probability distribution over the alternatives (so the resource to be divided is the probability of 1). Then the amount that is assigned to an alternatives is that alternative's probability of being selected post-division. While we mostly mention those that most closely related to justified representation, we also note that fairness notions that veer from justified representation have also been studied such as individual fair share (Bogomolnaia et al., 2005). Duddy (2015) studied a *proportional sharing* axiom that is reminiscent of justified representation. Specifically, this requires that for every subgroup of voters that is an α fraction of the population, there is at least α probability that one member in said group approves the final outcome.

Approval-based Social Choice in General

Recent work by Masarík et al. (2023) has looked towards providing proportional representation from a more general viewpoint. They introduce a general model for collective decision-making and their framework captures not only multiwinner voting, but its generality allows for the modelling of other interesting domains (some of which are mentioned later in this chapter). Thus, their proportionality results touch on various settings. Their model can be described as one of approval-based multiwinner voting with a *variable* number of winners (so not some fixed k) along with the key addition of *feasibility constraints*. These feasibility constraints fix which of the possible committees are valid outcomes for the process.

We highlight two of their EJR adaptations. The first is what they call *Base EJR (BEJR)*. This axiom states that a group of voters N' deserves ℓ candidates if for every feasible committee $W \in \mathcal{C}$ there is exists a set $X \subseteq \bigcap_{i \in N'} A_i$ with $|X| \geqslant \ell$ such that either $W \cup X$ is feasible, or $|N'|/n > \ell/(|W|+\ell)$. They showed that BEJR can always be satisfied by a rule that is similar in spirit to GCR and this holds for any constraint. They then proposed a stronger notion, that they simply refer to as EJR, where the difference from BEJR is that instead of a voter group N' considering any feasible outcome $W \in \mathcal{C}$, they only consider the subsets of the outcome at hand. It remains open whether their EJR can always be satisfied in general but they have shown that it can be satisfied for restricted classes of constraints. In terms of voting rules, they extend voting rules such as PAV and seq-Phragmén to their general constrained setting.

We now finish this part mentioning work by Mavrov et al. (2017) who studied a problem that is similar via the presence of constraints: how does one obtain proportional committees in an approval-based multiwinner model with diversity constraints. In this work, they also adapt EJR. We direct readers to (Masařík et al., 2023) for an extensive discussion of the EJR proposed by Mavrov et al. (2017) as well as a comparison between this EJR and the axioms that Masařík et al. (2023) study.

Public Decisions

The *public decisions* setting (Conitzer et al., 2017) is one of the models captured by the aforementioned general social choice introduced by Masařík et al. (2023). In this setting, a group of voters is presented with a set of issues for which they are expected to make a binary choice: typically, deciding to either *accept* or *reject* each issue. This setting has recently been studied with the aim of ensuring proportional representative public decisions. From the perspective of justified representation, there have been two interpretations of the notion in the public-decision setting.

The first approach intuitively states that 'a group of voters that agree on a set of issues T and represent an α fraction of the voter population, should control a $\alpha \cdot |T|$ number of the total issues in $\mathcal{I}$' (Skowron and Górecki, 2022). We refer to it as *agreement-EJR*. This notion aligns with the EJR notions used by Masařík et al. (2023) and as this public-decision model without constraints can be captured in their model, the positive proportionality results from there carry over.

The next EJR interpretation was defined by Freeman at al. (2020). They first defined a notion cohesiveness for public decisions that is analogous to that of multiwinner voting. They use this to define justified-representation notions that are more faithful translations of those from multiwinner voting. Their EJR axiom roughly states that 'a cohesive group of voters that agree on an α fraction of the issues and represent an α fraction of the voter population, should control $\alpha \cdot m$ of the issues in $\mathcal{I}$' (Freeman at al., 2020). We refer to this version of EJR as *cohesiveness-EJR*. Freeman at al. (2020)'s work shows that cohesiveness-EJR is much harder to achieve in this public-decision setting than it is in the multiwinner case. For example, they demonstrate that a variant of MES fails their cohesiveness-EJR axiom.

Proportionality has also drawn attention in settings that are similar to public decisions while also dealing with constraints. This includes the public decisions model with issues being interdependent via the use of *conditional ballots* (Brill at al., 2023c). Also, Haret et al. (2020) introduced justified representation to the model of *belief merging*.

Sequential Decision Making

There has been recent focus placed on *sequential decision making* where temporal elements have been introduced to the models of decision-making. Those that we consider can be seen as generalisations of multiwinner voting (Lackner, 2020;

Lackner and Maly, 2021b) and the public decisions model (Bulteau et al., 2021; Chandak et al., 2024). Now, while some work has studied various aspects such as potential for voters' strategic behaviour (Lackner et al., 2023), there has also been extensive focus on proportionality.

The recent trend was initiated by (Lackner, 2020) with a model of *perpetual voting*. Here, the approval-based multiwinner voting task is repeated over multiple rounds. This work looked more at ensuring fair outcomes for individual voters instead of voter groups as prescribed by the justified-representation notions.

Lackner and Maly (2023) studied how to make proportionally representative decisions in the perpetual voting model and for example, studied a perpetual version of seq-Phragmén. They also define a *lower quota for closed groups* axiom for the sequential setting that is reminiscent of the EJR.

We follow with focusing our attention to the recent work by Chandak et al. (2024) who look at the task of selecting m alternatives in m rounds. we use the terminology used by Chandak et al. (2024). They define two EJR axioms for their model in a manner that aligns with the agreement EJR notion that we discussed earlier in this section with respect to the public-decision model. The weaker EJR says if there is a voter group N' with $|N'| \geqslant \ell \cdot n/m$ and N' agrees in every round, there must be a voter that approves the decision in at least ℓ rounds. The stronger axiom, which they call *Strong EJR*, is more demanding and says that a voter group N' that agrees in $k \leqslant m$ rounds with $|N'| \geqslant \ell \cdot n/k$, there must be a voter that approves the decision in at least ℓ rounds. In terms of rules, they differentiate between rules that are *online* (round by round) such as perpetual seq-Phragmén, *semi-online* (round by round but the total number of rounds is known in advance) and *offline* (decisions made over all rounds simultaneously). They then adapt MES (considered a semi-online rule) and show that it satisfies EJR. They also adapt LS-PAV to their setting (considered an offline) and show that it satisfies Strong EJR. They go a step further and show that stronger proportionality guarantees (similar to the aforementioned cohesiveness EJR notion from the public-decision model) are not always achievable in their model. Specifically, strengthenings that are more faithful to the multiwinner definitions of EJR and PJR are not always satisfiable in their setting. Note that this holds in their setting when there are more than two alternatives available per round.

Also, Bulteau et al. (2021) analysed justified representation in the model of perpetual voting and proposed axioms that are the PJR analogues of the EJR and Strong EJR axioms defined by Chandak et al. (2024). We also mention the recent work by Elkind et al. (2024) who operate in this sequential setting and study the computational complexity of verifying whether an outcome is proportional, i.e., whether the outcome provides a proportionality axiom such as EJR. Note that there is also work done in ensuring fairness for participatory budgeting instances run over time (Lackner et al., 2021).

Research of a similar line is that of *online* decision-making where candidates appear one at a time and the decision-maker must make the irrevocably decide to include the candidate or not, before continuing to other candidates. In this online setting, proportionality was studied for online versions of committee elections (Do et al., 2022) and participatory budgeting (Banerjee et al., 2023).

Miscellaneous

This section details more examples of complex domains that have had proportionality studied within them.

Rankings as Outcomes. Here, voters present approval ballots but instead of returning some subset of the candidates as an outcome, the goal is to return a ranking of all the candidates as an outcome. Skowron et al. (2017) initiated the study of proportionality in this setting. Their notion states that a voter group that is cohesive and large enough should be entitled to control of a certain portion of a top segment of the collective ranking (with this top segment being determined by the groups cohesiveness and size). They then introduced a quantitative measure of proportionality that looks at the guarantees that sequential rules—such as seq-Phragmén—ensure in their rankings. Israel and Brill (2021) took it a step further and assessed *dynamic* proportional rankings. In their setting, during the sequential construction of the final collective ranking, the rankings are dynamically recomputed after every new candidate is added to the ranking. For variants of sequential rules such as seq-Phragmén, they also investigate how the rules fare against a quantitative measure of proportionality.

Approval-based Facility Location. Recently, there was the initiation of the study of the heterogeneous facility location problem where there are not enough locations to house each facility (Deligkas et al., 2023). With the work of Deligkas et al. (2023) focusing on axiomatic issues such as strategyproofness, Elkind et al. (2022) took this setting and analysed the possibilities of importing justified representation to it. They find that even their adaptation of JR is not always satisfiable. Thus, they proposed a weaker, related axiom that is better suited to their setting and is always satisfiable.

Trichotomous Ballots. The general model of Masarík et al. (2023) can also capture scenarios of the committee election problem but with the voters able to submit trichotomous ballots. These are ballots indicating the candidates that they either (*i*) approve of, (*ii*) disapprove of, or (*iii*) are neutral/indifferent on. The most notable attempt in representing voters proportionally in this setting was done by Talmon and Page (2021) who mostly found negative results, further exemplifying the difficulty of this task of transferring these proportionality notions to richer domains.

Cake Sharing. Bei et al. (2022) take the classical *cake cutting* problem of fair division (Procaccia, 2016) and adapt it to a public-good variant that they call *cake sharing*. Here, voters vote to decide on the subset of the cake that they will share amongst themselves. In the paper by Bei et al. (2022), they adapt EJR and PJR to this setting in order to test the proportionality qualities of their cake cutting mechanisms of interest.

Mixed Goods. Lu et al. (2023) generalise both the aforementioned work on cake sharing as well as the approval-based multiwinner voting model. In this work, they investigate how notions of proportionality can be imported to model where the ‘candidates’ that are voted on are resources that may be divisible or indivisible. They then refer to this as an approval-based multiwinner voting model with *mixed goods*. In terms of their analysis, they define a version of EJR for their mixed goods model (as well as a “up to one” weakening of it) and assess whether these can always be satisfied. In particular, they check whether their mixed-goods adaptations of rules such as MES and PAV satisfy any of these axioms.

Possibly Unavailable Candidates. Brill et al. (2023a) enrich the committee election task by asking the question: how to adapt proportionality if candidates that some voting rule chooses to be part of the committee are unavailable to join said committee? Their study focuses on *query policies* where candidates are sequentially queried on whether they are available to join the committee. They investigate the proportional multiwinner voting rules as well as axioms such as EJR and PJR within the context of this model and in particular, the notion of query policies.

Incomplete Votes. Halpern et al. (2023) tackle the problem of multiwinner voting where voters submit approval ballots that may be *incomplete*, i.e., a voter may not submit an opinion (either approve or disapprove) for some candidates. Specifically, they look to ensure some sort of proportionally representative committees are produced despite this (potential) lack of information. Their focused applications are platforms whose sets of candidates that are too large for individuals to assess in full. Thus, their approach is to study a model that allow queries for voters to be queried (at random) on their opinion of some subset of candidates. Then in this work, they study issues such as the number of queries needed to achieve JR as well as bounds on the number of queries that certain algorithms (such as one based on LS-PAV) require to achieve justified representation (in the form of EJR, PJR and JR).

Matching Markets. We also mention work by Boehmer et al. (2022) that takes on the problem of finding multiple matchings between voters who have approval preferences over each other. Formally, they model the matching problem using the multiwinner voting model by taking the candidate set to be the set of matchings between candidates. They focus on producing matchings that are proportionally fair with this fairness in this case being represented by importing the notion of justified representation. On their results, they show that more positive proportionality results are possible in their model. For example, their adaptation of PAV is polynomial-time computable in their model in contrast to PAV in the multiwinner voting model.

Priority Candidates. There is also the study of multiwinner voting where there are candidates who are given *priority* (Huising, 2023). Any outcome in this setting is considered a valid committee if the number of priority candidates within the committee reaches a predefined quota. Huising (2023) then adapted the axioms of justified representation to this model and tested variants of rules such as α-CC and seq-Phragmén against these axiom adaptations.

3.3 Takeaways

This chapter serves as a brief survey on work done on taking proportionality from apportionment and approval-based multiwinner voting, and lifting it to domains that we deem to as more 'complex' version of these standard models. Now, we conclude by mentioning some points that stem from this chapter that we deem worthy of keeping in mind moving forward through the thesis.

First, it is clear that lifting proportionality to what we call these complex domains has drawn, and continues to draw, significant attention from the computational social choice community. This further motivates the study of other such complex domains as done in the remainder of the thesis.

Then, note that in some parts of this chapter there were justified-representation adaptations that were presented in greater detail such as with participatory budgeting and sequential decision-making. The hope is that this provided readers with a sense of *how* proportionality has been lifted to the richer domains and thus, prepare them for how we do so for the domains that we study.

Also, from the work that has been conducted, it is evident that lifting strong proportionality requirements such as EJR is a task that is not only difficult, but that requires non-trivial conceptual work in importing justified representation in a manner that is well-suited to the complex domain in question but also results in achievable axioms.

Finally, in much of the work we have outlined, the computational complexity of the various rules is an aspect that was assessed. This matches our outlined goal that, ultimately, we wish to develop rules that can be used for these more complex domains *in practice*.

Part One

Some Seats Have More Value Than Others

Chapter 4

Apportionment with Weighted Seats

We delve into our first complex domain and it uses the apportionment model as a foundation. The merits of different apportionment methods are well understood, due to an elegant mathematical theory developed for the political realm (Balinski, 1982; Pukelsheim, 2014). However, existing work on apportionment remains restricted by the assumption—often not met in practice—that all seats are of equal value. In this chapter, we put forward an enriched model in which seats may have different weights reflecting their (objective) values.

There are numerous scenarios that fit this richer model: from the distribution of non-liquid assets in bankruptcies to beneficiaries with different entitlements, to the assignment of positions on news websites to editorial domains (such as politics, business, or sports), based on the readership's relative levels of interest in those domains. A particularly salient illustration is offered by the way special-purpose committees are constituted in the *Bundestag*, Germany's national parliament (Bundestag, 2023). These are committees with specific responsibilities (e.g., Budget or Defence), established anew in every political cycle. Usually, which party gets to nominate the head of each committee is the result of a negotiation—but when no consensus can be found, which happened in 8 out of 20 parliamentary sessions since 1949, standard apportionment methods are used. But different committees have different size and influence, so positions end up differing in value: the role of chair of the Budget Committee will be valued more highly than, say, that of chair of the Tourism Committee. As the values of positions will factor into the satisfaction of parties, a standard apportionment method (treating all positions the same) cannot possibly do justice to this scenario.

As noted in Chapter 2, the central problem in apportionment is finding an assignment of seats to parties that is as proportional as possible, given that perfect proportionality is often not feasible (e.g., for a parliament with 100 seats, there is no perfect apportionment for three parties that each obtained exactly one third of the votes). In our enriched model we associate each seat with a weight representing its objective value, and approach proportionality through the lens of the total weight available. We find that generalisations of the two central proportionality axioms of apportionment, *lower quota* and *upper quota*, are impossible to satisfy in general, but that apportionment methods that faithfully extend well-known standard methods satisfy natural relaxations of these axioms. The relaxations, based on the concept of satisfaction 'up to one' or 'up to any' seat, utilise an idea commonly seen in fair division (Budish, 2011; Caragiannis et al., 2019), that has also made its way to participatory budgeting (Peters et al., 2021b; Brill et al., 2023b). Additionally, we study *envy-freeness*, the central axiom in fair division (Amanatidis et al., 2022), which turns out to be related to upper quota. We show that envy-freeness up to any seat, an axiom that remains elusive in the general fair division setting, is satisfiable in our setting. Finally, we find that a direct generalisation of the *house monotonicity* axiom is prohibitively demanding, but weaker forms are readily satisfied by weighted counterparts of well-known apportionment methods.

The main take-away of this analysis is that achieving good solutions to the apportionment problem is harder in the presence of weights, but mild relaxations of the generalisations of standard properties of interest to our enriched model often lead to positive results. In particular, one can achieve stronger guarantees than those achievable for scenarios with subjective weights. Specifically, we show that these improved guarantees can be achieved by using objective weights to define weighted variants of the apportionment methods, which would not be possible in the fair division setting.

Additional Related Work. Our weighted apportionment model is located between standard apportionment and weighted fair division. In fair division (Amanatidis et al., 2022), we also aim to assign goods of different values to voters, but while in our model the weights of the seats are the same for all parties, in fair division voters usually have different valuations for the goods being allocated. The focus of the fair division literature, and certainly the part of the literature considering relaxations of classical fairness notions, is the case where all voters deserve the same utility, while in apportionment the central concerns directly stem from the fact that parties may have different entitlements. Nevertheless, there is a growing literature on fair division for voters with different entitlements (Suksompong, 2025), see e.g. the works of Farhadi et al. (2019), Aziz et al. (2020), and Babaioff et al. (2021), who study the possibility of achieving proportional allocations for voters with (positive) cardinal utilities. This can be seen as a generalisation of our model.

Most closely related to our work is a paper by Chakraborty et al. (2021), who investigate how apportionment methods can be used to produce *picking sequences* that guarantee fair allocations in this setting, and we shall refer to it often. While this richer model makes it possible to represent a wider range of scenarios, is also has downsides, notably in view of preference elicitation: requiring each party to declare its utility for each seat may be infeasible in practice. More importantly, the added generality of allowing voters to have different valuations for goods makes it much harder to achieve proportionality. Thus, as we shall see, the fairness guarantees provided by the fair division literature are significantly weaker than the ones that are achievable in the apportionment setting.

At first glance, the manner in which participatory budgeting generalises multi-winner voting by adding weights (when moving from candidates to projects with costs) may look similar to the generalisation that we propose here. But there are crucial differences: in participatory budgeting, weights are attached to party members rather than to seats. We will see later (both in this chapter and even more so in Chapter 5) that this difference also manifests itself in the normative properties one can achieve.

Also, in Chapter 3, we noted some extensions of the apportionment model as well as models such as those studied by Chandak et al. (2024) and Masarík et al. (2023)) that can be seen as generalisations of the apportionment model. However, all of these mentioned works still assume that all seats are of equal value. Thus, they extend the apportionment model in quite different ways.

Finally, our model can be seen as a special case of the public-decision model of Conitzer et al. (2017) where each issue's alternatives are exactly the set of parties and with a party's utility for a seat being that seat's weight.

Chapter Outline. Section 4.1 introduces the weighted apportionment model. Sections 4.2 and 4.3 offer an axiomatic and computational analysis of lower and upper quota properties, respectively, while Section 4.4 focuses on house monotonicity. Section 4.5 presents the experiments and contains both a case study of the German Bundestag as well as an analysis on synthetic data. We conclude with a summary of the chapter with Section 4.6.

4.1 The Model

In our weighted-seat apportionment model there are n voters, each voting for one of m parties, and the goal is to fill k seats of varying (objective) value with members of those parties, based on the votes.

We make two mild assumptions throughout. First, each party is approved by at least one voter. Second, we assume full supply holds: recall from Chapter 2 that this means that each party has at least k members so it can fill all of the k seats by itself.

We again have a vote vector $\boldsymbol{v} = (v_1, \dots, v_m) \in [n]^m_{\geqslant 1}$ that specifies how many votes party $p \in [m]$ garnered with the total number of votes $\sum_{p\in[m]} v_p$ equal to the number n of voters. Each seat $t \in [k]$ is associated with a weight w_t indicating how valuable t is. Thus, the environment in which the election takes place can be described by a *weight vector* $\boldsymbol{w} = (w_1, \dots, w_k) \in \mathbb{N}^k_{\geqslant 1}$, listing these weights in non-increasing order. Let $\omega = \sum_{t\in[k]} w_t$ be the total weight. A *seat* *assignment* is a vector $\boldsymbol{s} = (s_1, \dots, s_k) \in [m]^k$, where $s_t = p$ means that party $p \in [m]$ is assigned seat $t \in [k]$ with weight w_t. Given a seat assignment $\boldsymbol{s}$, we write $\boldsymbol{s}(p) = (t)_{s_t=p}$ for the vector of seats, in increasing order of index, assigned to party p under seat assignment $\boldsymbol{s}$. A *weighted apportionment instance* is a pair $(\boldsymbol{v}, \boldsymbol{w})$ of a vote vector $\boldsymbol{v}$ and a weight vector $\boldsymbol{w}$. We speak of a *unit-weight* instance in case $w_t = w_{t'}$ for all $t, t' \in [k]$.

As was discussed in Chapter 2, proportionality in the apportionment task is typically formalised in terms of the quota of a party. We do the same in our weighted setting, with the important caveat that the quota, in this case, is construed in terms of the total weight. So for the weighted setting, we set the quota of party p to be defined as $q_\omega(p) = \omega \cdot {v_p}/{n}$. In order to judge whether a party satisfies their quota, we need to reason about the weights accrued by a party via its seat assignment. This leads us to the notion of *weighted representation*. Now, the representation of party p derived from seat assignment $\boldsymbol{s}$ is formally defined as $r_\omega(p, \boldsymbol{s}) = \sum_{t \in \boldsymbol{s}(p)} w_t$, i.e., the sum of the weights of the seats assigned to p according to s. For a weight vector $\boldsymbol{w}$, the set of all possible representation values that a party can obtain from occupying at most $h \in [k]$ seats can be computed as follows:

$$R(\boldsymbol{w})_{[h]} = \left\{ \sum_{t \in T} w_t \mid T \in \mathcal{P}([k]) \text{ with } |T| \leqslant h \right\}.$$

We use the below example to illustrate the various parts of the model.

Example 4.1. Consider the following weighted apportionment instance with three parties and four seats:

	# of votes ($n = 100$)	$q_\omega(p)$
Party 1	60	13.2
Party 2	30	6.6
Party 3	10	2.2

$$\boldsymbol{w} = (10, 6, 4, 2),\ \omega = 22$$

Let us now turn to the methods we will use to assign seats to parties. A *weighted-seat apportionment method* (WSAM) M_ω takes a weighted apportionment instance $(\boldsymbol{v}, \boldsymbol{w})$ as input and maps it to a winning seat assignment $M_\omega(\boldsymbol{v}, \boldsymbol{w})$. We focus on two types of WSAMs, which generalise the most prominent methods for standard apportionment: *divisor methods* and the *largest remainder method (LRM)* (Balinski, 1982). Divisor methods afford several equivalent definitions—the sequential version proves most useful for our setting.

Definition 4.1 (Weighted divisor method)**.** *Given a weighted apportionment instance* $(\boldsymbol{v}, \boldsymbol{w})$ *and a function* $f_\omega : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$*, the weighted divisor method for* f_ω *works in* k *rounds, as follows. In round* $t \in [k]$*, seat* t *is assigned to the party* p *with the highest value for:*

$$ratio_p = \begin{cases} \frac{v_p}{f_\omega(g_p^\omega(t), w_t)} & \text{if } f_\omega(g_p^\omega(t), w_t) \neq 0 \\ \infty & \text{if } f_\omega(g_p^\omega(t), w_t) = 0, \end{cases}$$

where $g_p^\omega(t)$ *is the sum of the weights of the seats assigned to party* p *in earlier rounds. If required, a tie-breaking rule is used to choose between parties with equal ratio.*

Intuitively, weighted divisor methods allocate the available seats sequentially, starting with the most valuable seat, based on the ratio between values v_p and $f_\omega(g_p^\omega(t), w_t)$. It is, of course, possible to allocate the seats in a different (fixed) order but, to anticipate results to come, starting with the most valuable seat leads to particularly nice axiomatic properties.

Only certain choices for the function f_ω lead to reasonable weighted divisor methods. As mentioned in Chapter 2, the divisor methods that we hone in on are Adams, the unique divisor method satisfying upper quota, and D'Hondt, the unique divisor method satisfying lower quota (Balinski and Young, 1975). These two rules can be generalised to our weighted-seat setting as follows.

Definition 4.2 (Adams$_\omega$ Method)**.** *Adams*$_\omega$ *is the weighted divisor method defined by* $f_\omega(g_p^\omega(t), w_t) = g_p^\omega(t)$*.*

Definition 4.3 (D'Hondt$_\omega$ Method)**.** *D'Hondt*$_\omega$ *is the weighted divisor method defined by* $f_\omega(g_p^\omega(t), w_t) = g_p^\omega(t) + w_t$*.*

Example 4.2 (Adams$_\omega$ and D'Hondt$_\omega$)**.** Consider again the following weighted apportionment instance with three parties and four seats:

	# of votes ($n = 100$)	$q_\omega(p)$
Party 1	60	13.2
Party 2	30	6.6
Party 3	10	2.2

$$\boldsymbol{w} = (10, 6, 4, 2),\ \omega = 22$$

The Adams$_\omega$ method maximises the ratio ${}^{v_p}/_{g_p^\omega(t)}$. Since $g_p^\omega(t) = 0$ before a party receives any seats, each party gets a seat after the first three rounds; we assume tie-breaking assigns party 1, 2 and 3 seat 1, 2 and 3, respectively, for the partial assignment $\boldsymbol{s} = (1, 2, 3, \_)$. At round $t = 4$, $ratio_p$ is maximised by party 1, with $ratio_1 = {}^{60}/_{10}$ versus $ratio_2 = {}^{30}/_{6}$ and $ratio_3 = {}^{10}/_{4}$. The final assignment is $\boldsymbol{s} = (1, 2, 3, 1)$.

The D'Hondt$_\omega$ method maximises the ratio ${}^{v_p}/_{(g_p^\omega(t)+w_t)}$. Assigning the first seat to party 1 gives a ratio of ${}^{60}/_{(0+10)}$, versus ${}^{30}/_{(0+10)}$ and ${}^{10}/_{(0+10)}$ for parties 2 and 3, respectively, so this seat goes to party 1. The second seat goes to party 2. For the third seat we calculate $ratio_1 = {}^{60}/_{(10+4)}$, $ratio_2 = {}^{30}/_{(6+4)}$ and $ratio_3 = {}^{10}/_{(0+4)}$, so this seat goes to party 1. The final assignment is $\boldsymbol{s} = (1, 2, 1, 3)$.

Let us now turn to LRM. Intuitively, LRM assigns each party their lower quota of seats, as defined below, and then assigns the remaining seats based on the fractional remainder of each party's quota. As we will see in Theorem 4.2, this idea cannot work in the weighted setting. Instead, we put forward the following procedure.

Definition 4.4 (Greedy$_\omega$ Method). *In each round $t \in [k]$, the seat t with weight $w_t \in \boldsymbol{w}$ is assigned to the party p for which the value $q_\omega(p) - g_p^\omega(t)$ is maximal, with ties broken arbitrarily when arising.*

Intuitively, Greedy$_\omega$ always assigns a seat to the party p furthest away from its quota $q_\omega(p)$. Though not the obvious way of generalising LRM to the weighted setting, it can be checked that, without weights, Greedy$_\omega$ is equivalent to LRM. In the absence of a formal proof of this claim, we shall provide explicit proofs related to Greedy$_\omega$ where otherwise the result would follow directly from results for LRM in unit-weight instances.

Example 4.3 (Greedy$_\omega$). Consider again the following weighted apportionment instance with three parties and four seats:

	# of votes ($n = 100$)	$q_\omega(p)$
Party 1	60	13.2
Party 2	30	6.6
Party 3	10	2.2

$$\boldsymbol{w} = (10, 6, 4, 2),\ \omega = 22$$

For Greedy$_\omega$, note that the quotas $q_\omega(p) = \omega \cdot {}^{v_p}/_n$ for the parties are $q_\omega(1) = 13.2$, $q_\omega(2) = 6.6$ and $q_\omega(3) = 2.2$. And thus Greedy$_\omega$ returns the seat assignment $\boldsymbol{s} = (1, 2, 1, 3)$, the same as D'Hondt$_\omega$.

We note that the WSAMs defined above take the parties' previously assigned weights into account when deciding on a seat assignment. This is in contrast to Chakraborty et al. (2021), who use the standard apportionment methods to set the order. And since they consider a setting where items do not have objective values, it would not even be possible to define natural equivalents of our rules in their setting.

Also note that these three WSAM all have a common feature, that being that the seats are assigned in non-increasing order of weight. This seems the most natural choice with the most 'valuable' seats being given priority in getting assigned.

The axioms we present in the following sections are defined as properties of seat assignments, and we say that a WSAM M_ω *satisfies property* $\mathcal{P}$ if for every weighted apportionment instance $(\boldsymbol{v}, \boldsymbol{w})$ it is the case that the seat assignment $M_\omega(\boldsymbol{v}, \boldsymbol{w})$ satisfies property $\mathcal{P}$. Throughout, we work with weighted apportionment instances $(\boldsymbol{v}, \boldsymbol{w})$ with n voters, m parties, and k seats.

4.2 Weighted Lower Quota

We will now adapt the LQ axiom to our setting that weighted seats. For our weighted-seat setting it might seem natural to define the *weighted lower quota* as $\lfloor q_\omega(p) \rfloor = \lfloor \omega \cdot {}^{v_p}/_n \rfloor$. However, the following example shows that such a quota is not guaranteed to be achievable, even in the simplest case of two parties and two seats.

Example 4.4. Consider the following weighted apportionment instance with two parties and two seats:

	# of votes ($n = 100$)	$\lfloor \omega \cdot {}^{v_p}/_n \rfloor$
Party 1	50	50
Party 2	50	50

$$\boldsymbol{w} = (99, 1),\ \omega = 100$$

Observe that there is no way for both parties to get at least 50 in representation, given that one seat is significantly more valuable than the other.

Intuitively, the problem is that there may be no combination of seats that actually give each party its weighted lower quota. Let us thus try to restrict the lower quota of a party to the values that the party can achieve with the number of seats that it deserves, which is $l^{\#}(p) = \lfloor k \cdot {}^{v_p}/_{n} \rfloor$ for a party p. We now use this quantity to determine the party's *obtainable lower quota of weights*, $l^o(p)$ as follows:

$$l^o(p) = \max \left\{ w \in R(\boldsymbol{w})_{[l^{\#}(p)]} \mid w \leqslant q_\omega(p) \right\}.$$

We can now define our first proportionality property.

Definition 4.5 (Obtainable Weighted Lower Quota, WLQ^o)**.** *A seat assignment $\boldsymbol{s}$ provides WLQ^o, if for every party p, it is the case that $r_\omega(p, \boldsymbol{s}) \geqslant l^o(p)$.*

Note that with unit weights, WLQ^o is equivalent to the standard lower quota. In the weighted setting, however, computing $l^o(p)$ for any party p requires solving a Subset Sum problem and can hence not be done in polynomial time unless $\mathsf{P} = \mathsf{NP}$.[1] Note, as well, that WLQ^o is similar to the Extended Justified Representation (EJR) axiom in participatory budgeting: despite it being computationally difficult to compute the amount of representation deserved by a group of voters in this participatory budgeting setting, EJR can always be satisfied (Peters et al., 2021b). The same holds, now, for WLQ^o—in the case of two parties.

Proposition 4.1. *For every weighted apportionment instance with two parties there exists a seat assignment that provides WLQ^o.*

Proof. Consider a weighted apportionment instance with parties 1 and 2. By definition of the obtainable weighted lower quota, there exists a set T of seats such that $\sum_{t \in T} w(t) = l^o(1)$. Let $\boldsymbol{s}$ now be a seat assignment such that party 1 is assigned all seats in T and party 2 gets all of the other seats. By definition, $r_\omega(1, \boldsymbol{s}) = l^o(1)$ holds for party 1. For party 2 we have:

$$\begin{aligned} r_\omega(2, \boldsymbol{s}) &= \omega - r_\omega(1, \boldsymbol{s}) = \omega - l^o(1) \\ &\geqslant \omega - \omega \cdot \frac{v_1}{n} = \omega \cdot \frac{n - v_1}{n} = \omega \cdot \frac{v_2}{n} \geqslant l^o(2). \end{aligned}$$

Hence, WLQ^o is satisfied. □

Observe that computing this assignment takes exponential time in the general case, as we cannot compute the set of seats T in polynomial time (unless $\mathsf{P} = \mathsf{NP}$). Proposition 4.3 shows that this is unavoidable. However, in contrast to EJR, WLQ^o is not satisfiable in general with more than two parties.

[1] Note that this only holds when the weights are represented in binary within the input.

Theorem 4.2. *There are weighted apportionment instances for which there exists no seat assignment that provides WLQ^o.*

Proof. Consider the following weighted apportionment instance with three parties and three seats:

	# of votes ($n = 3$)	$l^o(p)$
Party 1	1	2
Party 2	1	2
Party 3	1	2

$$\boldsymbol{w} = (3, 2, 1),\ \omega = 6$$

We get $l^o(p) = 2$ for each party $p \in [3]$ but there exists no seat assignment that provides at least a weight of 2 to all three parties. □

Though WLQ^o cannot always be satisfied, as per Theorem 4.2, one might still ask for a WSAM that delivers an allocation satisfying WLQ^o on instances where this is possible. Unfortunately, the following result shows that such a requirement is not computationally tractable. The proof of this result involves a reduction from the NP-complete problem PARTITION, i.e., the problem of deciding, given a multiset $X = \{x_1, \ldots, x_k\}$ of k positive integers, whether there exists a partition of X into two subsets X_1 and X_2 such that $\sum_{x \in X_1} x = \sum_{x \in X_2} x$.

PARTITION

Given:	A multiset $X = \{x_1, \ldots, x_k\}$ of k positive integers.
Question:	Is there a partition of X into two subsets X_1 and X_2 such that $\sum_{x \in X_1} x = \sum_{x \in X_2} x$?

Proposition 4.3. *If there exists a polynomial-time algorithm $\mathbb{A}$ that finds a seat assignment $\boldsymbol{s}$ that provides WLQ^o whenever such a seat assignment exists, then* P = NP. *This holds even when restricted to the case where there are only two parties.*

Proof. Assume such an algorithm $\mathbb{A}$ exists and let $X = \{x_1, \ldots, x_k\}$, be an instance of the PARTITION problem. We create the following weighted apportionment instance $(\boldsymbol{v}, \boldsymbol{w})$, as follows. Set the weight vector $\boldsymbol{w} = (x)_{x \in X}$ to be the non-increasing vector of the k elements in X. Thus, we have $\omega = \sum_{x \in X} x$. Take two parties with $\boldsymbol{v} = (1, 1)$, so each party $p \in [2]$ receives half of the total votes. Now let $\boldsymbol{s}$ be the seat assignment produced by α on input $(\boldsymbol{v}, \boldsymbol{w})$. We claim X is a positive instance of Partition if and only if $r_\omega(p, \boldsymbol{s}) = {}^\omega/_2$ for every $p \in [2]$.

($\Rightarrow$) Assume that X is a positive instance of Partition. Thus, there exist subsets X_1 and X_2 such that $\sum_{x \in X_1} x = \sum_{x \in X_2} x$. In particular, this means that $\sum_{x \in X_t} x = {}^\omega/_2$, for each $t \in [2]$. Consider the constructed weighted apportionment instance $(\boldsymbol{v}, \boldsymbol{w})$. Each party deserves ${}^k/_2$ seats and has a weighted lower quota of ${}^\omega/_2$. As $\min(|X_1|, |X_2|) \leqslant {}^k/_2$, there exists a way to receive weight ${}^\omega/_2$ with ${}^k/_2$ seats. It follows that both parties have an obtainable lower quota of ${}^\omega/_2$. Finally, WLQ^o is satisfiable, as the seat assignment that gives all seats that correspond to an element of X_1 to party 1 and every seat that corresponds to an element of X_2 to party 2, satisfies WLQ^o. It follows that in the seat assignment produced by α each party must have representation ${}^\omega/_2$.

($\Leftarrow$) Assume that $r_\omega(p, \boldsymbol{s}) = {}^\omega/_2$ for both parties p. Let X_1 be the set of all elements that correspond to a seat that is allocated to party 1 and let X_2 be the set of all elements that correspond to a seat that is allocated to party 2. Then we must have $\sum_{x \in X_1} x = {}^\omega/_2 = \sum_{x \in X_2} x$, and hence X is a positive instance of Partition.

However, that means we can solve Partition in polynomial time by transforming it into the weighted apportionment instance $(\boldsymbol{v}, \boldsymbol{w})$, running α on that instance and checking whether $r_\omega(p, \boldsymbol{s}) = {}^\omega/_2$ holds for both parties $p \in [2]$. As Partition is NP-complete, this implies that P = NP. □

But, with some additional assumptions, we obtain a more positive outlook.

Proposition 4.4. *For a constant number of parties and weights in* $\boldsymbol{w}$ *that are polynomial in the input size, finding a seat assignment* $\boldsymbol{s}$ *that provides WLQ*^{o} *can be done in polynomial time, assuming such a seat assignment exists.*

Proof. We now describe a dynamic programming algorithm that finds a seat assignment $\boldsymbol{s}$ that provides WLQ^o whenever one exists. Consider a weighted apportionment instance with m parties.

The algorithm works as follows. For each $i \in [k]$ (where i represents the number of seats assessed thus far), it computes $\mathcal{W}_i$ which is a set of tuples of the form $(W_1, \ldots, W_m)$. Here, W_p indicates the sum of seat weights of the seats assigned to party $p \in [m]$.

Each $\mathcal{W}_{i+1}$ can be computed using $\mathcal{W}_i$ and the weight w_{i+1}, by looking at every combination of some tuple in $\mathcal{W}_i$ and some choice of party to assign the weight-w_{i+1} seat to. Once $\mathcal{W}_k$ is computed, we can check every tuple in $\mathcal{W}_k$ and for each tuple, assess whether it satisfies WLQ^o (which can be done in polynomial time). Specifically, this check can be done for each tuple $(W_1, \ldots, W_m)$ by assessing whether $W_p \geqslant l^o(p)$ for every party $p \in [m]$. From the assumption on the weights in $\boldsymbol{w}$ and the observation that computing l^o requires solving an instance of Subset Sum, we can apply a dynamic programming algorithm for the latter problem to compute $l^o(p)$ in polynomial time. And finally, there are at most ω^{2m} such tuples, which is polynomially many in the input size due to the assumptions on the weights in $\boldsymbol{w}$ and the number of parties being constant. Thus, this algorithm runs in polynomial time. □

Proposition 4.4's proof makes use of a dynamic programming algorithm. The assumptions of Proposition 4.4 may be restrictive, but they fit the scenarios envisioned for our model, which are not likely to feature a number of parties, or weight values, exponential in the input size.

We have, as of yet, made no inroads towards our goal of finding an achievable lower quota property for the weighted setting. To do so, it is helpful to look at lower quota in the unit-weight setting a bit differently: instead of thinking of the lower quota as the closest value to the quota that can be obtained in practice, we interpret it as guaranteeing that each party p is at most one seat away from *surpassing* its quota. To make this interpretation of lower quota work with weighted seats one must specify which seat, amongst those not assigned to it, a party has to additionally receive in order to surpass its quota. We parse this in three ways.

Definition 4.6 (WLQ up to one seat, WLQ-1)**.** *A seat assignment $\boldsymbol{s}$ provides WLQ-1 if, for every party p, either we have $r_\omega(p, \boldsymbol{s}) \geqslant q_\omega(p)$ or there exists some seat $t \in [k] \setminus \{t' \in \boldsymbol{s}(p)\}$ such that $r_\omega(p, \boldsymbol{s}) + w_t > q_\omega(p)$.*

Definition 4.7 (WLQ up to any seat, WLQ-X)**.** *A seat assignment $\boldsymbol{s}$ provides WLQ-X if, for every party p, either we have $r_\omega(p, \boldsymbol{s}) \geqslant q_\omega(p)$ or for every seat $t \in [k] \setminus \{t' \in \boldsymbol{s}(p)\}$, it holds that $r_\omega(p, \boldsymbol{s}) + w_t > q_\omega(p)$.*

Definition 4.8 (WLQ up to any seat from an overrepresented party, WLQ-X-r)**.** *A seat assignment $\boldsymbol{s}$ provides WLQ-X-r if, for every party p, either we have $r_\omega(p, \boldsymbol{s}) \geqslant q_\omega(p)$ or for every seat $t \in \{t' \in \boldsymbol{s}(p^*) \mid p^* \in [n] \setminus \{p\}, r_\omega(p^*, \boldsymbol{s}) > q_\omega(p^*)\}$, it holds that $r_\omega(p, \boldsymbol{s}) + w_t > q_\omega(p)$.*

WLQ-1 states that for each party p, there exists a seat it can additionally receive so as to surpass $q(p)$; WLQ-X states that each party p would surpass $q(p)$ if they would receive any of the additional seats. WLQ-X-r can then be seen as a weakening of WLQ-X where not all seats are considered, but only the seats that have been assigned to parties that have exceeded their representation quota. The intuition behind this requirement is that if one party receives more than their quota, then this is justified by the fact that we could give none of their seats to another party without that party exceeding their quota. Observe that all three axioms are equivalent to lower quota if restricted to unit-weight instances.

Clearly, WLQ-X implies WLQ-X-r, which in turn implies WLQ-1. It turns out that WLQ^o is incomparable with WLQ-X-r (so also WLQ-X).

Example 4.5 (WLQ^o does not imply WLQ-X-r)**.** Consider the following weighted apportionment instance with two parties and four seats:

	# of votes ($n = 2$)	$q_\omega(p)$	$l^o(p)$
Party 1	1	50	2
Party 2	1	50	2

$$\boldsymbol{w} = (97, 1, 1, 1),\ \omega = 100$$

Observe that the seat assignment $\boldsymbol{s} = (1, 1, 2, 2)$ satisfies WLQ^o but not WLQ-X-r since party 1 is sufficiently represented but party 2 could also receive seat 2 (that was assigned to party 1) without surpassing its quota $q_\omega(2) = 50$.

In the other direction, Proposition 4.3 and (the upcoming) Proposition 4.6 give us that WLQ-X does not imply WLQ^o assuming $\mathsf{P} \neq \mathsf{NP}$. In general, We also can show that WLQ-X does not imply WLQ^o, so these axioms are incomparable.

Example 4.6 (WLQ-X does not imply WLQ^o)**.** Consider the following weighted apportionment instance with three parties and three seats:

	# of votes ($n = 6$)	$q_\omega(p)$	$l^o(p)$
Party 1	3	3	3
Party 2	2	2	2
Party 3	1	1	0

$$\boldsymbol{w} = (3, 2, 1),\ \omega = 6$$

Now, take a seat assignment $\boldsymbol{s} = (3, 2, 1)$. This clearly does not satisfy WLQ^o with $r_\omega(1, \boldsymbol{s}) = 1 < l^o(1) = 3$. Note however that the addition of any of the two seats that party 1 did not get assigned would suffice in helping it cross $q_\omega(1) = 3$. So it follows that WLQ-X is satisfied.

We follow up by investigating whether WLQ^o implies WLQ-1.

Proposition 4.5. *WLQ^o implies WLQ-1.*

Proof. Consider some party p with $l^{\#}(p) = \lfloor k \cdot {}^{v_p}/_{n} \rfloor$. Now, suppose we iterate through the seats in non-increasing order of weight and assign h seats to party p such that the following two conditions hold:

1. $h \leqslant l^{\#}(p)$.
2. $\sum_{t \in [h]} w_t \leqslant q_\omega(p) < \sum_{t \in [h+1]} w_t$.

We know that $\sum_{t\in[h]} w_t \leqslant q_\omega(p)$ and that it is obtainable with at most $l^{\#}(p)$ seats, i.e., $\sum_{t\in[h]} w_t \in R(\boldsymbol{w})_{[l^{\#}(p)]}$. Since $l^o(p)$ is the maximal value amongst such weights, we get that $l^o(p) \geqslant \sum_{t\in[h]} w_t$ holds.

Now assume a seat assignment $\boldsymbol{s}$ provides WLQ^o, so we have $r_\omega(p, \boldsymbol{s}) \geqslant l^o(p)$ for party p. Assume that $r_\omega(p, \boldsymbol{s}) < q_\omega(p)$ and consider the seats assigned to party p in $\boldsymbol{s}$. If they are exactly the same h seats as selected above, then seat $h + 1$ works so that:

$$r_\omega(p, \boldsymbol{s}) + w_{h+1} = l^o(p) + w_{h+1} \geqslant \sum_{t\in[h]} w_t + w_{h+1} > q_\omega(p).$$

Otherwise, party p did not receive one of those h seats. Now, since we have $w_t \geqslant w_{h+1}$ for all those seats $t \in [h]$, then $r_\omega(p, \boldsymbol{s}) + w_t > q_\omega(p)$ holds for any seat $t \in [h]$. Thus, WLQ-1 is satisfied. □

We now question whether these axioms are easier to satisfy. First, for the two-party cases, not only can WLQ-X always be provided, but it is possible to do so efficiently.

Proposition 4.6. *For two parties, a seat assignment providing WLQ-X always exists and can be found in polynomial time.*

Proof. Recall that the weight vector $\boldsymbol{w} = (w_1, \ldots, w_k)$ is non-increasing (so w_k is minimal). We devise a method to find a seat assignment $\boldsymbol{s}$ that provides WLQ-X: let $t \in [k]$ be the minimal value such that $\sum_{i=t}^{k} w_i > q_\omega(1)$, and assign seats $t + 1$ to k to party 1 (so $r_\omega(1, \boldsymbol{s}) \leqslant q_\omega(1)$). Then assign the remaining seats to party 2 to obtain seat assignment $\boldsymbol{s}$.

Observe that amongst the seats that party 1 was not assigned, the seat t has the lowest weight w_t. Moreover, we have $r_\omega(1, \boldsymbol{s}) + w_t > q_\omega(1)$. Hence, WLQ-X is satisfied with respect to party 1. Now, let us assess for party 2. Since $r_\omega(1, \boldsymbol{s}) \leqslant q_\omega(1)$ holds by the choice of t and we know $q_\omega(1) + q_\omega(2) =$

$r_\omega(1, \boldsymbol{s}) + r_\omega(2, \boldsymbol{s}) = \omega$ (as $\boldsymbol{s}$ is complete), it must be the case that $r_\omega(2, \boldsymbol{s}) \geqslant q_\omega(2)$. Hence, WLQ-X is also satisfied with respect to party 2. □

In other words, for two parties, we find that WLQ-X is easier to satisfy than WLQ^o. Unfortunately, this does not extend to the case of more than two parties.

Proposition 4.7. *For weighted apportionment instances with three or more parties, a seat assignment providing WLQ-X may not exist.*

Proof. Consider the following weighted apportionment instance with three parties and six seats:

	# of votes ($n = 3$)	$q_\omega(p)$
Party 1	1	33
Party 2	1	33
Party 3	1	33

$$\boldsymbol{w} = (63, 30, 3, 1, 1, 1),\ \omega = 99$$

To achieve WLQ-X, the seats of weight 63 and 30 must be assigned to different parties, say parties 1 and 2. Now, if party 3 is not assigned all of the remaining four seats, then it cannot reach 33 in representation with the addition of any one of these four seats. But, if party 3 is assigned all four of these seats, then party 2 cannot reach 33 in representation by receiving any of the weight-1 seats. Thus, there is no way to provide WLQ-X. □

Observe that, in this example, $q_\omega(p) = l^o(p)$ for all parties p. Hence, we cannot even guarantee that each party is only one seat away from their obtainable lower quota of weights.

While knowing how difficult it is to provide WLQ-X is left for future work, a minor adjustment to the dynamic programming algorithm of Proposition 4.4 delivers the following.

Proposition 4.8. *Given a constant number of parties and the weights in $\boldsymbol{w}$ being polynomial in the input size, finding a seat assignment $\boldsymbol{s}$ that provides WLQ-X can be done in polynomial time, assuming such a seat assignment exists.*

Proof. Consider the dynamic programming algorithm from Proposition 4.4. If we alter the algorithm to also keep track of the smallest seat weight l_p that is assigned to each party p (alongside the sum of seat weights W_p assigned to party p) within the tuples in $\mathcal{W}_i$, then we can use the modified, final set of tuples $\mathcal{W}_k$ to check in polynomial time if WLQ-X is satisfied. □

Off the back of the mostly negative results regarding WLQ^o and WLQ-X, we take aim at the weaker requirement of WLQ-X-r, and find the following positive result using the Greedy$_\omega$ method.

Theorem 4.9. *The Greedy$_\omega$ method satisfies WLQ-X-r.*

Proof. Assume that WLQ-X-r is violated by some seat assignment $\boldsymbol{s}$ returned by the Greedy$_\omega$ method. Let party p_x be the party that witnesses it, i.e., it holds that $r_\omega(p_x, \boldsymbol{s}) < r_\omega(p_x, \boldsymbol{s}) + w_t \leqslant q_\omega(p_x)$ for some seat $t \in \{t' \in \boldsymbol{s}(p^*) \mid p^* \in [n] \setminus \{p\}, r_\omega(p^*, \boldsymbol{s}) > q_\omega(p^*)\}$.

As party p_x has less than $q_\omega(p_x)$ in representation, there must be a party p_y where $r_\omega(p_y, \boldsymbol{s}) > q_\omega(p_y)$. Let h be the round after which party p_y has more than $q_\omega(p_y)$ in representation (so party p_y was assigned seat h). By choice of round h, we have $g^\omega_{p_y}(h) + w_h > q_\omega(p_y)$ and hence, it holds that $w_h > q_\omega(p_y) - g^\omega_{p_y}(h)$. So we have that $w_h > q_\omega(p_y) - g^\omega_{p_y}(h)$, and since party p_x was not assigned seat h, we know that $w_h > q_\omega(p_y) - g^\omega_{p_y}(h) \geqslant q_\omega(p_x) - g^\omega_{p_x}(h)$. It then follows that $q_\omega(p_x) < g^\omega_{p_x}(h) + w_h \leqslant r_\omega(p_x, \boldsymbol{s}) + w_h$.

So this seat h is enough for party p_x to reach their quota with the same holding for all seats assigned to party p_y in prior rounds (as seats are assigned in non-increasing order). □

Recall that in standard apportionment, LRM is known to satisfy LQ (Balinski, 1982). Theorem 4.9 further justifies the Greedy$_\omega$ method as a weighted proxy of LRM. Recall, also, that in the standard setting D'Hondt satisfies LQ as well; we find that D'Hondt$_\omega$, now, satisfies WLQ-X-r.

Theorem 4.10. *The D'Hondt$_\omega$ method satisfies WLQ-X-r.*

Proof. For a seat assignment $\boldsymbol{s}$ returned by D'Hondt$_\omega$, for the sake of contradiction, assume that there is a party p_x such that there exists some $t \in \{t' \in \boldsymbol{s}(p^*) \mid p^* \in [n] \setminus \{p\}, r_\omega(p^*, \boldsymbol{s}) > q_\omega(p^*)\}$ such that $r_\omega(p_x, \boldsymbol{s}) < r_\omega(p_x, \boldsymbol{s}) + w_t \leqslant q_\omega(p_x)$.

Thus, we know that:

$$\frac{v_{p_x}}{g^\omega_{p_x}(k) + w_t} \geqslant \frac{v_{p_x}}{q_\omega(p_x)}$$

for some $t \in \{t' \in \boldsymbol{s}(p^*) \mid p^* \in [n] \setminus \{p\}, r_\omega(p^*, \boldsymbol{s}) > q_\omega(p^*)\}$, where $g^\omega_{p_x}(k)$ is the total weight assigned to party p_x at D'Hondt$_\omega$'s conclusion. This gives us the following:

$$\frac{v_{p_x}}{g^\omega_{p_x}(k) + w_t} \leqslant \frac{v_{p_x}}{q_\omega(p_x)} = \frac{v_{p_x}}{\omega \cdot {}^{v_{p_x}}\!/_{n}} = \frac{n}{\omega} \tag{4.1}$$

During D'Hondt$_\omega$, there must be some round h where some party $p_y \neq p_x \in [m]$ is assigned weight w_h such that:

$$\frac{n}{\omega} > \frac{v_{p_y}}{g^\omega_{p_y}(h) + w_h}.$$

Assume otherwise, and that for every party $p \in [m] \setminus \{p_x\}$, it holds that ${}^{v_p}\!/_{g^\omega_p(k)} \geqslant {}^{n}\!/_{\omega}$ after D'Hondt$_\omega$'s k rounds. Then we have that $\omega \cdot {}^{v_p}\!/_{n} \geqslant g^\omega_p(k)$ for all $p \in [m] \setminus \{p_x\}$. Summing over all parties with $\omega \cdot {}^{v_{p_x}}\!/_{n} > g^\omega_{p_x}(k)$ for party p_x, we get $\sum_{p \in [m]} \omega \cdot {}^{v_p}\!/_{n} = \omega > g^\omega_{p_x}(k) + \sum_{p \in [m] \setminus \{p_x\}} g^\omega_p(k)$ which means D'Hondt$_\omega$ did not assign all of the weight, contradicting its definition. So, there must exist some round h where for some party p_y, we have:

$$\frac{n}{\omega} > \frac{v_{p_y}}{g^\omega_{p_y}(h) + w_h} \tag{4.2}$$

Since weight w_h was assigned to party p_y in round h, and not party p_x, then we have that:

$$\frac{v_{p_y}}{g^\omega_{p_y}(h) + w_h} \geqslant \frac{v_{p_x}}{g^\omega_{p_x}(h) + w_h}$$

where $h \in \{t' \in \boldsymbol{s}(p^*) \mid p^* \in [n] \setminus \{p\}, r_\omega(p^*, \boldsymbol{s}) > q_\omega(p^*)\}$. And also considering the fact that $g^\omega_{p_x}(h) \leqslant g^\omega_{p_x}(k)$, it follows that:

$$\frac{v_{p_y}}{g^\omega_{p_y}(h) + w_h} \geqslant \frac{v_{p_x}}{g^\omega_{p_x}(h) + w_h} \geqslant \frac{v_{p_x}}{g^\omega_{p_x}(k) + w_h} \tag{4.3}$$

Putting equations (4.1), (4.2) and (4.3) together, it follows that:

$$\frac{n}{\omega} > \frac{v_{p_y}}{g^\omega_{p_y}(h) + w_h} \geqslant \frac{n}{\omega}.$$

This is a contradiction, so no such party p_x can exist. Note that we considered a seat weight w_h assigned to some party p_y in round h, such that p_y surpasses its quota. And such a weight w_h is sufficient in aiding party p_x in reaching $q_\omega(p_x)$. This holds for all seats assigned to party p_y before round h (as such seats h^* have weight $w_{h^*} \geqslant w_h$), and also those seats assigned to party p_y after round h (as such seats h^* are only assigned to party p_y, and not some party p_x below its quota $q_\omega(p_x)$ in that round, if the weight w_{h^*} would lead to party p_x reaching said quota). □

This improves on a result of Chakraborty et al. (2021) stating that D'Hondt satisfies an axiom, which is weaker than WLQ-X-r, called *WPROP1* (see Theorem 4.9 by Chakraborty et al. (2021)).[2] This is important as the "up to any" properties are, in many scenarios, much stronger than the equivalent "up to one" properties, in particular if the values of objects vary a lot. Consider the application of our rules to the allocation of non-liquid assets in a bankruptcy. We may use the approximate monetary value of the assets as a weight. In this case, we might have a few very valuable assets (e.g., a house or other property), and other assets of much lower values (e.g., furniture). In such a case "up to one" properties can become essentially meaningless, while "up to any" properties are still meaningful. This scenario shall be illustrated with the upcoming Example 4.7.

Crucially, our stronger result does not only stem from our restricted setting but also from our use of a version of D'Hondt that takes weights into account as the standard D'Hondt used by Chakraborty et al. (2021) does not satisfy WLQ-X-r in our setting. This can be seen in the following example.

Example 4.7 (Use of standard D'Hondt method fails WLQ-X-r)**.** Consider the following weighted apportionment instance with two parties and three seats:

	# of votes ($n = 12$)	$q_\omega(p)$
Party 1	10	10
Party 2	2	2

$$\boldsymbol{w} = (10, 1, 1),\ \omega = 12$$

The standard D'Hondt method assigns all three seats to party 1 as, in the D'Hondt method's three respective rounds, party 1 has the ratios 10, 5, and 2.5 versus party 2's ratio of 2 in all three rounds. Thus, party 2 has representation of 0 from the resulting seat assignment and none of weight-1 seats are enough to add so that party 2 exceeds its quota of $q_\omega(2) = 2$. However, observe that the seat assignment determined by standard D'Hondt provides WLQ-1 while our WSAM D'Hondt$_\omega$ returns the seat assignment $\boldsymbol{s} = (1, 2, 2)$ for this same election instance and this seat assignment $\boldsymbol{s}$ not only provides WLQ-X-r, but it is a much fairer outcome from an intuitive standpoint.

This further illustrates the importance of our focus on the weighted apportionment model as we need to use the specific properties of our model to define new

[2] WPROP1 is similar to WLQ-1, but WPROP1 is defined with a weak inequality in the condition on the existence of some seat of sufficient weight.

voting rules, namely the WSAMs, that can satisfy the stronger "up to any" properties. We also point out that Theorems 4.9 and 4.10 rely on the seats being assigned using a non-increasing order of seat weight, which contrasts with the results for the weaker axiom WLQ-1, where an arbitrary ordering would suffice.

Now, regarding the other divisor method, Adams$_\omega$ fails WLQ-1 (and thus WLQ-X-r), as it is known to violate LQ, which, as mentioned above, is equivalent to WLQ-1 in the unit-weight case (Balinski, 1982).

Example 4.8 (Adams$_\omega$ and rev-Adams$_\omega$ fail WLQ-1)**.** Consider the following weighted apportionment instance with four parties and four seats:

	# of votes ($n = 12$)	$q_\omega(p)$
Party 1	9	3
Party 2	1	$1/3$
Party 3	1	$1/3$
Party 4	1	$1/3$

$$\boldsymbol{w} = (1, 1, 1, 1),\ \omega = 4$$

Adams$_\omega$ gives each party $p \in [4]$ an initial ratio of $ratio_p = \infty$, so each party receives exactly one seat. Consequently, party 1 is two seats away from its weighted quota.

4.3 Weighted Upper Quota

While lower quota guarantees that each party receives at least as many seats as it deserves, we also want to prevent parties form getting more seats (and thus influence) than appropriate, a notion formalised by the *upper quota* (UQ) property. In the standard apportionment setting, UQ states that a party p amassing v_p of the n votes should receive *at most* $\lceil k \cdot v_p/n \rceil$ of the k seats (Balinski, 1982).

As with the lower quota, it is easy to check that there is no hope of satisfying the naïve weighted upper quota notion defined by $\lceil \omega \cdot v_p/n \rceil$. Consequently, we approach the definition of a *weighted upper quota* in a similar fashion to our definition of the obtainable weighted lower quota $l^o(p)$.

Specifically, we incorporate 'upper quota of seats' in defining a weighted upper quota for party p that is obtainable. So for a party p, we may define its weighted upper quota to be the following where $u^{\#}(p) = \lceil k \cdot v_p/n \rceil$:

$$u_o^*(p) = \max\{r \in R(\boldsymbol{w})_{[u^{\#}(p)]} \mid r \leqslant \min\{w \in R(\boldsymbol{w})_{[k]} \mid w \geqslant q_\omega(p)\}\}.$$

Note that this definition is more involved than, and not exactly symmetric to, that of $l^o(p)$. To see why, observe that to obtain a more obvious counterpart to the definition of $l^o(p)$, we should fix $u^*_o(p)$ to be:

$$\min\{r' \in R(\boldsymbol{w})_{[u^{\#}(p)]} \mid r' \geqslant q_\omega(p)\}$$

However, we cannot do so due to the fact that such a value may not exist, as all values in $R(\boldsymbol{w})_{[u^{\#}(p)]}$ may be less than $q_\omega(p)$. Thus, we define $u^*_o(p)$ to be the largest value in $R(\boldsymbol{w})_{[u^{\#}(p)]}$ that does not exceed an upper bound of $\min\{w \in R(\boldsymbol{w})_{[k]} : w \geqslant q_\omega(p)\}$ which always exists.

This seems a natural counterpart to the obtainable weighted lower quota $l^o(p)$, but we can show that there exist weighted apportionment instances where the sum of obtainable weighted upper quotas is less than the total weight, and therefore, any reasonable upper quota axiom based on $u^*_o(p)$ will be violated. The following details such a weighted apportionment instance.

Example 4.9. Consider the following weighted apportionment instance with five parties and six seats:

	# of votes ($n = 100$)	$u^{\#}(p)$	$q_\omega(p)$	$u^*_o(p)$
Party 1	40	2	4	2
Party 2	15	1	1.5	1
Party 3	15	1	1.5	1
Party 4	15	1	1.5	1
Party 5	15	1	1.5	1

$$\boldsymbol{w} = (5, 1, 1, 1, 1, 1),\ \omega = 10$$

Party 1 has an upper quota of seats of $u^{\#}(1) = 2$ and all other parties have an upper quota of seats of only 1. Thus, the obtainable upper quotas taking this ‘upper quota of seats’ into account would be $u^*_o(1) = 2$ for party 1 and $u^*_o(p) = 1$ for the other parties $p \in \{2, \ldots, 5\}$, which sums to $\sum_{p\in[m]} u^*_o(p) = 7 < 10 = \omega$.

We thus define the obtainable weighted upper quota, that does not incorporate the ‘upper quota of seats’, as follows:

$$u^o(p) = \min\left\{r \in R(\boldsymbol{w})_{[k]} \mid r \geqslant q_\omega(p)\right\}.$$

Now we can define the axiom WUQ^o based on this notion u^o of an obtainable upper quota value.

Definition 4.9 (Obtainable Weighted Upper Quota, WUQ^o)**.** *A seat assignment* $\boldsymbol{s}$ *provides WUQ*^{o} *if for every party* p *it is the case that* $r_\omega(p, \boldsymbol{s}) \leqslant u^o(p)$*.*

Note that for unit-weight election instances WUQ^o reduces to the standard UQ property. For this definition of the obtainable upper quota, we achieve essentially the same results as for WLQ^o. The following four results can be seen as direct counterparts to those shown for WLQ^o.

Proposition 4.11. *For every weighted apportionment instance with two parties there exists a seat assignment that provides WUQ*^{o}*.*

Proof. Consider a weighted apportionment instance with parties 1 and 2. By the definition of the obtainable weighted upper quota, there exists a set of seats T such that $\sum_{t \in T} w(t) = u^o(1)$. Now, let $\boldsymbol{s}$ be a seat assignment such that party 1 gets all seats in T, and party 2 all other seats. By definition, $r_\omega(1, \boldsymbol{s}) = u^o(1)$ for party 1. Moreover, for party 2, we have:

$$\begin{aligned} r_\omega(2, \boldsymbol{s}) &= \omega - r_\omega(1, \boldsymbol{s}) = \omega - u^o(1) \\ &\leqslant \omega - \omega \cdot \frac{v_1}{n} = \omega \cdot \frac{n - v_1}{n} = \omega \cdot \frac{v_2}{n} \leqslant u^o(2). \end{aligned}$$

Hence, WUQ^o is satisfied. □

Proposition 4.12. *There are weighted apportionment instances where no complete seat assignment provides WUQ*^{o}*.*

Proof. We use the weighted apportionment instance familiar from the proof of Theorem 4.2. Consider the following weighted apportionment instance with three parties:

	# of votes ($n = 3$)	$u^o(p)$
Party 1	1	2
Party 2	1	2
Party 3	1	2

$$\boldsymbol{w} = (3, 2, 1),\ \omega = 6$$

Providing each party with at most 2 in representation cannot be achieved if we need to assign all three seats. □

Proposition 4.13. *If there exists a polynomial-time algorithm $\mathbb{A}$ that finds a seat assignment $\boldsymbol{s}$ that provides WUQ^o whenever such a seat assignment exists, then* P = NP. *This holds even when restricted to the case where there are only two parties.*

Proof (sketch). Consider the reduction from the Partition problem in Proposition 4.3. Observe that in the weighted apportionment instance constructed in this reduction, the obtainable weighted lower quota $l^o(p)$ for each party $p \in [2]$ is exactly equal to the obtainable weighted upper quota $u^o(p)$. So the same arguments provided for WLQ^o work for WUQ^o. □

Proposition 4.14. *Given a constant number of parties and the weights in $\boldsymbol{w}$ being polynomial in the input size, finding a seat assignment $\boldsymbol{s}$ that provides WUQ^o can be done in polynomial time, assuming such a seat assignment exists.*

Proof. Consider the dynamic programming algorithm from Proposition 4.4. Observe that the same algorithm can be deployed to find WUQ^o-providing seat assignments with the following, simple modification: once $\mathcal{W}_k$ is computed, the algorithm checks each tuple $(W_1, \dots, W_m)$ to determine whether, for every $p \in [m]$, it holds that $W_p \leqslant u^o(p)$ instead of whether $W_p \geqslant l^o(p)$. And we can compute $u^o(p)$ in polynomial time due to (i) the assumption on the weights in $\boldsymbol{w}$, and (ii) observing this task is also equivalent to solving the Subset Sum problem. □

Given that the results for WUQ^o are as negative as the results for WLQ^o, we next consider the up to one/any relaxations that allowed us to define satisfiable lower-quota axioms.

Definition 4.10 (WUQ up to one seat, WUQ-1)**.** *A seat assignment $\boldsymbol{s}$ provides WUQ-1 if, for every party p, either $r_\omega(p, \boldsymbol{s}) \leqslant q_\omega(p)$ or there exists some seat $t \in \boldsymbol{s}(p)$ such that $r_\omega(p, \boldsymbol{s}) - w_t < q_\omega(p)$.*

Definition 4.11 (WUQ up to any seat, WUQ-X)**.** *A seat assignment $\boldsymbol{s}$ provides WUQ-X if, for every party p, either $r_\omega(p, \boldsymbol{s}) \leqslant q_\omega(p)$ or for every seat $t \in \boldsymbol{s}(p)$, it holds that $r_\omega(p, \boldsymbol{s}) - w_t < q_\omega(p)$.*

WUQ-X states that, for every party p, disregarding any seat it received would take it below $q(p)$ while for WUQ-1, there need only exist one such seat assigned to party p to take it below $q(p)$. With unit weights, these axioms are exactly UQ. Observe that there is no natural way of defining a counterpart to WLQ-X-r, as the items we remove from each party must be assigned to them.

WUQ-X clearly implies WUQ-1, but how do these two axioms relate to WUQ^o? Here, we see the first difference between upper and lower quota, as WUQ^o does not only imply WUQ-1 but even WUQ-X (whereas Example 4.5 showed that WLQ^o does not imply WLQ-X).

Proposition 4.15. *WUQ^o implies WUQ-X.*

Proof. Assume WUQ-X is violated by a seat assignment $\boldsymbol{s}$. Then there exists a party p such that for every seat $t \in \boldsymbol{s}(p_x)$, it holds that $r_\omega(p, \boldsymbol{s}) > r_\omega(p, \boldsymbol{s}) - w_t \geqslant q_\omega(p)$. But then this means that the value $r_\omega(p, \boldsymbol{s}) - w_t$ is an achievable weight representation value that is at least as large as $q_\omega(p)$, i.e., $r_\omega(p, \boldsymbol{s}) - w_t \in \{r \in R(\boldsymbol{w})_{[k]} \mid r \geqslant q_\omega(p)\}$.

Since $u^o(p)$ is the smallest of the weights in $\{r \in R(\boldsymbol{w})_{[k]} \mid r \geqslant q_\omega(p)\}$, we get $r_\omega(p, \boldsymbol{s}) > r_\omega(p, \boldsymbol{s}) - w_t \geqslant u^o(p)$ and hence, WUQ^o is also violated. □

In the other direction, here is an example that shows that WUQ-X does not imply WUQ^o.

Example 4.10 (WUQ-X does not imply WUQ^o)**.** Consider the following weighted apportionment instance with three parties and three seats:

	# of votes ($n = 6$)	$q_\omega(p)$	$u^o(p)$
Party 1	3	3	3
Party 2	2	2	2
Party 3	1	1	1

$$\boldsymbol{w} = (3, 2, 1),\ \omega = 6$$

Now, take a seat assignment $\boldsymbol{s} = (3, 2, 1)$. This does not satisfy WUQ^o as we get that $r_\omega(3, \boldsymbol{s}) = 3 > u^o(3) = 1$, but note that the removal of the seat that party 3 was assigned would suffice in helping it go below $q_\omega(3) = 1$.

Now, we ask whether we can satisfy these weaker upper-quota axioms. A natural candidate for doing so is the Adams$_\omega$ method as it is known to satisfy UQ in unit-weight cases (Balinski, 1982). Indeed, Adams$_\omega$ even satisfies the stronger notion of WUQ-X. This is in stark contrast to WLQ-X, which was not satisfiable in general.

Theorem 4.16. *Adams$_\omega$ satisfies WUQ-X.*

Proof. Assume there exists a party p_x that receives more representation than $q_\omega(p_x)$, otherwise, WUQ-X is satisfied by definition. Let t be the round of Adams$_\omega$ such that party p_x was assigned seat t and after which $g^\omega_{p_x}(t) > q_\omega(p_x)$, and thus, we have that:

$$g^\omega_{p_x}(t) - w_t < q_\omega(p_x) = \omega \cdot \frac{v_{p_x}}{n}.$$

We now show that party p_x does not receive any more seats. Consider some round $t^* > t$ where seat t^* is to be assigned. We know that $\sum_{p \in [m]} q_\omega(p) = \omega$ holds alongside the following:

$$g^\omega_{p_x}(t^*) \geqslant g^\omega_{p_x}(t) > \omega \cdot \frac{v_{p_x}}{n}.$$

As at least one seat, namely seat t^*, has not been assigned yet, there must be a party p_y such that $g^\omega_{p_y}(t^*) < q_\omega(p_y) = \omega \cdot {}^{v_{p_y}}\!/_{n}$. And it follows that:

$$\frac{v_{p_x}}{g^\omega_{p_x}(t^*)} < \frac{v_{p_x}}{\omega \cdot {}^{v_{p_x}}\!/_{n}} = \frac{n}{\omega} = \frac{v_{p_y}}{\omega \cdot {}^{v_{p_y}}\!/_{n}} < \frac{v_{p_y}}{g^\omega_{p_y}(t^*)}.$$

Hence, party p_y has a strictly better ratio than party p_x, so Adams$_\omega$ does assign seat t^* to the latter party. So, we know that removing seat t would be enough for party p_x to fall below their weighted quota $q_\omega(p_x)$, and thus, up to this point, we have proven that Adams$_\omega$ satisfies WUQ-1.

To show that Adams$_\omega$ satisfies WUQ-X, we need the following additional argument: for all the seats $j < t$ that Adams$_\omega$ assigned to party p_x prior to it being assigned seat t, we have that $w_j \geqslant w_t$ and hence, $g^\omega_{p_x}(t) - w_j \leqslant g^\omega_{p_x}(t) - w_t < q_\omega(p_x)$. So removing any one of these seats will suffice to ensure that party p_x does not exceed $q_\omega(p_x)$. □

We now expand our focus to the following *envy-freeness* axioms (envy-freeness is a well-known fairness notion in the fair division literature).

Definition 4.12 (Weighted envy-freeness up to any seat, WEF-X)**.** *A seat assignment $\boldsymbol{s}$ provides WEF-X if for any two parties p_x, p_y, it holds for every seat $t \in \boldsymbol{s}(p_y)$ that*

$$\frac{r_\omega(p_x, \boldsymbol{s})}{v_{p_x}} \geqslant \frac{r_\omega(p_y, \boldsymbol{s}) - w_t}{v_{p_y}}.$$

We can also define the following weakening of WEF-X.

Definition 4.13 (Weighted envy-freeness up to one seat, WEF-1)**.** *A seat assignment $\boldsymbol{s}$ provides WEF-1 if for any two parties p_x, p_y, there exists some seat*

$t \in \boldsymbol{s}(p_y)$ *such that:*

$$\frac{r_\omega(p_x, \boldsymbol{s})}{v_{p_x}} \geqslant \frac{r_\omega(p_y, \boldsymbol{s}) - w_t}{v_{p_y}}.$$

In our setting, both WEF-X and WEF-1 ensure that no party prefers the representation afforded to another party. Conceptually, this is similar to the upper-quota notion that states that no party is represented more than it truly deserves. To provide a formal connection between envy-freeness and upper quota, we prove that WUQ-X follows from WEF-X.

Proposition 4.17. *WEF-X implies WUQ-X.*

Proof. Observe that if WUQ-X is violated, then there exists a party p_x such that $r_\omega(p_x, \boldsymbol{s}) - w_t \geqslant q_\omega(p_x) = \omega \cdot {}^{v_{p_x}}/_n$ for some $t \in \boldsymbol{s}(p_x)$. On the other hand, for WEF-X to hold, for every party $p_y \in [n] \setminus \{p_x\}$ and every $t \in \boldsymbol{s}(p_x)$ it must be the case that

$$\frac{r_\omega(p_y, \boldsymbol{s})}{v_{p_y}} \geqslant \frac{r_\omega(p_x, \boldsymbol{s}) - w_t}{v_{p_x}}.$$

These inequalities imply that $r_\omega(p, \boldsymbol{s}) \geqslant \omega \cdot {}^{v_p}/_n = q_\omega(p)$ for every party $p \in [n]$, which is not possible if party p_x exceeded its quota $q_\omega(p_x)$. □

And then it is clear to see that WEF-1 implies WUQ-1 and as the proof is analogous to that of Proposition 4.17, it is omitted here.

Proposition 4.18. *WEF-1 implies WUQ-1.*

On the other hand, it is straightforward to see that this weaker WEF-1 does not imply WUQ-X.

Example 4.11 (WEF-1 does not imply WUQ-X)**.** Consider the following weighted apportionment instance with two parties and three seats:

	# of votes ($n = 2$)	$q_\omega(p)$
Party 1	1	7
Party 2	1	7

$$\boldsymbol{w} = (11, 2, 1),\ \omega = 14$$

Now, take a seat assignment $\boldsymbol{s} = (1, 2, 1)$. This does not satisfy WUQ-X as we get that $r_\omega(1, \boldsymbol{s}) - w_3 = 12 - 1 > 7$, but note that the removal of seat 1 from party 1 would remove any envy from party 2.

We also show that an axiom that is studied by Chakraborty et al. (2021), and is a weakening of WEF-1, does not imply WUQ-1. First, we define this weaker axiom.

Definition 4.14 (Weak weighted envy-freeness up to one seat, WWEF-1)**.** *We say a seat assignment $\boldsymbol{s}$ provides WWEF-1 if for any two parties p_x, p_y, there exists some seat $t \in \boldsymbol{s}(p_y)$ such that at least one of inequalities 4.4 and 4.5, given below, hold:*

$$\frac{r_\omega(p_x, \boldsymbol{s})}{v_{p_x}} \geqslant \frac{r_\omega(p_y, \boldsymbol{s}) - w_t}{v_{p_y}}. \tag{4.4}$$

$$\frac{r_\omega(p_x, \boldsymbol{s}) + w_t}{v_{p_x}} \geqslant \frac{r_\omega(p_y, \boldsymbol{s})}{v_{p_y}}. \tag{4.5}$$

Now, here is an example illustrating that WWEF-1 does not imply WUQ-1.

Example 4.12 (WWEF-1 does not imply WUQ-1)**.** Consider the following weighted apportionment instance with 101 parties and four seats:

	# of votes ($n = 200$)	$q_\omega(p)$
Party 1	100	2
Party 2	1	0.02
⋮	⋮	⋮
Party 101	1	0.02

$$\boldsymbol{w} = (1, 1, 1, 1),\ \omega = 4$$

D'Hondt$_\omega$ assigns all seats to party 1 which violates WUQ-1. Now, to see that WWEF-1 is satisfied, observe that condition 4.5 of the definition of WWEF-1 is satisfied. This holds as for all parties $p \in \{2, \dots, 101\}$, assigning them a weight-1 seat is enough such that ${(r_\omega(p,\boldsymbol{s})+w_t)}/{v_p} = {(0+1)}/{1} = 1 \geqslant {r_\omega(1,\boldsymbol{s})}/{v_1} = {4}/{100}$.

Now, we show that Adams$_\omega$ satisfies WEF-X.

Theorem 4.19. *The Adams$_\omega$ method satisfies WEF-X.*

Proof. Suppose there are two parties p_x, p_y with:

$$\frac{r_\omega(p_y, \boldsymbol{s})}{v_{p_y}} < \frac{r_\omega(p_x, \boldsymbol{s})}{v_{p_x}}$$

That is to say that party p_y envies party p_x. Now, consider the last seat t that was assigned to party p_x by Adams$_\omega$ in round h. Since this seat was assigned to party p_x, we have that:

$$\frac{v_{p_x}}{g^{\omega}_{p_x}(h)} \geqslant \frac{v_{p_y}}{g^{\omega}_{p_y}(h)}.$$

And as seat t was the last seat assigned to party p_x, we get:

$$\frac{r_\omega(p_y, \boldsymbol{s})}{v_{p_y}} \geqslant \frac{g^{\omega}_{p_y}(h)}{v_{p_y}} \geqslant \frac{g^{\omega}_{p_x}(h)}{v_{p_x}} = \frac{r_\omega(p_x, \boldsymbol{s}) - w_t}{v_{p_x}}.$$

So, removing seat t from party p_x leads to party p_y no longer envying party p_x, and since all seats assigned to party p_x prior to seat t have weight at least as large as seat t, removing any of these seats is sufficient to remove party p_y's envy. □

This result improves on that of Chakraborty et al. (2021) that shows that the standard Adams method can be used to achieve WEF-1. It would be interesting to see if one can generalise our result to their setting. In this regard, note that finding a rule that satisfies WEF-X in the setting of Chakraborty et al. (2021) would imply the existence of *EF-X* allocations in the standard fair division setting, which is considered one of the major open questions in fair division (Amanatidis et al., 2022).

Since Adams$_\omega$ does not satisfy WLQ-1, can envy-freeness and lower quota be satisfied at the same time? This is not possible, as WEF-1 and WLQ-1 are incompatible.

Proposition 4.20. *WEF-1 and WLQ-1 are incompatible.*

Proof. Consider the following weighted apportionment instance with three parties and three seats:

	# of votes ($n = 6$)	$q_\omega(p)$
Party 1	8	≈ 15.2
Party 2	2	≈ 3.8
Party 3	1	≈ 1.9

$$\boldsymbol{w} = (8, 7, 6),\ \omega = 21$$

To satisfy WLQ-1, two of the seats must be assigned to party 1. However, this leads to one of parties 2 and 3 having no seats: the party without a seat will envy party 1 and be a witness to a violation of WEF-1. □

As Adams$_\omega$ does not satisfy WLQ-1, we ask whether upper- and lower-quota axioms can be satisfied at the same time. Let us turn to D'Hondt$_\omega$ and Greedy$_\omega$ since they satisfy WLQ-1. D'Hondt$_\omega$ does not satisfy UQ in the unit-weight case so it cannot satisfy WUQ-1 (the same holds for rev-D'Hondt$_\omega$). Example 4.12 shows this well-known fact that D'Hondt$_\omega$ violates UQ even in the unit-weight setting by an arbitrary number of seats (Balinski, 1982).

Example 4.13 (D'Hondt$_\omega$ and rev-D'Hondt$_\omega$ fail WUQ-1)**.** Consider again the following weighted apportionment instance with 101 parties and four seats:

	# of votes ($n = 200$)	$q_\omega(p)$
Party 1	100	2
Party 2	1	0.02
⋮	⋮	⋮
Party 101	1	0.02

$$\boldsymbol{w} = (1, 1, 1, 1),\ \omega = 4$$

Both D'Hondt$_\omega$ and rev-D'Hondt$_\omega$ assign all seats to party 1 which violates WUQ-1.

The Greedy$_\omega$ method however, is a contender with its connection to LRM, known to satisfy UQ (Balinski, 1982). As it satisfies WLQ-1 it cannot satisfy WEF-1, but we find that it does satisfy WUQ-X.

Theorem 4.21. *Greedy$_\omega$ satisfies WUQ-X.*

Proof. Take a seat assignment $\boldsymbol{s}$ constructed by Greedy$_\omega$. Assume there is a party p_x that received more representation than $q_\omega(p_x)$, i.e., $r_\omega(p_x, \boldsymbol{s}) = g^\omega_{p_x}(k) > q_\omega(p_x)$, otherwise, WUQ-X is satisfied. Let t be the round after which $g^\omega_{p_x}(t) > q_\omega(p_x)$ holds, and so $g^\omega_{p_x}(t) - w_t < q_\omega(p_x)$ also holds. We argue that party p_x does not get assigned any seat $t^* > t$. Observe that in every round $t' \in [k]$, there always exists a party p_y such that $q_\omega(p_y) - g^\omega_{p_y}(t') \geqslant 0$ as we have $\sum_{p \in [m]} q_\omega(p) = \omega$. It then follows, for every round $t^* > t$, that:

$$q_\omega(p_y) - g^\omega_{p_y}(t^*) \geqslant 0 > q_\omega(p_x) - g^\omega_{p_x}(t^*).$$

Hence, party p_x cannot be assigned seat t^*. Given this we know that Greedy$_\omega$ satisfies WUQ-1. However, for all seats $j < t$ assigned to party p_x up to the assignment of seat t, we have $w_j \geqslant w_t$. So, it follows that $g^{\omega}_{p_x}(t) - w_j \leqslant g^{\omega}_{p_x}(t) - w_t < q_{\omega}(p_x)$ and this means that removing any seat assigned to arty p suffices for this party to not be above its weighted quota and thus, WUQ-X is also satisfied. □

These results highlight the importance of the order that seats are assigned, with the non-increasing order yielding more positive results than the arbitrary ordering for the WSAMs.

4.4 House Monotonicity for Weights

Why focus on D'Hondt$_{\omega}$ or Adams$_{\omega}$ if the Greedy$_{\omega}$ method satisfies both WLQ-X-r and WUQ-X? The answer lies with *house monotonicity.*

Beyond the aim of proportional representation, this is an additional desideratum in the apportionment literature, that states that an increase in the number of available seats should not harm any party in terms of the number of seats that it is assigned. This property is important in order to avoid situations such as the often cited *Alabama Paradox*, and its failure has, historically, been the cause of much political animosity (Szpiro, 2010).

Definition 4.15 (House Monotonicity, HM)**.** *Take apportionment instances* $(\boldsymbol{v}, k)$ *and* $(\boldsymbol{v}, k+1)$*. An apportionment method* M *is house monotone if for seat assignments* $\boldsymbol{s} = M(\boldsymbol{v}, k)$ *and* $\boldsymbol{s}^* = M(\boldsymbol{v}, k+1)$*, it holds that* $r(p, \boldsymbol{s}^*) \geqslant r(p, \boldsymbol{s})$ *for every party* $p \in [m]$*.*

It is known that LRM satisfies LQ and UQ in standard apportionment but fails HM, whereas divisor methods satisfy it (this was also mentioned in Chapter 2) (Balinski, 1982; Pukelsheim, 2014). Note that Chakraborty et al. (2021) also study house monotonicity in their setting. As our WSAMs differ from the rules studied by Chakraborty et al. (2021), their results on house monotonicity do not apply to our methods.

We look at adaptations of HM to our weighted setting. The first is most faithful to the intuition of HM. We first consider the following strong generalisation of house monotonicity.

Definition 4.16 (Full House Monotonicity, full-HM)**.** *We say a WSAM* M_{ω} *satisfies full-HM if for every weighted apportionment instance* $(\boldsymbol{v}, \boldsymbol{w})$ *and every* $w^* \in \mathbb{N}_{\geqslant 1}$ *such that* $\boldsymbol{w}^* = (w')_{w' \in W}$ *is a non-increasing weight vector where* $W = \{w \in \boldsymbol{w}\} \cup \{w^*\}$*, it holds for* $\boldsymbol{s} = M_{\omega}(\boldsymbol{v}, \boldsymbol{w})$ *and* $\boldsymbol{s}^* = M_{\omega}(\boldsymbol{v}, \boldsymbol{w}^*)$ *that* $r_{\omega}(p, \boldsymbol{s}^*) \geqslant r_{\omega}(p, \boldsymbol{s})$ *for every party* $p \in [m]$*.*

Firstly, we can always satisfy full-HM in a trivial manner by always assigning all of the seats to the largest party (with lexicographic tie-breaking for example). Now, full-HM can also always be satisfied in a more natural fashion by using, for example, the non-weighted divisor methods (Chakraborty et al., 2021). Of course, as we have seen before, these methods do not satisfy the strongest possible quota axioms WLQ-X-r and WUQ-X. Of course, this is arguably not a desirable WSAM and, in particular, violates all quota axioms we have considered so far.

As with much of the work up to this point, we see that dealing with weights provides additional difficulty in importing properties to our setting. For unit-weight elections, all divisor methods satisfy house monotonicity (Balinski, 1982; Pukelsheim, 2014) but unfortunately, our adapted versions of Adams and D'Hondt fail to do the same for full-HM.

Proposition 4.22. *Adams$_\omega$ and D'Hondt$_\omega$ fail full-HM.*

Proof. Let us first consider Adams$_\omega$ and consider an instance with three parties and there initially being four seats to be assigned:

	# of votes ($n = 12$)
Party 1	5
Party 2	5
Party 3	2

For a weight vector $\boldsymbol{w} = (8, 8, 3, 2)$, the seat assignment returned by Adams$_\omega$ is $\boldsymbol{s} = (1, 2, 3, 3)$, assuming that ties are broken according to the ordering of $\boldsymbol{v}$. Thus, we have $r_\omega(3, \boldsymbol{s}) = 5$ for party 3. Suppose a weight-4 seat is added to $\boldsymbol{w}$ so as to obtain the weight vector $\boldsymbol{w}^* = (8, 8, 4, 3, 2)$. Then Adams$_\omega$ returns $\boldsymbol{s}^* = (1, 2, 3, 1, 2)$. So party 3 go from $r_\omega(3, \boldsymbol{s}) = 5$ to $r_\omega(3, \boldsymbol{s}^*) = 4$.

Now, consider the following weighted apportionment instance with three parties and only two seats to begin with:

	# of votes ($n = 41$)
Party 1	21
Party 2	10
Party 3	10

The seat assignment returned by D'Hondt$_\omega$ for a weight vector $\boldsymbol{w} = (2, 2)$ is $\boldsymbol{s} = (1, 1)$, giving both seats to party 1 who obtain 4 in representation. Suppose a weight-3 seat is added to $\boldsymbol{w}$ so as to obtain $\boldsymbol{w}^* = (3, 2, 2)$. Then D'Hondt$_\omega$ returns $\boldsymbol{s}^* = (1, 2, 3)$ that assigns a seat to each party. So party 1 go from $r_\omega(1, \boldsymbol{s}) = 4$ to $r_\omega(1, \boldsymbol{s}^*) = 3$. □

We leave open the question whether there is a WSAM that satisfies full-HM along with some of our proportionality axioms. Instead, we try to achieve positive results by restricting the weight associated with an election's additional seat. This consideration leads us to the following, weaker axiom.

Definition 4.17 (Minimal House Monotonicity, min-HM)**.** *We say a WSAM M_ω satisfies min-HM if for every weighted apportionment instance $(\boldsymbol{v}, \boldsymbol{w})$ and every $w^* \in \mathbb{N}_{\geqslant 1}$ such that $w^* \leqslant w_k$ and $\boldsymbol{w}^* = (w')_{w' \in W}$ is a non-increasing weight vector where $W = \{w \in \boldsymbol{w}\} \cup \{w^*\}$, it holds for $\boldsymbol{s} = M_\omega(\boldsymbol{v}, \boldsymbol{w})$ and $\boldsymbol{s}^* = M_\omega(\boldsymbol{v}, \boldsymbol{w}^*)$ that $r_\omega(p, \boldsymbol{s}^*) \geqslant r_\omega(p, \boldsymbol{s})$ for every party $p \in [m]$.*

This axiom provides another difference between our work and that of Chakraborty et al. (2021) as we note that min-HM can only be defined with objective values for items (seats in our case) and thus would not make much sense in the setting of Chakraborty et al. (2021).

Now, back to our analysis of the WSAMs. Positively, divisor methods clearly satisfy min-HM as all seats prior to an extra seat are assigned in the same way.

Proposition 4.23. *All divisor methods satisfy min-HM.*

While much weaker than full-HM, min-HM is enough to further distinguish between our WSAMs as Greedy$_\omega$, given LRM's failure of min-HM with unit weights, also fails min-HM (so fails full-HM).

Proposition 4.24. *Greedy$_\omega$ fails min-HM.*

Proof. Consider the following weighted apportionment instance with three parties and three seats to assign:

	# of votes ($n = 10$)
Party 1	5
Party 2	4
Party 3	1

The seat assignment returned by Greedy$_\omega$ for the weight vector $\boldsymbol{w} = (4, 3, 2)$ (so $\omega = 9$) is $\boldsymbol{s} = (1, 2, 3)$.

Now we add a weight-1 seat and for the weight vector $\boldsymbol{w}^* = (4, 3, 2, 1)$ (so $\omega = 10$), the Greedy$_\omega$ method assigns the first two seats to parties 1 and 2, respectively. In round 3, note that $q_\omega(p) - g_p^\omega(3) = 1$ for parties $p \in [3]$. Suppose that party 1 is assigned the third seat via tiebreaking. For the next round, parties 2 and 3 remain equally entitled to the last seat. Suppose that tiebreaking leads to this seat being assigned to party 2. The method then returns the seat assignment $\boldsymbol{s}^* = (1, 2, 1, 2)$ with party 3 receiving less representation than in the original weighted apportionment instance. □

We leave to future work the task of looking at other, natural notions of monotonicity. Next, we instead focus on a real-world application of weighted apportionment. For an overview of the axioms that we have studied thus far and how they relate with each other, see Figure 4.1

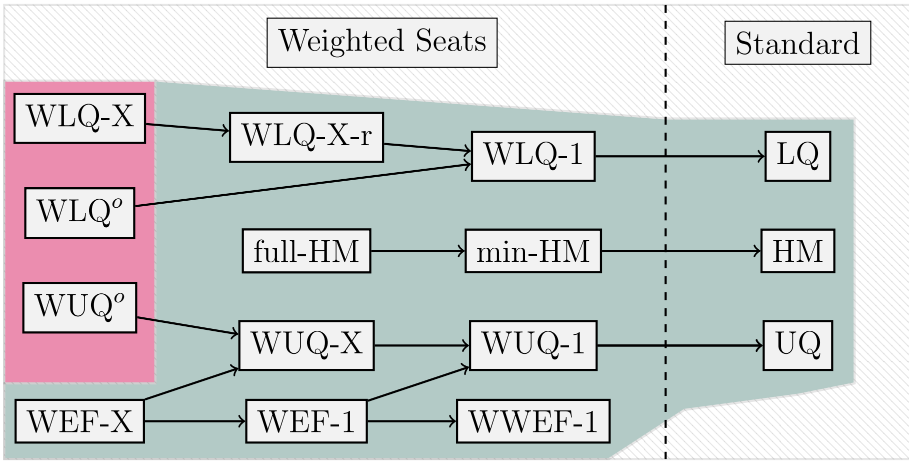

Figure 4.1: The relations between the weighted apportionment axioms studied within this chapter and the apportionment axioms mentioned in Chapter 2. In this figure, an arrow (⟶) from axiom A pointing towards axiom B indicates that axiom A implies axiom B. The backgrounds represent whether an axiom can, or cannot, always be satisfied. A green background (■) represents the former while the red background (■) represent the latter.

4.5 Experiments

In this section, we present the experimental analysis that we have conducted on our WSAMs. The section consists of two parts. The first details the case we performed on the data of the German Bundestag committee allocations (mentioned in this chapter's introduction). The second part of this section then deals with the analysis on apportionment instances that we artificially generate.

Bundestag Case Study. We present here a case study based on the allocation of chair positions to parties in Bundestag committees, first mentioned in this chapter's introduction. Our objective is to compare the results produced by our weighted apportionment methods with the historical results, full details of which are publicly available (Feldkamp, 2023; Wikipedia, 2024). In doing so, we interpret the size of a committee as a proxy for its importance. We acknowledge that this can only be an approximation of the true value of a committee, but we believe it is accurate enough to give a first impression of the performance of different WSAMs. Moreover, we wish to emphasise that our case-study does

	Adams_ω	D'Hondt_ω	Greedy_ω
WLQ-X-r	✗	✓	✓
WLQ-1	✗	✓	✓
WUQ-X	✓	✗	✓
WUQ-1	✓	✗	✓
WEF-X	✓	✗	✗
WEF-1	✓	✗	✗
min-HM	✓	✓	✗

Table 4.1: This table indicates whether a WSAM satisfies (✓) an axiom or not (✗). We only consider axioms that are satisfied by at least one of our WSAMs

not concern the distribution of seats *inside* the committees (this is settled using standard apportionment methods for each committee independently). Instead, our rules are used to decide which party can choose the head of each committee (of which there is only one per committee). While noting the presence of exceptional cases, it does not seem implausible to us that a party would consider leading two smaller committees nearly as good as leading one bigger committee.

The existing data covers all 20 legislative periods in Germany between 1949 and 2021. For each of these periods, between 4 and 7 parties entered parliament, between 19 and 28 committees were formed, and each committee had between 9 and 49 members.[3] To construct a weighted apportionment instance for a given legislative period, we take the members of parliament to be the voters,[4] we take the chair positions for the committees of that period to be the seats to be filled, and we use the sizes of those committees as the weights of the seats.

For each of the 20 weighted apportionment instances thus created, we are interested in how the historical *Bundestag* seat assignment fares in terms of representing parties proportionally and how that assignment compares to the assignments returned by our WSAMs (Adams_ω, D'Hondt_ω, and Greedy_ω). Given a seat assignment $\boldsymbol{s}$, we first ask which of our nine proportionality axioms it satisfies. For testing WLQ^o and WUQ^o, we encoded the computations of the obtainable weighted quotas $l^o(p)$ and $u^o(p)$ into Integer Linear Programs (ILP) and employed an ILP solver to compute them efficiently. As the binary measure of axiom satisfaction can only provide limited insight, we introduce three further measures to allow for a more fine-grained analysis:

[3]In case any relevant data points (such as the size of a committee) changed over the course of a legislative period, we always used the start of that period as our point of reference.

[4]Data taken from (Feldkamp, 2023)

- *average distance to the weighted quota:*

$$\delta(\boldsymbol{s}) = \frac{1}{n} \cdot \sum_{p \in [m]} |r_\omega(p, \boldsymbol{s}) - q_\omega(p)|;$$

- *average distance below the weighted lower quota:*

$$\delta^-(\boldsymbol{s}) = \frac{1}{|P^-|} \cdot \sum_{p \in P^-} l^o(p) - r_\omega(p, \boldsymbol{s}),$$

  where $P^- = \{p \in [m] \mid r_\omega(p, \boldsymbol{s}) < l^o(p)\}$;

- *average distance above the weighted upper quota:*

$$\delta^+(\boldsymbol{s}) = \frac{1}{|P^+|} \cdot \sum_{p \in P^+} r_\omega(p, \boldsymbol{s}) - u^o(p),$$

  where $P^+ = \{p \in [m] \mid r_\omega(p, \boldsymbol{s}) > u^o(p)\}$.

This gives us a total of twelve measures we can use to determine how well a given seat assignment does in terms of providing proportional representation.

Our results are summarised in Table 4.2 (the script used to generate these Bundestag case study results (along with synthetic experiment results) is available online at github.com/julianchingy/weighted-seat-experiments.git). Note that the historical seat assignments perform reasonably well in terms of our measures of quality, but both D'Hondt$_\omega$ and Greedy$_\omega$ do markedly better. This is borne out by the rate at which the axioms are satisfied, and the results of the distance measures. Notably, not only do the Bundestag seat assignments yield worse median and maximum distances than all the WSAMs, but D'Hondt$_\omega$ and Greedy$_\omega$ significantly outperform the Bundestag assignments in all three metrics (with significance level $\alpha = 0.05$). Of the latter group, Greedy$_\omega$ produces lower distance results across the board, even compared to D'Hondt$_\omega$. However, these differences between Greedy$_\omega$ and D'Hondt$_\omega$ are not statistically significant (for significance level $\alpha = 0.05$). Interestingly, Adams$_\omega$ returned the overall poorest results, particularly with regards to the distance measures, with it only standing out (positively) in testing the envy-freeness axioms. Despite the interesting results, we note that the small sample size is a drawback of this Bundestag analysis and towards addressing this concern, we continue our experimental investigation by looking at a larger pool of apportionment instances that we generate ourselves.

Synthetic Data. As we only studied 20 instances in the Bundestag case, we turn to generated data in an effort to make a more confident assessment of our WSAMs' performance. For three sets of parameter combinations, we generated 1000 weighted apportionment instances where all values were generated uniformly

	Bundestag	Adams_ω	D'Hondt_ω	Greedy_ω
WLQ^o (%)	0	0	25	30
WLQ-X (%)	20	5	95	100
WLQ-X-r (%)	20	35	100	100
WLQ-1 (%)	45	80	100	100
WUQ^o (%)	0	0	5	5
WUQ-X (%)	20	100	90	100
WUQ-1 (%)	50	100	100	100
WEF-X (%)	0	100	10	25
WEF-1 (%)	10	100	30	50
Median δ	17.1	16.6	4.7	2.5
Maximum δ	66	30.8	6.6	5.9
Median δ^-	25.2	17	5	2.2
Maximum δ^-	101	44	9	6
Median δ^+	32.2	15.2	6.2	3.6
Maximum δ^+	163	31	11	10

Table 4.2: Summary of results for the 20 Bundestag committee election instances. For each of the four seat assignments, the table shows: (*i*) for each axiom, the percentage of election instances for which the axiom is satisfied, and (*ii*) for each distance measure, the median and maximum distances across the 20 election instances.

at random from the chosen ranges:[5] 10 parties (votes ranging from 5 to 300) and 25 committee seats (weights ranging from 10 to 50);

1. 10 parties (votes ranging from 5 to 300) and 25 committee seats (weights ranging from 10 to 50)

2. 10 parties (votes ranging from 5 to 1000) and 100 seats (weights ranging from 1 to 1000)

3. 10 parties (votes ranging from 5 to 1000) and 100 seats (weights ranging from 1 to 101)

The first set of parameters (*i*) sees us trying to mimic the structure of the Bundestag instances while looking at more than just 20 instances. The latter two sets use the same number of parties with an increase in seats with the intended benefit being greater flexibility for the WSAMs to satisfy the axioms. The difference in their range for the seat weights—one ranging between $1 - 1000$ and the other between $1 - 101$—is made to see if any discrepancies in performances occur

[5]Note that when we generate the weights of seats, we use integer values.

between the two 'extremes': one where the seat weights may vary considerably and another where the seat weights have smaller differences between them).

We then proceeded as in the Bundestag case study and applied the twelve measures used in the Bundestag case study to compare the seat assignments produced by our three WSAMs on the three sets of generated instances. The results for the data that most resembles that of the Bundestag committee data can be seen in Table 4.3.

	Adams_ω	D'Hondt_ω	Greedy_ω
WLQ^o (%)	0	2.3	2.3
WLQ-X (%)	35.9	99.2	97.8
WLQ-X-r (%)	72.5	100	100
WLQ-1 (%)	100	100	100
WUQ^o (%)	0	0	0.4
WUQ-X (%)	100	96.4	100
WUQ-1 (%)	100	100	100
WEF-X (%)	100	6	23.4
WEF-1 (%)	100	56	72.8
Median δ	7.9	3.3	3.1
Maximum δ	23.9	7.2	6.9
Median δ^-	6.8	3.5	3.3
Maximum δ^-	23	10	10.7
Median δ^+	9.8	3.8	3
Maximum δ^+	28	12	8.2

Table 4.3: Summary of results for the 1000 generated election instances with 10 parties (votes ranging from 5 to 300) and 25 committee seats (weights ranging from 10 to 50). For each of the three seat assignments, the table shows: (*i*) for each axiom, the percentage of election instances for which the axiom is satisfied, and (*ii*) for each distance measure, the median and maximum distances across the 1000 election instances.

With Table 4.3's results, we find that the performances of our WSAMs on these generated instances mostly mirror their performances on the Bundestag committee instances. Specifically, Adams_ω only outperformed the other WSAMs, at least to a significant degree, in the envy-freeness measures while D'Hondt_ω and Greedy_ω once again bore the most positive results with regard to the remaining measures However, we do find that Adams_ω fares much better on the lower-quota-like axioms as well as the median values for the three distance than it did in the Bundestag case study. Thus, when looking at more instances, a more positive picture is painted for Adams_ω. And much like the Bundestag case, our checks showed no statistical significance (for significance level $\alpha = 0.05$) between the distance measure results of D'Hondt$_\omega$ and Greedy$_\omega$.

We move onto the results on the two data generated sets with 10 parties (with votes ranging from 5 to 1000) and 100 seats can be found in Tables 4.4 (weights ranging from 1 to 1000) and 4.5 (weights ranging from 1 to 101).

	Adams$_\omega$	D'Hondt$_\omega$	Greedy$_\omega$
WLQ^o (%)	0	0	0
WLQ-X (%)	0.8	25.3	21.8
WLQ-X-r (%)	73.6	100	100
WLQ-1 (%)	100	100	100
WUQ^o (%)	0	0	0
WUQ-X (%)	100	99.1	100
WUQ-1 (%)	100	100	100
WEF-X (%)	100	11.8	12.2
WEF-1 (%)	100	99.4	96.5
Median δ	67.5	9.2	10.5
Maximum δ	457.4	35.4	47
Median δ^-	51.2	8.4	9.8
Maximum δ^-	336.3	36.7	42.5
Median δ^+	102	12	11.8
Maximum δ^+	818	93	78

Table 4.4: Summary of results for the 1000 generated election instances with 10 parties (votes ranging from 5 to 1000) and 100 seats (weights ranging from 1 to 1000). For each of the three seat assignments, the table shows: (*i*) for each axiom, the percentage of election instances for which the axiom is satisfied, and (*ii*) for each distance measure, the median and maximum distances across the 1000 election instances.

We find that no notable difference between the WSAM's performances arise when we consider the set of instances where the seat weights may vary by large amounts (results in Table 4.4) and compare them to how the WSAMs fared on the Bundestag-like instances (results in Table 4.3). The two most noteworthy changes are (*i*) the worsened performance of Adams$_\omega$ in satisfying WLQ-X and (*ii*) the improvements of the other WSAMs in providing envy-freeness.

Now, if we look at the WSAMs' behaviour on the instances where there are many seats that do not have large differences in weights (results in Table 4.5), we find interesting results. In particular, D'Hondt$_\omega$ outperforms both Adams$_\omega$ and Greedy$_\omega$ (although only a small margin above Greedy$_\omega$) when it comes to

	Adams_ω	D'Hondt_ω	Greedy_ω
WLQ^o (%)	0	28	14.5
WLQ-X (%)	0.1	30.6	13.9
WLQ-X-r (%)	76	100	100
WLQ-1 (%)	100	100	100
WUQ^o (%)	0	9.9	8.8
WUQ-X (%)	100	99.5	100
WUQ-1 (%)	100	100	100
WEF-X (%)	100	14.2	17.9
WEF-1 (%)	100	99	98.5
Median δ	6.3	0.8	0.9
Maximum δ	43.7	2	2.1
Median δ^-	4.7	1	1
Maximum δ^-	40.2	3	2.3
Median δ^+	10.7	1	1
Maximum δ^+	85	4	2.2

Table 4.5: Summary of results for the 1000 generated election instances with 10 parties (votes ranging from 5 to 1000) and 100 seats (weights ranging from 1 to 101). For each of the three seat assignments, the table shows: (*i*) for each axiom, the percentage of election instances for which the axiom is satisfied, and (*ii*) for each distance measure, the median and maximum distances across the 1000 election instances.

satisfying WUQ^o despite not satisfying any upper-quota-like axiom in general.[6] That along with the fact that D'Hondt_ω outperforms both of the other WSAMs with respect to WLQ^o, suggests that D'Hondt_ω may be a well-rounded option for providing proportionality in our setting. However, more extensive experimental analysis would be required to make any definitive claims in this regard.

4.6 Chapter Summary

We studied a model of apportionment with weighted seats, and generalised apportionment methods and central axioms from the apportionment literature to this model. Direct generalisations of the axioms, we found, yield (mostly) negative results, but mild relaxations are amenable to positive results. The positive outlook is further justified, in particular for the D'Hondt_ω and Greedy_ω methods, by an experimental case study on Bundestag committee assignments.

[6]Note that it only matched Greedy_ω in Table 4.2.

These examples illustrate that extra care needs to be taken when generalising our work on apportionment with weighted seats to more complex settings.

Future Work. We finish off by pointing out some natural directions for future research that follow from our work. Amongst the open questions regarding house monotonicity (in Section 4.4) such as looking at the alternative adaptations of house monotonicity, we acknowledge the following as the most pressing future work stemming from this section: investigating the existence of a natural WSAM that satisfies the full-HM axiom. Naturally, given the rich apportionment literature, it would be interesting to study other prominent apportionment properties and rules, e.g., population monotonicity and the Sainte-Laguë method. Also, the experimental results in Section 4.5 suggest that our WSAMs—especially D'Hondt$_\omega$ and Greedy$_\omega$—are worth further experimental investigation such as looking at other sources of data, or testing the WSAMs against different quantitative measures of proportionality.

Finally, we highlight that a natural follow-up to the results on apportionment would be to lift the positive weighted-seat results to a more general setting, e.g., multiwinner voting. One step in this direction would be to forego the assumption that we made of full supply, i.e., that each party has enough members to fill all available k seats. In concluding, we briefly touch on this assumption and specifically, we want to make the case that it is a necessary one. In fact, we find that, once making this step of dropping full supply, even the weakest of our axioms fail to be satisfied. We start with WLQ-1.

Example 4.14 (WLQ-1 without full supply)**.** Consider this weighted apportionment instance with four parties and 72 seats:

	# of votes ($n = 10$)	$q_\omega(p)$
Party 1	3	30
Party 2	3	30
Party 3	3	30
Party 4	1	10

$$\boldsymbol{w} = (15, 15, 1, \ldots, 1),\ k = 72,\ \omega = 100$$

Suppose each party $p \in [3]$ can receive two seats while party 4 can receive 66 seats. So some party $p \in [3]$ with $q_\omega(p) = 30$ must be assigned two weight-1 seats and so WLQ-1 cannot be provided.

The same can also be shown for WUQ-1 with the following example.

Example 4.15 (WUQ-1 without full supply)**.** Consider this weighted apportionment instance with three parties and five seats:

	# of votes ($n = 3$)	$u^o(p)$
Party 1	1	3
Party 2	1	3
Party 3	1	3

$$\boldsymbol{w} = (3, 3, 3, 3, 3),\ \omega = 15$$

Suppose each party $p \in [2]$ can receive one seat while party 3 can receive three seats. So, party 3 with $q_\omega(p) = 3$ must be assigned three seats of weight 3 and this violates WUQ-1.

These preliminary thoughts show that such an extension to multiwinner voting requires care, as our positive apportionment results fail without the mild assumption of full supply. In the chapter that follows, we take initial steps towards this task of extending our model to multiwinner voting.

Chapter 5

Approval-based Multiwinner Voting with Weighted Seats

Over recent years, significant progress has been made on research on multiwinner voting and the attention that this model has drawn (and continues to draw) comes as no surprise given the extensive real-world scenarios that it can capture. Of the research conducted on this model, that of proportionality has attracted plenty of focus (see Chapter 2 for some of this work). Here, we are off the back of Chapter 4 where we tackled the apportionment task with the added complication of the parliamentary seats being valued differently by voters. A similar drawback presents itself for the standard multiwinner model with that drawback being that all seats on the committee are usually considered to be of equal value. In practice, this is often not the case, as the following examples demonstrate.

1. *Participatory Budgeting*: A local municipality is tasked with running an instance of participatory budgeting where not only the selection of public projects must be made, but the municipality must also select where (or when within some set timeline) the projects are to be implemented. In this case, it would be reasonable to assume that voters would prefer that their approved projects be situated closer to them (or are realised sooner) instead of these projects to be some considerable distance from them (or having low priority regarding when they are implemented).

2. *Television Slot Assignment*: The executives of a television network wish to assign movies/shows to certain timeslots. Beyond some executives having preferences amongst the movies/shows, it may be that some timeslots are more preferable as they ordinarily attract more viewers, for example.

3. *Committee Election*: Seats on a project committee within Company A's tech department must be filled, with each seat representing the head of some

project. It is natural to assume that some projects hold more importance amongst the voters than others. Hence, the voters wish for their most preferred candidates to head these projects, e.g., being the head of Project A may be considered more valuable than being head of Project B.

In this chapter, we take the first steps in generalising the work of Chapter 4 and pose the following research question.

What proportionality guarantees can we provide within the multiwinner voting model when the committee seats are of differing value?

So in this chapter, we enrich the multiwinner voting model with weights over the seats. But how do we determine such weights? In the example of Company A, how does the decision-maker quantify the value that a voting employee has for a particular project? Ideally, the decision-maker can elicit the voters' own subjective utilities associated with each committee seat. However, this poses its own set of challenges (e.g., it is non-trivial for voters to place an exact value on each seat). Towards getting a handle on this model, we make a simplifying assumption, much like in Chapter 4, that each seat can be assigned some number that represents its true value. Moreover, we assume that this value is the same for all of the voters (so voters are in agreement on every seat's value). How does this assumption affect the potential practical use of our results? Let us go back to our examples from above to show that these objective values occur naturally.

1. *Participatory Budgeting*: Suppose there is a participating budgeting process with designated periods for the municipality to implement the chosen projects. Thus, any selected projects must also be assigned to a specific period where they are to be realised. Now suppose that these available implementation periods start at varying times after the project selection process concludes: some occur a few weeks afterwards while other periods are scheduled to happen after over a year. In such a case, the start date of each implementation period is the same for every voter. Thus, the start dates can used as objective measures of every voters value of the implementation periods.

2. *Television Slot Assignment*: Suppose that the television network holds viewership data for each timeslot and this serves as a reliable estimation of the number of eyes a movie/show will obtain if assigned to a particular timeslot. Such data can serve as a proxy for how every voting executive values the timeslots.

3. *Committee Election*: Suppose that each of Company A's project leaders shall be afforded a certain amount of company funds to see the project through to completion. It is not unreasonable to assume that the supporters

of a particular candidate would prefer said candidate to take the lead on more expensive projects (as more is at stake for the company as a whole). In this way, each project's associated funds can represent its some objective value for the voters.

Of course, there are many scenarios where determining the proxy of each seat's value is not as straightforward. However, we argue that our restricted scope is sensible for this initial, exploratory step towards achieving proportional representation in this domain.

Additional Related Work. As already mentioned in the introduction of Chapter 4, the model of participatory budgeting bears some resemblance to ours (Rey and Maly, 2023). Results from that chapter make clear that this seeming resemblance does not lead to similar results with there being a clear difference between the models. Here, we use the EJR axiom to drive this point home further. EJR, which corresponds to lower quota in the apportionment setting, is satisfiable in the participatory budgeting setting (Peters et al., 2021b), while we saw in Chapter 4 that our lower quota adaptation is not generally satisfiable for weighted-seat apportionment. Thus, one can expect the results from our weighted-seat multiwinner model to also differ from that of participatory budgeting. In later sections, this expectation shall be confirmed.

We also note the resemblance to the following settings (that have bee briefly discussed in Chapter 3). Take the budget division (or probabilistic social choice) model (Aziz et al., 2020; Brandt, 2017). Here, voters may not only care about having their approved projects receiving some funding, but also care about the amount of funding assigned to their approved projects (which may reflect the weight of a seat). So one can see the similarity if we consider our setting as a budget division task where exactly k projects are to be assigned some funding with the division of the budget into k parts being fixed from the outset (Brandl et al., 2022).

Our model also resembles that of approval-based facility location (Deligkas et al., 2023; Elkind et al., 2022), with this chapter's running participatory budgeting example being illustrative of the similarities between the models. Recall that the goal of approval-based facility location is to place a k-sized subset of facilities on some locations. Here, the voters' approvals of certain facilities are dependent on whether facilities are placed in locations within a certain distance from them. So not only does the selection of the candidate (facility) have importance, but so does the seat (location) that they are assigned to. Thus, we find a similarity with our idea of committee seats having different values for each voter (at least regarding what they approve of).

Chapter Outline. In Section 5.1, we present the multiwinner voting model with weighted seats as well as some rules for this setting. Section 5.2 details how we adapt the cohesiveness notion and EJR to our model. In this section, we also explore various relaxations of our weighted EJR axiom and, specifically, we ascertain which rules satisfy these axioms. Section 5.3 sees us import the priceability notion to our weighted-seat setting and test the rules against this notion. Section 5.4 offers our summary of the chapter as well as potential paths for future work.

5.1 The Model

Take a set of m candidates $C = \{a, b, c, \ldots\}$ and a set of n voters $N = \{1, \ldots, n\}$. Each candidate is in contention for one of k seats with $k \leqslant m$. Each voter $i \in N$ submits an approval ballot $A_i \subseteq C$, indicating the candidates they approve of. For a candidate c, we denote the set of voters that approve of a as $N(c) = \{i \in N \mid c \in A_i\}$. An approval profile $\boldsymbol{A} = (A_1, \ldots, A_n)$ is a vector of the n voters' approval ballots.

Each seat is associated with a weight and we again use a weight vector $\boldsymbol{w} = (w_1, \ldots, w_k) \in \mathbb{R}^k_{\geqslant 1}$ to denote the non-increasing vector of the seats' weights. We denote the total weight of the seats as $\omega = \sum_{w_j \in \boldsymbol{w}} w_j$.

Formally, an outcome is a seat assignment $\boldsymbol{s} = (s_1, \ldots, s_k) \in C^k$ of k candidates such that $s_i \neq s_j$ for all $i \neq j \in [k]$ (each candidate is assigned *at most* one seat). Given a seat assignment $\boldsymbol{s} = (s_1, \ldots, s_k)$, we say a candidate $s_t \in \boldsymbol{s}$ is elected to the seat t with weight $w_t \in \boldsymbol{w}$.

A *weighted multiwinner election instance* is denoted by the pair $(\boldsymbol{A}, \boldsymbol{w})$. We refer to a weighted election instance $(\boldsymbol{A}, \boldsymbol{w})$ as a *unit-weight election instance* when $w_t = w_{t'}$ for all $t, t' \in [k]$ (which represents standard approval-based multiwinner election instances).

Consider some set of candidates $A \subseteq C$. The satisfaction with a seat assignment $\boldsymbol{s}$ is given this set A is determined by the weight of the seats occupied by candidates from A. Formally, the *weighted satisfaction* from seat assignment $\boldsymbol{s} = (s_1, \ldots, s_k)$ with respect to candidate set A is denoted as:

$$sat_\omega(A, \boldsymbol{s}) = \sum_{j=1}^{k} \mathbb{1}_{s_j \in A} \cdot w_j.$$

For a weight vector $\boldsymbol{w}$, the set of all possible weighted satisfaction values:

$$\mathcal{S}(\boldsymbol{w}) = \left\{ \sum_{t \in T} w_t \;\middle|\; T \in \mathcal{P}([k]) \right\}.$$

Note that $sat_\omega(A, \boldsymbol{s}) \in \mathcal{S}(\boldsymbol{w})$ for any set of candidates $A \subseteq C$ and any seat assignment $\boldsymbol{s}$.

Example 5.1. Here is an example with two voters and five candidates with there being four seats that must be assigned:

	a	b	c	d	e
Voter 1	✓	✓	✓	✗	✓
Voter 2	✓	✗	✗	✗	✗

$$\boldsymbol{w} = (10, 5, 4, 1), \omega = 20$$

For this example, observe that the set of possible satisfaction values is $\mathcal{S}(\boldsymbol{w}) = \{0, 1, 4, 5, 6, 9, 10, 11, 14, 15, 16, 19, 20\}$. As voter 2 only approves of one candidate then this voter can only obtain satisfaction in $\{0, 1, 4, 5, 10\}$ while voter 1 could get a satisfaction of all values in $\mathcal{S}(\boldsymbol{w}) \setminus \{20\}$.

A *weighted-seat approval-based multiwinner voting (WMWV) rule* F_ω will take a weighted election instance $(\boldsymbol{A}, \boldsymbol{w})$ as input, and map it to a non-empty set of winning seat assignments denoted as $F_\omega(\boldsymbol{A}, \boldsymbol{w})$ (so rules may be *irresolute*). We now introduce some of the WMWV rules that we will study in this chapter.

To start with, inspired by the work on picking sequences (Bouveret et al., 2016; Chakraborty et al., 2021), we look to employ standard approval-based multiwinner voting rules to create sequential assignment rules in our setting. Specifically, we use a sequential approval-based multiwinner voting rule F_α such as MES to set an ordering of candidates to be assigned. In this way, the ordering does not depend on the weight vector $\boldsymbol{w}$ of the election instance. This still leaves the specific weights to be assigned to the k winning candidates. As candidates are selected sequentially, we shall assume that the seats are assigned in non-increasing order of weight (since we assume that more weight is more valuable). We refer to such rules as *seat-based* WMWV rules. We focus on MES[arbitrary] and MES[seq-Phragmén] as the seat-based rules due to their strong proportionality guarantees in the standard multiwinner setting.

We now define WMWV rules take the weight vector into account to a greater extent than the seat-based rules. We adapt two approval-based multiwinner rules, namely seq-Phragmén and MES and refer to these adaptations as the *weight-based* WMWV rules.[1]

Definition 5.1 (seq-Phragmén$_\omega$)**.** *During the run of this method, voters continuously earn funds and may pay to assign a seat to a candidate that they collectively approve. Assigning a candidate to seat t comes with a price of w_t. Each voter*

[1]Note that for these rules, the seats are not necessarily assigned in non-increasing order of weight.

begins with a personal budget of 0 *and this budget continuously grows. If at some moment, the voters that approve of an unelected candidate* $c \in C$ *collectively hold enough funds to pay the price of* w_t *of the unassigned seat currently under consideration, then this candidate* c *is assigned to seat* t*, the funds of all voters approving of candidate* c *are set to* 0*, and then the process continues until all* k *seats have been assigned to create a seat assignment* $\boldsymbol{s}$*. Any ties are broken first in favour of the groups that were able to afford the price at an earlier moment, and otherwise, they are broken arbitrarily.*

When operating on unit-weight multiwinner elections, the seq-Phragmén$_\omega$ rule is exactly seq-Phragmén. We continue with our weighted variant of MES.

Definition 5.2 (MES$_\omega$)**.** *The rule will occur in some rounds* $r \in \{1, \dots, k\}$ *where voters may pay to assign candidates to seats with associated weights in the weight vector* $\boldsymbol{w} = (w_1, \dots, w_k)$*. In each round* r*, the assignment to the* r*-th committee seat of weight* w_r *incurs a price* w_r*. With* $b_i(r)$ *being voter* i*'s budget to start round* r*. We set* $b_i(1) = {}^{\omega}/_{n}$ *as the initial budget of every voter* $i \in N$*. The rule starts at round* 1 *with an empty seat assignment* $\boldsymbol{s}^0 = ()$*, and will assign seats to candidates over rounds. In round* r*, we say a pair* (c, w_r) *is* q-affordable *for some* $q \in \mathbb{R}_{\geqslant 0}$*, with* $c \in C \setminus \boldsymbol{s}^{r-1}$ *and with seat* r *being unassigned, if:*

$$\sum_{i \in N(c)} \min(q, b_i(r)) \geqslant w_r.$$

If there exists no pair that is q*-affordable in round* r*, the rule stops and returns the seat assignment* $\boldsymbol{s}^r$*.*

Otherwise, for a q*-affordable pair* (c, w_r) *for a minimum* q *(use some arbitrary tie-breaking if there are* q*-affordable pairs), the rule will fix* $\boldsymbol{s}^r$ *by setting the* r*-th element in* $\boldsymbol{s}^{r-1}$ *to be candidate* c*. The budget of each voter* $i \in N(c)$ *is then set to* $b_i(r+1) = b_i(r) - \min(q, b_i(r))$ *while for voters* $i \notin N(c)$*, we set* $b_i(r+1) = b_i(r)$*.*

As with the standard MES, we must deal with how we complete a potentially partial seat assignment provided by MES$_\omega$. Using the same notation as in the approval-based setting without weights, we write MES$_\omega$[arbitrary] for completion using arbitrary assignment and MES$_\omega$[seq-Phragmén$_\omega$] for completion using the seq-Phragmén$_\omega$ rule. The latter uses seq-Phragmén$_\omega$ to complete a partial MES$_\omega$ seat assignment in an analogous fashion to the way in which seq-Phragmén completes the partial committees of MES (as detailed in Chapter 2).

5.2 Weighted Justified Representation

The work on the obtainable weight lower quota WLQ^o in the previous chapter leads to the following definition of cohesiveness so as to take into account both the number of seats and the seats' weights.

Definition 5.3 (ℓ_ω-cohesiveness)**.** *For an integer $\ell_\omega \in \mathcal{S}(\boldsymbol{w})$, a set of voters $N' \subseteq N$ is ℓ_ω-cohesive if both of the following conditions hold:*

- $|N'| \geqslant \ell_\omega \cdot n/\omega$.
- *There exists a set of candidates $C' \subseteq \bigcap_{i \in N'} A_i$ with size $|C'| = t$ such that there exists a seat assignment $\boldsymbol{s}$ where $sat_\omega(C', \boldsymbol{s}) \geqslant \ell_\omega$ and $|N'| \geqslant n \cdot t/k$.*

Intuitively, a group of voters is ℓ_ω-cohesive if they can find a seat assignment that holds t candidates that they commonly agree on, and yields ℓ_ω satisfaction for these t candidates, such that they are large enough to demand t seats in the seat assignment, and large enough to demand ℓ_ω of the total weight ω. This leads to a modified definition of the multiwinner EJR for this weighted setting.

Definition 5.4 (Weighted EJR, WEJR)**.** *A seat assignment $\boldsymbol{s}$ provides WEJR if for every ℓ_ω-cohesive group $N' \subseteq N$ such that there exists a candidate $c \in \bigcap_{i \in N'} A_i$ with $c \notin \boldsymbol{s}$, there exists a voter $i \in N'$ such that $sat_\omega(A_i, \boldsymbol{s}) \geqslant \ell_\omega$.*

The requirement for an ℓ_ω-cohesive group to have some collectively approved candidate, that does not have a seat, can be thought of as a multiwinner analogue of the full supply assumption made in Chapter 4 (full supply was discussed in detail in the summary of Chapter 4). With this observation, it is clear that WEJR implies WLQ^o for election instances with party-list profiles. Thus, the negative weighted-apportionment result that a seat assignment providing WLQ^o may not exist (see Theorem 4.2) implies that we may not be able to satisfy WEJR.

Corollary 5.1. *A seat assignment $\boldsymbol{s}$ that provides WEJR may not exist.*

Moreover, as WEJR implies WLQ^o, Proposition 4.3's hardness result for WLQ^o carries over for WEJR as we can use the same reduction from the Partition problem for WEJR.

Corollary 5.2. *If there exists a polynomial-time algorithm $\mathbb{A}$ that finds a seat assignment $\boldsymbol{s}$ that provides WEJR whenever such a seat assignment exists, then* $\mathsf{P} = \mathsf{NP}$.

Furthermore, as WEJR implies EJR from standard multiwinner voting and, as mentioned in Chapter 2, as checking whether a multiwinner committee provides EJR is a coNP-complete problem, it follows that checking whether a weighted seat assignment provides WEJR is also coNP-hard.

The result of Proposition 4.7 gives us no hope for an 'up-to-any' relaxation that is always satisfiable in the multiwinner setting so we skip over to the 'up-to-one' version instead.

5.2.1 Up-to-one Weakening of WEJR

Given that WLQ-1 is always satisfiable in the weighted-seat apportionment setting, we turn towards an 'up-to-one' weakening of WEJR. Before doing so, here is some extra notation that we require. For a seat assignment $\boldsymbol{s} = (s_1, \ldots, s_k)$ and a set of candidates A, let $\boldsymbol{s}(A) = (t)_{t\in[k], s_t\in A}$ denote the increasing vector of the positions within $\boldsymbol{s}$ of candidates from A (e.g., for a set $A = \{a, b, c\}$ and a seat assignment $\boldsymbol{s} = (b, d, a, e)$, we have $\boldsymbol{s}(A) = (1, 3)$).

Definition 5.5 (WEJR up to one seat, WEJR-1)**.** *A seat assignment $\boldsymbol{s}$ provides WEJR-1 if for every ℓ_ω-cohesive group $N' \subseteq N$ such that there exists a candidate $c \in \bigcap_{i\in N'} A_i$ with $c \notin \boldsymbol{s}$, there exists a voter $i \in N'$ and some seat $t \in \boldsymbol{s}(C \setminus A_i)$ such that $sat_\omega(A_i, \boldsymbol{s}) + w_t > \ell_\omega$.*

In the unweighted multiwinner model, WEJR-1 reduces to EJR. Now, we move onto the usual question of whether this axiom can always be provided. To answer this in the positive, we introduce the following rule that is an adaptation of the Greedy Cohesive Rule that was devised for multiwinner elections (see Definition 2.13), and is known to satisfy EJR.

Definition 5.6 (Weighted Greedy Greedy Cohesive Rule, GCR$_\omega$)**.** *The rule starts with an empty seat assignment, all voters being active, set a variable $x = 1$ and recall that the weight vector $\boldsymbol{w} = (w_1, \ldots, w_k)$ is in non-increasing order. In iterations, the rule looks for the largest ℓ_ω-cohesive group of active voters N', assigns h candidates approved by voters in N' to the h unassigned seats in $\{x, \ldots, x+h\}$ with the largest weights such that:*

$$\sum_{t\in(w_i)_{w_i\in\boldsymbol{w}, x\geqslant i\geqslant x+h}} w_t \;\leqslant\; \ell_\omega \;<\; w_{x+h+1} + \sum_{t\in(w_i)_{w_i\in\boldsymbol{w}, x\geqslant i\geqslant x+h}} w_t.$$

The rule then sets the variable to $x = t + 1$, makes the voters in N' inactive and continues to the next iteration, unless there are either (i) no more active voters, or (ii) no more seats to be assigned. In case either (i) or (ii) occurs, the rule terminates.

Much like the first phase of MES$_\omega$, this rule may return a seat assignment that is not complete. If necessary, we complete partial GCR$_\omega$ seat assignments in an arbitrary manner and refer to it as GCR$_\omega$[arbitrary]. We do not consider completion using a specific WMWV rule such as seq-Phragmén$_\omega$, as doing similarly in the standard multiwinner setting does not yield results that are notably more positive. For example, completing GCR with seq-Phragmén does not lead to a priceable rule but EJR remains satisfied (Peters and Skowron, 2020). However, the method of completion has no bearing on the following result regarding WEJR-1.

Theorem 5.3. *GCR_ω satisfies WEJR-1.*

Proof. Assume to the contrary that GCR_ω outputs a seat assignment $\boldsymbol{s}$ such that there exists an ℓ_ω-cohesive group of voters N' where there exists a candidate $c \in \bigcap_{i\in N'} A_i$ with $c \notin \boldsymbol{s}$ and where no seat $t \in \boldsymbol{s}(C \setminus A_i)$ of weight w_t exists such that $sat_\omega(A_i, \boldsymbol{s}) + w_t > \ell_\omega$ for some voter $i \in N'$. We can focus on a potential swap using the unelected candidate $c \in \bigcap_{i\in N'} A_i$ where $c \notin \boldsymbol{s}$. Let us assess the situation before completing the committee.

If this entire group of voters N' remained in the last round of GCR_ω, then GCR_ω would've made the required assignments of best seats for this group N'. To see that there is enough total weight available to assign for group N', observe that in each round of GCR_ω where some group of voters N^* is made inactive, at most $\omega \cdot {|N^*|}/{n} \geqslant \ell_\omega$ in weight is assigned in this round. So if voter group N' is remains in its entirety then at most $\omega \cdot {(n-|N'|)}/{n}$ in weight has been assigned in prior rounds. And N' is ℓ_ω-cohesive, we get that $|N'| \geqslant \ell_\omega \cdot {n}/{\omega}$ and therefore, we have that strictly less than the following in weight has been assigned in total:

$$\omega \cdot \frac{n - |N'|}{n} \leqslant \omega \cdot \frac{n - (\ell_\omega \cdot {n}/{\omega})}{n} = \omega - \ell_\omega.$$

Thus, there was enough weight available for GCR_ω to assign to N' such that this voter group is not a witness to the violation of WEJR-1.

Now, it must be the case N' must not have remained in its entirety during GCR_ω's last round, hence there must be some voter $i \in N'$ that was made inactive in an earlier round. Thus, voter i was part of an ℓ^*_ω-cohesive group N^* where $\ell^*_\omega \geqslant \ell_\omega$. But then there exists some seat t with weight w_t, specifically the unassigned seat with largest weight after voter group N^* was made inactive, such that assigning this weight-w_t seat to an unelected candidate approved by voter i, gives voter i strictly more than $\ell^*_\omega \geqslant \ell_\omega$ in satisfaction. So in the end, we get that N' cannot exist and that GCR_ω satisfies WEJR-1, even with a potentially partial seat assignment. □

Indeed, this result is a positive one. However, it does come with the downside that the seat assignments returned by GCR_ω are hard to compute (as GCR_ω must iterate through all ℓ_ω-cohesive groups). So with WEJR-1 always being satisfiable, it is of immediate interest whether it can always be satisfied by a rule that is polynomial-time computable. For one of our weight-based rules that can be computed in polynomial time, namely seq-Phragmén$_\omega$, we know that it does not satisfy WEJR-1 (as in standard multiwinner voting, seq-Phragmén does not satisfy EJR (see Chapter 2)). However, the following can be shown for our weight-based MES_ω.

Theorem 5.4. *MES_ω satisfies WEJR-1.*

Proof. Take a (possibly partial) seat assignment $\boldsymbol{s}$ returned by MES_ω and assume that there exists an ℓ_ω-cohesive group of voters N' such that there is a candidate $c \in \bigcap_{i \in N'} A_i$ with $c \notin \boldsymbol{s}$ and there is no seat $t \in \boldsymbol{s}(C \setminus A_i)$ so that for some voter $i \in N'$, we have $sat_\omega(A_i, \boldsymbol{s}) + w_t > \ell_\omega$. So each voter in N' ash at most $\ell_\omega - w_h$ in satisfaction where w_h is the largest seat not assigned to any candidate approved by a voter from N'.

We consider two cases. First, assume that for some round r of MES_ω, some voter $i \in N'$ contributed more than $1/|N'|$ per unit of weight for assigning the weight w_r. So, before round r, each voter in N' contributed at most $1/|N'|$ per unit of weight for a seat assignment, and since each voter is satisfied by at most $\ell_\omega - w_h$, they each had spent at most $1/|N'| \cdot (\ell_\omega - w_h)$ before round r. Thus, in round r, they each had at least the following in personal budget:

$$\frac{\omega}{n} - \frac{\ell_\omega - w_h}{|N'|} = \frac{|N'| \cdot \omega/n - \ell_\omega + w_h}{|N'|} \geqslant \frac{\ell_\omega - \ell_\omega + w_h}{|N'|} = \frac{w_h}{|N'|}.$$

Thus, as a collective, group N' has at least w_h in funds in round r so they could afford seat r with a price of w_r. Now, in this round r, we assume that voter i contributed to purchasing the assignment of the seat with weight w_r. Suppose that this purchase was w_r/α-affordable and we assumed that voter i contributed $1/\alpha$ per unit of weight with $1/\alpha > 1/|N'|$. But this means that $w_r/\alpha > w_r/|N'|$, so MES_ω would have rather let the voter group N' make the assignment to seat r with each voter $i \in N'$ contributing at most $1/|N'|$ per unit of weight, contradicting that voter i contributed more than this value.

Now, consider the other case, where each voter in N' spent at most $1/|N'|$ per unit of weight during the entire MES_ω process. Then from the reasoning above, the voter group N' has at least w_h in funds MES_ω's end. As there is money leftover, some seats were not assigned, and as weight w_h is the largest seat that N' was not assigned, they could afford to pay for some candidate that they approve of to be assigned to it (we know such a candidate exists by the assumption on N' witnessing a violation of WEJR-1). This contradicts that the rule terminated and thus, such a voter group N' cannot exist. □

This means that we can find seat assignments providing WEJR-1 in polynomial time via the use of MES_ω's first phase.

The following example then illustrates that we cannot get the same positive result using standard MES, regardless of the completion method used (amongst those that we consider).

Example 5.2 (MES, MES[arbitrary] and MES[seq-Phragmén] fail WEJR-1)**.** Consider this example with two voters, eight candidates and seven seats to be assigned:

	$\{c_1, \ldots, c_4\}$	$\{c_5, \ldots, c_8\}$
Voter 1	✓	✗
Voter 2	✗	✓

$$\boldsymbol{w} = (46, 30, 20, 1, 1, 1, 1), \omega = 100$$

Each voter begins MES with budget of $7/2$ and can afford to pay for three seat assignments. With seats assigned in non-increasing order, assume that the tie-breaking favours voter 1 and the seat assignment after the first phase of MES is $(c_1, c_2, c_3, c_5, c_6, c_7, \_)$. This does not provide WEJR-1 with the violation occurring with voter 2.

Now, for both completion methods, suppose that voter 1 is favoured once more and the final seat assignment for both MES[arbitrary] and MES[seq-Phragmén] is $\boldsymbol{s} = (c_1, c_2, c_3, c_5, c_6, c_7, c_4)$. Observe that voter 2 is a 50-cohesive group with a satisfaction of 3 and, thus, no seat that was assigned to voter 1 suffices to voter 2 getting more than their deserved satisfaction. So WEJR-1 remains violated.

From the previous chapter, we found that we could go a step further than the up-to-one notion in finding a weighted lower quota axiom that is always satisfiable. This was observed with the WLQ-X-r axiom (see Definition 4.8). We now generalise WLQ-X-r to this weighted-seat multiwinner model. To do so, we introduce the following notation. Given a seat assignment $\boldsymbol{s}$ for a weighted election instance $(\boldsymbol{A}, \boldsymbol{w})$, we denote the set of *overrepresented* ℓ_ω-cohesive groups as:

$$N_{>\ell_\omega} = \{N' \subseteq N \mid N' \text{ is } \ell_\omega\text{-cohesive and } sat_\omega(A_i, \boldsymbol{s}) > \ell_\omega \text{ for some } i \in N'\}.$$

We then denote the set of candidates that are approved by *some* overrepresented ℓ_ω-cohesive group with the following:

$$C_{\ell_\omega} = \{c \in C \mid c \in \bigcap_{i \in N'} A_i \text{ for some } N' \in N_{>\ell_\omega}\}.$$

With this notation, we now define the following axiom that is called *WEJR up to any seat from an overrepresented cohesive group (WEJR-X-r).*

Definition 5.7 (WEJR-X-r)**.** *A seat assignment $\boldsymbol{s}$ provides WEJR-X-r if for every ℓ_ω-cohesive group $N' \subseteq N$ such that there exists a candidate $c \in \bigcap_{i \in N'} A_i$ with*

$c \notin \boldsymbol{s}$, there exists a voter $i \in N'$ such that for every seat $t \in \boldsymbol{s}(C_{\ell_\omega} \setminus A_i)$, it holds that $sat_\omega(A_i, \boldsymbol{s}) + w_t > \ell_\omega$.

We follow by investigating what rules satisfy WEJR-X-r. First, we show that both MES[arbitrary] and MES[seq-Phragmén] do not satisfy this axiom.

Example 5.3 (MES's first phase fails WEJR-X-r)**.** This can be shown in the first phase of MES. Consider this example with two voters, 12 candidates and six available seats:

	$\{c_1, \dots, c_6\}$	$\{c_7, \dots, c_{12}\}$
Voter 1	☑	☒
Voter 2	☒	☑

$$\boldsymbol{w} = (3, 3, 1, 1, 1, 1), \omega = 10$$

In MES's initial setup, each voter has a budget of 3 so can afford to pay to assign candidates that they support to three of the seats. Given the assumption of arbitrary tie-breaking between voters within rounds, suppose the seat assignment returned by MES's first phase is $\boldsymbol{s} = (c_1, c_2, c_3, c_7, c_8, c_9)$. Each voter is a 5-cohesive group, voter 1 is an overrepresented group, and the weight-1 seat that was assigned to voter 1's approved candidate is not enough for voter 2 to exceed their deserved representation of 5.

Now, observe that the first phases of MES_ω and GCR_ω trivially satisfy this axiom as no ℓ_ω-cohesive groups can become overrepresented during the execution of those rules. Thus, it is of interest to study the WEJR-X-r axiom with respect to *complete* seat assignments. Negatively, we now show that MES_ω[seq-Phragmén_ω] fails WEJR-X-r (and so the same holds for MES_ω[arbitrary]).

Example 5.4 (MES_ω[seq-Phragmén_ω] fails WEJR-X-r)**.** Consider this example with three voters, 18 candidates and six seats:

	$\{c_1, \dots, c_6\}$	$\{c_7, \dots, c_{12}\}$	$\{c_{13}, \dots, c_{18}\}$
Voter 1	☑	☒	☒
Voter 2	☒	☑	☒
Voter 3	☒	☒	☑

$$\boldsymbol{w} = (63, 30, 3, 1, 1, 1), \omega = 99$$

Note that each voter is a 33-cohesive group and begins MES$_\omega$'s first phase with a personal budget of 33. For MES$_\omega$'s first phase, each voter can afford to fund any of the seats with weight less than 63 (with the lower weighted seats paid for in earlier rounds). Suppose that due to arbitrary tie-breaking between equally q-affordable assignments, MES$_\omega$'s first phase returns the partial seat assignment $(\_, c_{13}, c_7, c_1, c_2, c_3)$. At this point, WEJR-X-r cannot be provided as the only seat assignments that satisfy WEJR-X-r are those where both of the following hold: (i) the voter that is assigned the weight-63 seat is only assigned this single seat, and (ii) only a single voter has satisfaction greater than 33. We see that in completing the partial $(\_, c_{13}, c_7, c_1, c_2, c_3)$, it is not possible to satisfy both of these conditions.

For the sake of completion of the counterexample, we consider the MES$_\omega$ completion step using seq-Phragmén$_\omega$. Voters 1 and 2 start the seq-Phragmén$_\omega$ process with personal budgets of 30 while voter 3 begins with 3. Due to this, voters 1 and 2 will each accumulate 63 in funds before voter 3 and thus, one of these voters will pay for the final assignment to the weight-63 seat. Suppose that it is voter 1 due to tie-breaking and the returned seat assignment is $\boldsymbol{s} = (c_4, c_{13}, c_7, c_3, c_2, c_1)$. Then voter 1 would be an overrepresented 33-cohesive group but none of the weight-1 seats can be given to voters 2 and 3 such that they cross 33 in satisfaction. Thus, this seat assignment $\boldsymbol{s}$ does not provide WEJR-X-r.

It is worth noting that in the previous chapter, it was vital that the seats are assigned in non-increasing order so as to satisfy WLQ-X-r. So, it is reasonable to wonder whether a variant of MES$_\omega$, where seats are considered in this non-increasing order, may lead to a more positive result for WEJR-X-r. Unfortunately, this is not the case, with the same weighted election instance from Example 5.4 serving as a counterexample for such an MES$_\omega$ variant. Moreover, we can show the following for GCR$_\omega$[arbitrary] which, by design, assigns the seats in non-increasing order of weight.

Example 5.5 (GCR$_\omega$[arbitrary] fails WEJR-X-r)**.** This can be shown in the first phase of MES. Consider this example with two voters, 18 candidates and six seats:

	$\{c_1, \ldots, c_6\}$	$\{c_7, \ldots, c_{12}\}$	$\{c_{13}, \ldots, c_{18}\}$
Voter 1	☑	☒	☒
Voter 2	☒	☑	☒
Voter 3	☒	☒	☑

$$\boldsymbol{w} = (63, 30, 3, 1, 1, 1), \omega = 99$$

Note again that each voter is a 33-cohesive group. For GCR$_\omega$, the first phase returns the following partial seat assignment $(\_, c_1, c_2, c_7, c_8, c_9)$ (with tie-breaking in favour of voters 1 and 2). Suppose that arbitrary completion leads to the weight-63 seat being assigned to a candidate of voter 2 and thus, the seat assignment is $\boldsymbol{s} = (c_{10}, c_1, c_2, c_7, c_8, c_9)$. Then in this case, voter 2 is overrepresented and as above, them being assigned the weight-1 seats leads to a violation of WEJR-X-r.

Since we also know that seq-Phragmén$_\omega$ cannot satisfy WEJR-X-r (as it fails the weaker WEJR-1), we have exhausted the rules that we have defined thus far in our efforts to satisfy WEJR-X-r. So now, in showing that WEJR-X-r can always be satisfied by a rule that returns complete seat assignments, we turn to a WSAM that satisfies WLQ-X-r, namely the Greedy$_\omega$ method (see Definition 4.4), and we adapt it.

Definition 5.8 (GreedyX$_\omega$)**.** *This rule works in k rounds and assigns seats in non-increasing order. For each voter, we use $s_i(r)$ to denote the satisfaction that voter i gained from assignments to seats made in rounds before round r. Starting in round 1, in each round r, the rule finds the ℓ_ω-cohesive group N' with the largest value $d_r(N') = \ell_\omega - \max\{s_i(r) \mid i \in N'\}$, and assigns the seat r with weight w_r to some candidate approved by all voters in N'.*

So, whereas GCR$_\omega$ looks to satisfy the most underrepresented ℓ_ω-cohesive groups by assigning multiple seats to candidates approved by said groups, GreedyX$_\omega$ takes a more refined, seat-by-seat approach and with the following result, we find that this leads to improved proportionality guarantees.

Theorem 5.5. *GreedyX$_\omega$ satisfies WEJR-X-r.*

Proof. Take a seat assignment $\boldsymbol{s}$ returned by GreedyX$_\omega$ and assume that WEJR-X-r is violated. So there exists an ℓ_ω-cohesive group N' (where there exists a candidate $c \in \bigcap_{i \in N'} A_i$ with $c \notin \boldsymbol{s}$) such that there is no voter where $sat_\omega(A_i, \boldsymbol{s}) + w_t > \ell_\omega$ holds with w_t being the smallest weight of some seat assigned to another ℓ_ω-cohesive group that is overrepresented.

Take round h, where seat h with weight w_h is being assigned, to be the round where some ℓ^*_ω-cohesive group N^* became overrepresented. This means that with the addition of seat h we have that for some voter in N^*, their satisfaction crosses ℓ^*_ω and thus, $d_h(N^*) - w_h < 0$. But since the seat was assigned to a candidate of group N^* in round h, it must hold that $d_h(N^*) \geqslant d_h(N')$. From this, it follows that $0 > d_h(N^*) - w_h \geqslant d_h(N') - w_h$ and so with this seat h, adding it to the satisfaction of some voter in N' means that N' exceeds ℓ_ω. This also works for all seats assign for N^* prior to seat h as seats are assigned in non-increasing order. □

Thus, we find that GreedyX$_\omega$, up to this point of our study, is the best performing rule with respect to our weighted EJR adaptations. However, it still remains to find a WMWV rule that is polynomial-time computable and provides WEJR-X-r.

5.2.2 An Alternative Relaxation of WEJR

We can also offer an alternative relaxation and this is obtained by strengthening the ℓ_ω-cohesiveness requirement. This modification of ℓ_ω-cohesiveness further restricts the values of satisfaction a group may demand. For a weight vector $\boldsymbol{w}$, instead of restricting possible satisfaction to elements of the set $\mathcal{S}(\boldsymbol{w})$, we use the non-decreasing vector **low-**$\boldsymbol{\mathcal{S}}(\boldsymbol{w}) = (\ell_1, \ell_2, \dots, \ell_k)$ where $\ell_t = \sum_{j=1}^{t} w_j$. So, for every $t \in [k]$, the possible satisfaction is the sum of the t lowest weights.

Definition 5.9 (Lower ℓ_ω-cohesiveness, low-ℓ_ω-cohesiveness)**.** *For an integer $\ell_\omega \in$ **low-**$\boldsymbol{\mathcal{S}}(\boldsymbol{w})$, a set of voters $N' \subseteq N$ is low-ℓ_ω-cohesive if both of the following conditions hold:*

- $|N'| \geqslant \ell_\omega \cdot {}^{n}/_{\omega}$.
- *There exists a set of candidates $C' \subseteq \bigcap_{i \in N'} A_i$ with size $|C'| = t$ such that there exists a seat assignment $\boldsymbol{s}$ where $sat_\omega(C', \boldsymbol{s}) \geqslant \ell_\omega$ and $|N'| \geqslant n \cdot {}^{t}/_{k}$.*

The intuition this captures is that a group N' already represented by $t-1$ seats, that is demanding an additional t-th seat of certain weight, must not only be large enough to demand t of the k seats, but N' must already be large enough to demand the sum of the $t-1$ lowest weights.

Example 5.6 (Illustrating low-ℓ_ω-cohesiveness)**.** For example, the only low-ℓ_ω-cohesive groups that can form, when $\boldsymbol{w} = (5, 3, 3, 1)$ is the weight vector, are those for weight $\ell_\omega \in \{1, 4, 7, 12\}$.

This modified cohesiveness notion focused on the lowest weighted seats leads us to the following weaker version of the WEJR property.

Definition 5.10 (Lower Weighted EJR, low-WEJR)**.** *A seat assignment $\boldsymbol{s}$ provides low-WEJR if for every low-ℓ_ω-cohesive group $N' \subseteq N$ such that there exists a candidate $c \in \bigcap_{i \in N'} A_i$ with $c \notin \boldsymbol{s}$, there exists a voter $i \in N'$ such that $sat_\omega(A_i, \boldsymbol{s}) \geqslant \ell_\omega$.*

Example 5.7. Consider the counterexample from Theorem 4.2, now adapted to the weighted multiwinner setting. There are three voters, nine candidates and three seats to be assigned:

	$\{a, b, c\}$	$\{d, e, f\}$	$\{g, h, i\}$
Voter 1	☑	☒	☒
Voter 2	☒	☑	☒
Voter 3	☒	☒	☑

$$\boldsymbol{w} = (3, 2, 1), \omega = 6$$

Given the weight vector $\boldsymbol{w} = (3, 2, 1)$, we have that $\textbf{low-}\boldsymbol{\mathcal{S}}(\boldsymbol{w}) = (1, 3, 6)$. Note that the single voters cannot each, on their own, form a low-ℓ_ω-cohesive group for a weight $\ell_\omega > 1$, whereas, they each represent a 2-cohesive group.

Next, we show that seat assignments that provide low-WEJR can always be produced.

Proposition 5.6. *A seat-based WMWV rule based on an approval-based multiwinner voting rule F_α that satisfies EJR also satisfies low-WEJR.*

Proof. Take $\textbf{low-}\boldsymbol{\mathcal{S}}(\boldsymbol{w}) = (\ell_1, \ell_2, \dots, \ell_k)$ for some weight vector $\boldsymbol{w}$. Note that if a group of voters N' is low-ℓ_t-cohesive for some $\ell_t \in \textbf{low-}\boldsymbol{\mathcal{S}}(\boldsymbol{w})$, then they would deserve at least t seats. Thus, from the definition of low-WEJR, this group may demand exactly ℓ_t in weight. So if this voter group were represented by *any* t seats, then they would receive at least their ℓ_t in satisfaction, given that ℓ_t is the minimum possible weight that t committee seats can yield.

Now, consider a weighted election instance $(\boldsymbol{A}, \boldsymbol{w})$. Note that every low-$\ell_t$-cohesive group for some $\ell_t \in \textbf{low-}\boldsymbol{\mathcal{S}}(\boldsymbol{w})$ is also a t-cohesive group as in definition from standard approval-based multiwinner voting (see Definition 2.8). Thus, for any weighted election instance $(\boldsymbol{A}, \boldsymbol{w})$, to find a seat assignment $\boldsymbol{s}$ providing low-WEJR, we can use any approval-based rule that satisfies EJR, so it will provide t-cohesive groups with sufficient representation from t seats.

Note that the seats need not be assigned in non-increasing order to satisfy low-WEJR. □

In this respect, low-WEJR does not appear to be a strong axiom for our setting as we can use a standard approval-based multiwinner voting rule such as MES[arbitrary] in order to achieve it. However, we find with the following example that our weighted adaptation of MES[arbitrary] fails low-WEJR. And in the fact, the counterexample also shows that GCR_ω[arbitrary] fails this axiom as well.

Example 5.8 (MES_ω[arbitrary] and GCR_ω[arbitrary] fail low-WEJR)**.** Consider again the following example with three voters, nine candidates and three seats.

	$\{a,b,c\}$	$\{d,e,f\}$	$\{g,h,i\}$
Voter 1	☑	☒	☒
Voter 2	☒	☑	☒
Voter 3	☒	☒	☑

$$\boldsymbol{w} = (3, 2, 1), \omega = 6$$

Note that each voter represents a low-1-cohesive group. Suppose that the first phases of these rules return the partial seat assignment $(\_, a, d)$. Then with arbitrary completion, suppose that the rules return the seat assignment $\boldsymbol{s} = (b, a, d)$. Here, voter 3 falls short of their deserved satisfaction of 1.

It immediately follows that WEJR-1 does not imply low-WEJR despite both being equivalent to EJR in the standard, unweighted multiwinner setting.

Corollary 5.7. *WEJR-1 does not imply low-WEJR.*

We also point out that as low-WEJR is equivalent to EJR in the standard multiwinner setting, we have that seq-Phragmén_ω fails low-WEJR as the unweighted seq-Phragmén fails to satisfy EJR.

5.3 Priceability with Weights

Next, we introduce an adaptation of the priceability notion (introduced in the Background chapter).

Definition 5.11 (ω-Priceability)**.** *Suppose that voters have a personal budget of b and they spend funds to assign candidates that they approve of, to seats with weights from $\boldsymbol{w}$. A* weighted price system $\boldsymbol{ps}_\omega = (b, \{p_i\}_{i\in[n]})$ *is a pair where b*

is the individual budget *per voter and each voter* $i \in N$ *has a* payment function $p_i : C \to [0, b]$. *A weighted price system* $\boldsymbol{ps}_\omega = (b, \{p_i\}_{i\in[n]})$ *is* ω-affordable *for a seat assignment* $\boldsymbol{s} = (s_1, \ldots, s_k)$ *if the following conditions hold:*

C1. *If* $p_i(c) > 0$*, then* $c \in A_i$ *(a voter only pay for candidates she approves of).*

C2. $\sum_{c\in C} p_i(c) \leqslant b$ *(a voter does not spend more than her budget).*

C3. *For every* $s_t \in \boldsymbol{s}$, $\sum_{i\in[n]} p_i(s_t) = w_t$ *(payments for this candidate must be equal to the weight of the seat they were assigned to in* $\boldsymbol{s}$*).*

C4. $\sum_{i\in[n]} p_i(s) = 0$ *for every candidate* $s \notin \boldsymbol{s}$ *(candidates that were not assigned a seat in the seat assignment are not paid for).*

We then have two additional conditions for a price system $\boldsymbol{ps}_\omega$*:*

C5. *For every* $s \notin \boldsymbol{s}$*, it holds that:*

$$\sum_{i\in N(s)} \left(b - \sum_{s'\in \boldsymbol{s}} p_i(s')\right) \leqslant \max\{w_1, \ldots, w_k\}$$

(supporters of any unelected candidate do not collectively hold, in terms of their unspent budget, more than the highest price).

C6. *For every* $s \notin \boldsymbol{s}$*, it holds that:*

$$\sum_{i\in N(s)} \left(b - \sum_{s'\in \boldsymbol{s}} p_i(s')\right) \leqslant \min\{w_1, \ldots, w_k\}$$

(supporters of any unelected candidate do not collectively hold, in terms of their unspent budget, more than any price).

We then say that a weighted price system $\boldsymbol{ps}_\omega$ *that is* ω*-affordable* weakly supports *a seat assignment* $\boldsymbol{s}$ *if it also satisfies condition* C5 *and it* strongly supports *a seat assignment* $\boldsymbol{s}$ *if it satisfies also condition* C6*. If there exists a weighted price system* $\boldsymbol{ps}_\omega$ *that weakly (or strongly) supports a seat assignment* $\boldsymbol{s}$ *then we say this seat assignment* $\boldsymbol{s}$ *is weakly (or strongly)* ω-priceable *for weight vector* $\boldsymbol{w}$*.*

It is clear that a seat assignment that is strongly ω-priceable is also weakly ω-priceable, and both notions reduce to Priceability in the unweighted multiwinner model (see Definition 2.12). In their definition, we focused on the max and min operators as they capture intuitions similar to the 'up-to-one' and 'up-to-any' relaxations that we have touched on thus far. Moving forward, we consider this property for complete seat assignments.

For standard multiwinner voting, we know that priceable committees provide PJR (see Chapter 2). We now make a similar connection between strong ω-priceablility and the following PJR-like property.

Definition 5.12 (WPJR up to any seat, WPJR-X)**.** *A seat assignment $\boldsymbol{s}$ provides WPJR-X if for every ℓ_ω-cohesive group $N' \subseteq N$, it holds for every seat $t \in \boldsymbol{s}(C \setminus \bigcup_{i\in N'} A_i)$ that $sat_\omega(\bigcup_{i\in N'} A_i, \boldsymbol{s}) + w_t > \ell_\omega$.*

This next result shows that strong ω-priceablility implies WPJR-X.

Proposition 5.8. *A strongly ω-priceable seat assignment $\boldsymbol{s}$ provides WPJR-X.*

Proof. Take a seat assignment $\boldsymbol{s}$ that is strongly ω-priceable and assume that WPJR-X is violated. So there must exist a ℓ_ω-cohesive group $N' \subseteq N$ such that $sat_\omega(\bigcup_{i\in N'} A_i, \boldsymbol{s}) \leqslant \ell_\omega - w_t$ where w_t is the seat with smallest weight amongst the seats not assigned to any candidate approved by some member of N'. As $\boldsymbol{s}$ that is strongly ω-priceable, there exists a price system such that all seats are paid for. So the entire population can afford to spend ω, thus the personal budget b per voter is at least ${}^{\omega}/_{n}$.

Now, since N' has a sum of satisfaction of at most $\ell_\omega - w_t$, they spent at most this amount in funds in the price system. Thus, the remaining voters $N \setminus N'$ must have spent the remaining funds $\omega - (\ell_\omega - w_t)$. Then since N' is ℓ_ω-cohesive group, we know that $|N'| \cdot {}^{\omega}/_{n} \geqslant \ell_\omega$, we find then find that the budget per voter b of voters in $N \setminus N'$ is:

$$\frac{\omega - (\ell_\omega - w_t)}{n - |N'|} > \frac{\omega - \ell_\omega}{n - |N'|} \geqslant \frac{\omega - |N'| \cdot {}^{\omega}/_{n}}{n - |N'|} \geqslant \frac{\omega(n - |N'|)}{n(n - |N'|)} = \frac{\omega}{n}.$$

So each voter in $N \setminus N'$ had a personal budget b strictly greater than ${}^{\omega}/_{n}$, so this must hold of voters in N'. And since they spent at most $\ell_\omega - w_t$ collectively, they must have at least the following in unspent budget as a group:

$$b \cdot |N'| - (\ell_\omega - w_t) > \frac{\omega}{n} \cdot |N'| - (\ell_\omega - w_t) \geqslant \ell_\omega - \ell_\omega + w_t \geqslant w_t.$$

Thus, we have a contradiction that $\boldsymbol{s}$ is strongly ω-priceable as the voters in N' have more than $w_t \geqslant \min\{w_1, \dots, w_k\}$ in funds which violates condition C6 of strong ω-priceability. □

Given this result, the counterexample of Proposition 4.7 (that shows that a seat assignment providing WLQ-X may not exist) can also be used to show that WPJR-X is not always satisfiable.

Proposition 5.9. *A seat assignment $\boldsymbol{s}$ that provides WPJR-X may not exist.*

This spells bad news for the task of satisfying strong ω-priceability as we find that we cannot always return a seat assignment that is strongly ω-priceable (this is illustrated in Example 5.9 below).

Corollary 5.10. *There exists a weighted election instance where no strongly ω-priceable seat assignment exists.*

Example 5.9. [A strongly ω-priceable seat assignment may not exist] We adapt the aforementioned counterexample for Proposition 4.7 (that shows that a seat assignment providing WLQ-X may not exist) that is also used in Example 5.4.

	$\{c_1,\dots,c_6\}$	$\{c_7,\dots,c_{12}\}$	$\{c_{13},\dots,c_{18}\}$
Voter 1	☑	☒	☒
Voter 2	☒	☑	☒
Voter 3	☒	☒	☑

$$\boldsymbol{w} = (63, 30, 3, 1, 1, 1), \omega = 99$$

In order for a price system to be ω-affordable, each voter requires a personal budget of 63. However, this would lead to violation of condition C6 as the voter who does not pay for the seats of weight 63 and 30, will have more in unspent funds than the lowest price $\min\{w_1, \dots, w_k\}$. Thus, there is no seat assignment that is strongly ω-priceable for this weighted election instance.

From this result, we now know that WEJR-X-r does not imply WPJR-X. Thus, we move on by examining a weakening of WPJR-X.

Definition 5.13 (WPJR up to one seat, WPJR-1)**.** *We say that a seat assignment $\boldsymbol{s}$ provides WPJR-1 if for every ℓ_ω-cohesive group $N' \subseteq N$, there exists some seat $t \in \boldsymbol{s}(C \setminus \bigcup_{i\in N'} A_i)$ such that $sat_\omega(\bigcup_{i\in N'} A_i, \boldsymbol{s}) + w_t > \ell_\omega$.*

As WEJR-1 implies WPJR-1, we know that rules such as MES_ω[seq-Phragmén_ω] and GCR_ω[arbitrary] satisfy WPJR-1. Now, we move on to the weaker notion of weak ω-priceability and show that it implies the following property. We can then relate WPJR-1 to weak ω-priceability with the following result, the proof of which is omitted as it is analogous to that of Proposition 5.8.

Proposition 5.11. *A weakly ω-priceable seat assignment $\boldsymbol{s}$ provides WPJR-1.*

We now ask which of the rules that we have studied always return weakly ω-priceable seat assignments. As in the multiwinner voting literature, the method used to complete the partial outcomes is important when looking to return priceable outcomes. In particular, MES_ω[arbitrary] and GCR_ω[arbitrary] do not always return weakly ω-priceable seat assignments as can be seen in with the seat assignments they produced in Example 5.8. As a consequence, we have that weak ω-priceability does not follow from WEJR-1, which mirrors the relationship between EJR and priceability in the standard multiwinner context. And with reasoning that is analogous to that used to show that seq-Phragmén is priceable, we find that seq-Phragmén$_\omega$ is weakly ω-priceable. The proof is omitted as it is almost identical to that of Peters and Skowron (2020).

Proposition 5.12. *seq-Phragmén$_\omega$ always returns weakly ω-priceable seat assignments.*

In an immediate corollary, we find that seq-Phragmén$_\omega$ satisfies WPJR-1.

Corollary 5.13. *seq-Phragmén$_\omega$ satisfies WPJR-1.*

Proposition 5.12 not only gives us a rule that returns weakly ω-priceable outcomes, but we now also have a way to complete partial seat assignments returned by MES_ω in such a way as to ensure that they are weakly ω-priceable.

Proposition 5.14. *MES_ω[seq-Phragmén$_\omega$] always return weakly ω-priceable seat assignments.*

Now, with the connection between WEJR-1 and WPJR-1 being clear, the following example clarifies how low-WEJR relates to WPJR-1 by showing that the former does not imply the latter (and thus, low-WEJR does not imply weak ω-priceability).

Example 5.10 (low-WEJR does not imply WPJR-1)**.** Consider this example with two voters, ten candidates and five seats.

	$\{c_1, \dots, c_5\}$	$\{c_6, \dots, c_{10}\}$
Voter 1	☑	☒
Voter 2	☒	☑

$$\boldsymbol{w} = (3, 3, 2, 1, 1), \omega = 10$$

Each voter qualifies as a 5-cohesive group. Now, consider the seat assignment $\boldsymbol{s} = (c_1, c_2, c_3, c_6, c_7)$. This seat assignment provides low-WEJR as each voter reaches their deserved weight satisfaction of 2. However, since voter 2 has a satisfaction of 2, there is no unassigned seat such that they can surpass the value of 5 as is required by WPJR-1.

We now assess whether the rules that we have seen to satisfy low-WEJR, such as MES[arbitrary] and MES[seq-Phragmén], also return weakly ω-priceable seat assignments.

Example 5.11 (MES's first phase fails weak ω-priceability)**.** This can be shown in the first phase of MES. Consider this example with two voters, eight candidates and four seats.

	$\{c_1, \ldots, c_4\}$	$\{c_5, \ldots, c_8\}$
Voter 1	✓	✗
Voter 2	✗	✓

$$\boldsymbol{w} = (3, 3, 1, 1), \omega = 8$$

In MES's initial setup, each voter has a budget of 2 so can each afford to pay to assign candidates that they support to two of the seats. Given the assumption of arbitrary tie-breaking between voters within rounds, suppose the seat assignment returned by MES's first phase is $\boldsymbol{s} = (c_1, c_2, c_5, c_6)$. This is not weakly ω-priceable as voter 1 would be require a budget of 6 but then voter 2 would have a remaining budget of 4 which is strictly greater than the largest price of 3.

In fact, both these rules fail the weaker WPJR-1 as is shown in Example 5.2. We also find that GreedyX$_\omega$ does not meet the requirements of weak ω-priceability.

Example 5.12 (GreedyX$_\omega$ fails weak ω-priceability)**.** Consider this example, which is adapted from the standard multiwinner voting setting (Peters and Skowron, 2020). Here, we have 15 candidates, six voters and 12 seats available.

	$\{c_i\}_{i\in[3]}$	c_4	c_5	c_6	$\{c_{i+6}\}_{i\in[3]}$	$\{c_{i+9}\}_{i\in[3]}$	$\{c_{i+12}\}_{i\in[3]}$
Voter 1	✓	✓	✗	✗	✗	✗	✗
Voter 2	✓	✗	✓	✗	✗	✗	✗
Voter 3	✓	✗	✗	✓	✗	✗	✗
Voter 4	✗	✗	✗	✗	✓	✗	✗
Voter 5	✗	✗	✗	✗	✗	✓	✗
Voter 6	✗	✗	✗	✗	✗	✗	✓

$$\boldsymbol{w} = (1, 1, \ldots, 1), \omega = 12$$

Here, each voter is a 2-cohesive group while voters $\{1, 2, 3\}$ are a 3-cohesive group. For GreedyX$_\omega$, the first seat is assigned to candidate collectively approved by the voters $\{1, 2, 3\}$. Then next eight rounds will see two seats assigned to candidates supported by voters $\{1, 2, 3\}$, voter 4, voter 5 and voter 6. At this stage of GreedyX$_\omega$, voters $\{1, 2, 3\}$, voter 4, voter 5 and voter 6 are each equally entitled to an additional seat with there being 3 seats left to assign to. Suppose that the chosen tie-breaking mechanism favours the individual voters and suppose that the seat assignment returned by GreedyX$_\omega$ is $\boldsymbol{s} = (c_1, c_2, c_3, c_7, c_8, c_9, c_{10}, c_{11}, c_{12}, c_{13}, c_{14}, c_{15})$. The equivalent unweighted multiwinner committee that corresponds to this seat assignment $\boldsymbol{s}$ is not priceable (Lackner and Skowron, 2023), and thus, this seat assignment is not weakly ω-priceable.

This result leaves still searching for a WMWV rule that satisfies both WEJR-X-r and weak ω-priceability (under the assumption that both axioms can always be satisfied simultaneously). A taxonomy of the proportionality axioms studied in this chapter as well as the relations between them can be seen in Figure 5.1 below.

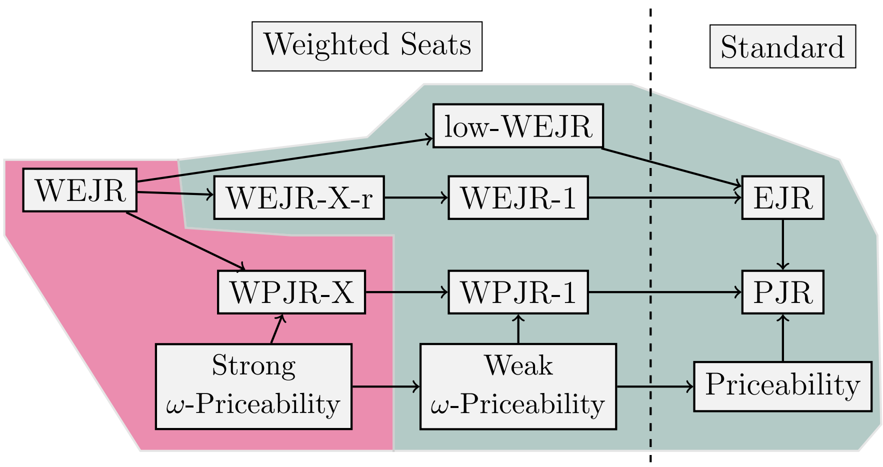

Figure 5.1: The relations between the weighted multiwinner voting axioms studied within this chapter along with the standard multiwinner voting axioms defined in Chapter 2. In this figure, an arrow ($\longrightarrow$) from axiom A pointing towards axiom B indicates that axiom A implies axiom B. The backgrounds represent whether an axiom can, or cannot, always be satisfied. A green background (■) represents the former while the red background (■) represent the latter.

	WEJR-X-r	WEJR-1	low-WEJR	WPJR-1	Weak ω-Priceability
MES[arbitrary]	✗	✗	✓	✗	✗
MES[seq-Phrag]	✗	✗	✓	✗	✗
seq-Phragmén$_\omega$	✗	✗	✗	✓	✓
MES$_\omega$[arbitrary]	✗	✓	✗	✓	✗
MES$_\omega$[seq-Phrag$_\omega$]	✗	✓	◇	✓	✓
GCR$_\omega$[arbitrary]	✗	✓	✗	✓	✗
GreedyX$_\omega$	✓	✓	◇	✓	✗

Table 5.1: Summary that indicates for certain axioms whether a WMWV rule satisfies it (✓), fails it (✗) or whether this question is still open (◇). Note that MES[seq-Phragmén] and MES$_\omega$[seq-Phragmén$_\omega$] were shortened to MES[seq-Phrag] and MES$_\omega$[seq-Phrag$_\omega$], respectively, to narrow the table.

5.4 Chapter Summary

This chapter saw us introduce a model of multiwinner voting that is enriched with weights over the committee seats. We subsequently explored various adaptations of the EJR axiom with the strongest, namely WEJR, proving too strong to always be satisfiable in general. However, with WEJR-1 and WEJR-X-r, we identified relaxations of WEJR that can always be satisfied with the former even being satisfiable by a polynomial-time computable rule. We followed with low-WEJR which represented an alternative path in relaxing WEJR. Despite finding mostly negative results regarding low-WEJR and the weight-based rules, it displayed that the seat-based rules do have a role to play in ensuring proportional representation in this setting. The chapter's analysis closed with the study of weighted variants of priceability. Our study of both strong and weak ω-priceability also resulted in an investigation of weighted PJR-like notions and this led us to increased insights on the nature of the WMWV rules. Notably, the weaker of the two ω-priceability axioms put a more positive light on our weighted versions of seq-Phragmén and MES (with regards to producing proportional outcomes).

Future Work. We finish off the chapter by detailing avenues for follow-up research. Of the technical questions that are left over, we deem the most pertinent to be the following two:

- Is there a polynomial-time computable rule that provides WEJR-X-r?
- Is there a weakly ω-priceable rule that provides WEJR-X-r?

Further work that advances on our research line could incorporate the study of other proportionality notions such as proportionality degree or look at adapting other multiwinner voting rules with PAV being an obvious candidate. We have focused on ensuring that cohesive groups are not underrepresented in this weighted multiwinner setting but have thus far neglected to study a multiwinner analogue of the apportionment notion of upper quota. While there have been some considerations made within the literature for dealing with *overrepresentation* (Cevallos and Stewart, 2021; Sánchez-Fernández et al., 2024; Boehmer et al., 2024), none generalise the upper quota to multiwinner voting. Let alone for our weighted setting, this would be interesting to study for standard multiwinner voting.

In Chapter 4, we also looked at house monotonicity (HM). The multiwinner analogue would be the axiom of *Committee Monotonicity (CM)* (Lackner and Skowron, 2023). We omitted the study of a weighted version of CM in this chapter because there is no known approval-based multiwinner voting rule that satisfies both CM and EJR, which is unlike the apportionment case where the two divisor methods we studied (D'Hondt and Adams) satisfy HM alongside a proportionality axiom (lower and upper quota, respectively).

Lastly, we point out that an experimental analysis may also be fruitful avenue for obtaining insights on the weighted proportionality properties and the WMWV rules that we have introduced. Such an analysis may consist of questions such as how often do the WMWV rules satisfy the proportionality properties in cases where the rule does not satisfy it, or where the property is not always satisfiable, in general.

Detour

Chapter 6

Simulating Multiwinner Voting Rules (Beyond Approval Ballots) in Judgment Aggregation

Up to this point in the thesis, the task of lifting proportionality to more complex domains has proven to be challenging. In this chapter, we stray from this aim and instead, look towards gaining a further understanding of multiwinner voting rules, proportional or otherwise. To improve our understanding of the mechanisms that have been proposed for different settings and to enable us to transfer some of the knowledge gained in one domain to another domain of collective decision making, it is important to isolate fundamental building blocks that are common to different solutions.

To do so, we take an approach used for single-winner voting rules (Endriss, 2018) and investigate the extent to which it is possible to simulate multiwinner voting rules within the framework of logic-based *judgment aggregation* (JA) (List and Pettit, 2002; Grossi and Pigozzi, 2014; Endriss, 2016). Moreover, our analysis shall not be limited to the approval-based multiwinner voting rules that we have explored thus far, but also to the more general multiwinner voting rules that use *preorders* as the ballots. The advantage of this is that we are able to extract structural insights on multiwinner voting rules that use not only approval ballots, but also *strict ordinal preferences* (where voters provide a strict ranking of the candidates) (Faliszewski et al., 2017). Regarding the multiwinner election outcomes, we focus on the standard scenario where a committee of fixed size k is to be elected (as seen often in this thesis).

Embedding a multiwinner voting rule into judgment aggregation, which is much more expressive than most voting frameworks, makes it very natural to study refinements of standard rules in a principled manner, e.g., by imposing additional constraints on outcomes or by varying the types of preferences voters can report. This not only permits us to clarify the commonalities (and differences) between rules originally developed for different purposes, but also allows for the development of new rules with particular properties.

For our embeddings of multiwinner voting rules, we use the particular framework of judgment aggregation with *rationality* and *feasibility constraints* (Endriss, 2018). The former constrains the range of admissible inputs to an aggregation problem, and the latter constrains its possible outputs. This leads to particularly simple and natural embeddings. Importantly though, the increased expressive power of judgment aggregation means that embeddings come at a cost: judgment aggregation is a framework that, generally speaking, is computationally much more demanding than voting. To address this challenge, we analyse the extent to which the feasibility constraints featuring in our embeddings can be encoded as Boolean circuits in *decomposable negation normal form*, which by a recent result allows for the design of tractable aggregation rules (De Haan, 2018).

Chapter Outline. We start in Section 6.1 by presenting the models of multiwinner voting with preorders as voters' ballots and also, we define some multiwinner voting rules that use these preorders. This is followed by introducing the model of judgment aggregation with rationality and feasibility constraints as well as some JA rules that we focus on (in Section 6.2). Next, in Section 6.3, we (*i*) propose a way to model preference aggregation in the JA model with rationality and feasibility constraints, and (*ii*) propose a way to simulate multiwinner voting rules using this particular JA model and (*iii*) detail various results related to this JA simulation of some multiwinner voting rules. We then follow with Section 6.4 where we use DNNF circuit encodings to obtain computational results on constraints and rules. We then provide the chapter's conclusion in the final section (Section 6.5).

6.1 Multiwinner Voting Beyond Approvals

Thus far, we have only considered collective-choice scenarios where participating voters submit their preferences as approval ballots. This was especially desirable given the simplicity of the ballots, but in this chapter, we take the more general approach of using preorders as ballots. As mentioned in this chapter's introduction, this allows us to consider, as special cases, not only approval-based multiwinner rules, but also rules for either strict and weak ordinal preferences. Another advantage of going beyond weak preferences and using preorders is that voters can express not only *indifference* between two candidates, but voters can also express the *incomparability* of a candidate pair.

6.1.1 The Model

In this model, our usual set of voters $N = \{1, \ldots, n\}$ submits their preferences over a m-sized set of candidates $C = \{a, b, c, \ldots\}$. Each voter $i \in N$ has a *preference preorder* $\succcurlyeq_i$ over the candidates, where for two candidates $c, d \in C$, we have that $c \succcurlyeq_i d$ indicates that voter i weakly prefers c over d. For a pair of candidates $c, d \in C$, a voter i is said to be *indifferent* between c and d if we have both $c \succcurlyeq_i d$ and $d \succcurlyeq_i c$. On the other hand, we say voter i finds candidates c and d to be *incomparable* if neither $c \succcurlyeq_i d$ and $d \succcurlyeq_i c$ hold. We then use $\succ_i$ to denote the strict part of $\succcurlyeq_i$, so $c \succ_i d$ holds if voter i strictly prefers c to d, i.e., $c \succcurlyeq_i d$ holds but not $d \succcurlyeq_i c$ holds. Let $\mathcal{R}(C)$ be the set of all possible preorders.

A *preference profile* is then a vector $\boldsymbol{P} = (\succcurlyeq_1, \ldots, \succcurlyeq_n)$ of the voters' n preferences preorders. Then for a fixed committee target size k, we say an election instance is a pair $(\boldsymbol{P}, k)$. The goal once again is to elect a committee W of k candidates. Now, we define rules for multiwinner elections that take preference preorders as input. So, a *multiwinner voting rule* is a function $F : (\mathcal{R}(C))^n \rightarrow \mathcal{P}_+(C)$ mapping a preference profile $\boldsymbol{P}$ to a set of winning committees (so, again, F is irresolute).

6.1.2 Ranking-based Multiwinner Rules

As we now work with voter preferences as preorders, we can expand our study of multiwinner voting rules to those that handle more than just approval ballots. Having already defined some approval-based multiwinner voting rules in the Background chapter, we now present some multiwinner voting rules designed to deal with weak ordinal preferences (Faliszewski et al., 2017; Aziz et al., 2017b).

Endriss (2018) simulates some prominent, single-winner voting rules, namely the *Borda* and *Copeland* rules (Zwicker, 2016). These results prompt us to aim our initial focus on some multiwinner analogues of two of these single-winner methods. The first is based on the Borda rule which, in the single-winner case, elects the candidate with the highest Borda score. The Borda score is defined for a candidate $x \in C$ as:

$$B(x, \boldsymbol{P}) = \sum_{i \in N} |\{y \in C \mid x \succ_i y\}|.$$

Definition 6.1 (k-Borda)**.** *The rule elects the committee/s of the k candidates with the highest Borda scores.*

The next two rules both represent ways to extend single-winner Copeland to the multiwinner setting. To define them, we must first define the *majority relation* $\succcurlyeq_M$ to be such that $x \succcurlyeq_M y$ holds for a pair of candidates $x, y \in C$ if and only if $|\{i \in N \mid x \succcurlyeq_i y\}| \geqslant {}^n\!/_2$. We then write $x \succ_M y$ if we have that $x \succcurlyeq_M y$ and

$y \not\succcurlyeq_M x$, while if both $x \succcurlyeq_M y$ and $y \succcurlyeq_M x$ hold, we write $x \sim_M y$. Now, for a candidate $x \in C$ and a value $\tau \in [0, 1]$, we define the Copeland score to be:

$$C^\tau(x, \boldsymbol{P}) = |\{y \in C \mid x \succ_M y\}| + \tau \cdot |\{y \in C \mid x \sim_M y\}|.$$

Intuitively, a candidate x's Copeland score is determined by their performance in majority pairwise contests against other candidates: they gain a point for every victory and for all tied pairwise contests, the value τ represents the score given in such cases. We shall consider two options for this value τ which leads to two different rules (Aziz et al., 2017b).

Definition 6.2 (k-Copeland^0)**.** *The rule elects the committee/s of the k candidates with the highest Copeland score/s for $\tau = 0$.*

Definition 6.3 (k-Copeland^1)**.** *The rule elects the committee/s of the k candidates with the highest Copeland score/s for $\tau = 1$.*

Note that if for every pair of candidates $x, y \in C$, exactly one of $x \succcurlyeq_M y$ and $y \succcurlyeq_M x$ holds (such as when n is odd), then k-Copeland^0 and k-Copeland^1 coincide. In such cases, we simply use the name k-Copeland to refer to both k-Copeland^0 and k-Copeland^1 simultaneously.

6.2 Judgment Aggregation

We shall use the framework of judgment aggregation in our analysis. Typically, judgment aggregation deals with collective choice on binary issues where the feasible outcomes on these are restricted by some logical constraint. This results in the framework being widely applicable and it has garnered research interest from a variety of research areas such as: computer science (Grossi and Pigozzi, 2014; Baumeister et al., 2016; Endriss, 2016); philosophy and economics (List and Pettit, 2002; Dietrich, 2006; Nehring and Puppe, 2008; Dokow and Holzman, 2010; Pauly and van Hees, 2006); and law with the classical doctrinal paradox originating from this legal sphere (Kornhauser and Sager, 1983). Now, we also turn to judgment aggregation for our aim of simulating a variety of multiwinner voting rules.

The idea of modelling problems of preference aggregation within the framework of judgment aggregation by means of a so-called preference agenda goes back to, at least, the work of Dietrich and List (2007). While a number of authors, such as Miller and Osherson (2009) and Dietrich (2014), have discussed parallels between specific voting rules and specific JA rules, Lang and Slavkovik (2013) were the first to systematically investigate the question of how to translate common voting rules into judgment aggregation. Endriss (2018) refined their approach and showed that explicitly distinguishing between rationality and feasibility constraints in judgment aggregation greatly simplifies the task of arriving at principled embeddings. We use the same basic approach also here.

6.2.1 The Model

Let $N = \{1, \ldots, n\}$ be a set of voters. We ask each voter to either accept or reject each of the issues in the *agenda* Φ, a finite set of propositional atoms which we refer to as *propositions*.

A *judgment* is a function $J : \Phi \to \{0, 1\}$, where acceptance of an agenda item is represented by 1 and rejection by 0. For any two judgments J and J', we use $\text{Agr}(J, J') = \{\varphi \in \Phi \mid J(\varphi) = J'(\varphi)\}$ to refer to the set of agenda items that they agree on and $\text{Dis}(J, J') = \{\varphi \in \Phi \mid J(\varphi) \neq J'(\varphi)\}$ to refer to the set of those they disagree on. A *judgment profile* is a vector $\boldsymbol{J} = (J_1, \ldots, J_n) \in (\{0, 1\}^\Phi)^n$ of judgments, one for each voter. The intensity of the *support* of an issue $\varphi \in \Phi$ in a profile $\boldsymbol{J}$ is denoted as $n_{(\boldsymbol{J},\varphi)} = |\{i \in N \mid J_i(\varphi) = 1\}|$. Given a judgment profile $\boldsymbol{J}$, we can define the *majority judgment* for each $\varphi \in \Phi$ as follows: $\text{Maj}(\boldsymbol{J})(\varphi) = 1$ if $n_{(\boldsymbol{J},\varphi)} \geqslant n/2$ and $\text{Maj}(\boldsymbol{J})(\varphi) = 0$ otherwise.

An *aggregation rule* is a function F_{JA} that takes as input a profile and outputs a judgment that is meant to represent a reasonable choice for a collective judgment. Although, ideally, this rule returns a single judgment, most rules allow for a tie between judgments, thereby possibly returning a set of judgments. Thus, formally, an aggregation rule is a function $F_{\text{JA}} : (\{0, 1\}^\Phi)^n \to \mathcal{P}_+(\{0, 1\}^\Phi)$ that maps any given profile to a nonempty set of judgments. An example of such a function is the majority rule defined as $F_{\text{JA}} : \boldsymbol{J} \mapsto \{\text{Maj}(\boldsymbol{J})\}$.

Let $\mathcal{L}(\Phi)$ denote the propositional language with the set of agenda items in Φ taking the role of propositional variables. We use formulas in this language to express constraints Γ regarding judgments: $J \models \Gamma$ holds if Γ is true under the truth assignment corresponding to J. The set of all judgments that satisfy a given constraint Γ is $\text{Mod}(\Gamma) = \{J \in \{0, 1\}^\Phi \mid J \models \Gamma\}$. Following Endriss (2018), we use such constraints both to express *rationality* constraints, i.e., constraints indicating an acceptable input to an aggregation rule, and *feasibility* constraints, i.e., constraints indicating an acceptable output. We then say that an aggregation rule F_{JA} *guarantees* Γ'*-feasible outcomes on* Γ*-rational profiles*, if $F_{\text{JA}}(\boldsymbol{J}) \subseteq \text{Mod}(\Gamma')$ holds for every profile $\boldsymbol{J} \in \text{Mod}(\Gamma)^n$.

6.2.2 Majoritarian Rules

We now detail some well-known JA rules that we will make use of. All of these rules shall guarantee, by definition, that the output will satisfy a given feasibility constraint Γ'. These particular rules are defined with this feature as we know that standard JA rules, such as the majority rule, do not ensure the return of Γ'-consistent outcomes (Kornhauser and Sager, 1983), unless the rationality and feasibility constraints satisfy certain conditions (Endriss, 2018; Grandi and Endriss, 2013).

We adopt the naming conventions used by Endriss (2018) but also mention some alternative names used to refer to these rules in the literature.

The *max-num* rule, also known as the *endpoint* rule (Miller and Osherson, 2009) and the *generalised Slater* rule (Nehring et al., 2014), selects judgments for which the number of agreements with the majority outcome is maximal:

Definition 6.4 (max-num)**.** *Given a judgment profile* $\boldsymbol{J}$ *and a feasibility constraint* Γ'*, the* max-num *rule is defined as:*

$$\text{max-num}(\boldsymbol{J}, \Gamma') = \underset{J \in \text{Mod}(\Gamma')}{\text{argmax}} |\text{Agr}(J, \text{Maj}(\boldsymbol{J}))|.$$

The *max-sum* rule, also known as the *prototype* rule (Miller and Osherson, 2009), the *median* rule (Nehring et al., 2014), and the *generalised Kemeny* rule (Dietrich, 2014), maximises the sum of agreements with the profile:

Definition 6.5 (max-sum)**.** *Given a judgment profile* $\boldsymbol{J}$ *and a feasibility constraint* Γ'*, the* max-sum *rule is defined as:*

$$\text{max-sum}(\boldsymbol{J}, \Gamma') = \underset{J \in \text{Mod}(\Gamma')}{\text{argmax}} \sum_{i \in N} |\text{Agr}(J, J_i)|$$

We refer to these rules $F_{\text{JA}} \in \{\text{max-sum}, \text{max-num}\}$ as *majoritarian* JA rules. Note that both of these two rules really constitutes an entire family of aggregation rules, one for each feasibility constraint Γ'. To refer to the aggregation rule from a family of aggregation rules F_{JA} for a given feasibility constraint Γ', we write $F_{\text{JA}}(\cdot, \Gamma')$. Observe that when $\text{Maj}(\boldsymbol{J}) \in \text{Mod}(\Gamma')$, then we have for any $F_{\text{JA}} \in \{\text{max-sum}, \text{max-num}\}$ that $F_{\text{JA}}(\cdot, \Gamma') = \{\text{Maj}(\boldsymbol{J})\}$, i.e., if the majority judgment $\text{Maj}(\boldsymbol{J})$ is consistent with the feasibility constraint then these majoritarian rules will return $\text{Maj}(\boldsymbol{J})$ as the unique judgment.

6.3 Simulating Multiwinner Voting Rules

We begin this section by detailing how we model preference aggregation within our JA model and we then define what we mean for a JA rule to simulate a multiwinner voting rule.

6.3.1 Preference Agenda and Constraints

When voting, every voter reports a preference, as either a strict ranking or an approval set. The outcome of an election can also be viewed as a preference: all of the candidates elected are (collectively) preferred to all those not elected. Next, we prepare the grounds for our embeddings of multiwinner voting rules into JA by defining a number of constraints that can be used to model relevant properties of such preferences.

Given a set of candidates C, we let $\Phi^C_{\succcurlyeq} = \{p_{x \succcurlyeq y} \mid x, y \in C\}$ be the *preference agenda* (Dietrich and List, 2007; Endriss, 2016). Now we can think of accepting the proposition $p_{x \succcurlyeq y}$ as expressing a (weak) preference of candidate x over candidate y. We furthermore write $p_{x \succ y}$ as a shorthand for $p_{x \succcurlyeq y} \wedge \neg p_{y \succcurlyeq x}$.

We now can express properties of binary relations as constraints in our logical language. Each constraint is defined for some set of candidates $W \subseteq C$. This includes, in particular, common properties of preference relations such as completeness, antisymmetry, and transitivity:

$$\textsc{complete}_W = \bigwedge_{x,y \in W} (p_{x \succcurlyeq y} \vee p_{y \succcurlyeq x})$$

$$\textsc{anti-sym}_W = \bigwedge_{x,y \in W \text{ s.t. } x \neq y} \neg(p_{x \succcurlyeq y} \wedge p_{y \succcurlyeq x})$$

$$\textsc{transitive}_W = \bigwedge_{x,y,z \in W} (p_{x \succcurlyeq y} \wedge p_{y \succcurlyeq z} \rightarrow p_{x \succcurlyeq z})$$

We can now formulate a constraint that is satisfied by a judgment on $\Phi^C_{\succcurlyeq}$ that corresponds to a strict ranking of all candidates in C:

$$\textsc{ranking} = \textsc{complete}_C \wedge \textsc{anti-sym}_C \wedge \textsc{transitive}_C$$

For two candidates, for which it is not the case that you strictly prefer one over the other, there are two possibilities: you either are *indifferent* between them or you consider them *incomparable*. The next two constraints express indifference and incomparability, respectively, between all the candidates in W:

$$\textsc{indiff}_W = \bigwedge_{x,y \in W} (p_{x \succcurlyeq y} \wedge p_{y \succcurlyeq x})$$

$$\textsc{incomp}_W = \bigwedge_{x,y \in W} \neg(p_{x \succcurlyeq y} \vee p_{y \succcurlyeq x})$$

Again for a given set $W \subseteq C$, the following constraint expresses a strict preference for all candidates in W over all those not in W:

$$\textsc{top}_W = \bigwedge_{x \in W} \bigwedge_{y \in C \setminus W} (p_{x \succ y})$$

This property is satisfied, for instance, by an voter's approval ballot in case the voter approves of exactly the candidates in W. But it also holds for the collective preference returned by a multiwinner voting rule in case A is the set of winning candidates.

We are going to require constraints to describe both that (i) there exists such a set W of most preferred candidates and that (ii), there exists such a set and that this set has size k. In both cases, we may assume either indifference between all the candidates within the same set, or all of these candidates being incomparable. We focus on five such constraints.

For the first two of them, the first part of the name indicates the structure of the top set W, while the second part indicates that of the bottom set $C \setminus W$.

$$\begin{aligned}
\textsc{indiff-indiff} &= \bigvee_{W \in \mathcal{P}_+(C)} \left(\textsc{top}_W \wedge \textsc{indiff}_W \wedge \textsc{indiff}_{C\setminus W}\right)\\
\textsc{indiff-incomp} &= \bigvee_{W \in \mathcal{P}_+(C)} \left(\textsc{top}_W \wedge \textsc{indiff}_W \wedge \textsc{incomp}_{C\setminus W}\right)\\
\textsc{incomp-incomp} &= \bigvee_{W \in \mathcal{P}_+(C)} \left(\textsc{top}_W \wedge \textsc{incomp}_W \wedge \textsc{incomp}_{C\setminus W}\right)
\end{aligned}$$

Note that we do not consider the constraint INCOMP-INDIFF. The reason we did not do so is that it lacks the same conceptual appeal as INDIFF-INCOMP and INCOMP-INCOMP (since we find it most natural that a voter cannot compare the alternatives not in her top-set) while for INDIFF-INDIFF, we include it (despite it lacking this conceptual appeal) for technical reasons as it yields mathematically interesting results. The remaining three constraints follow the same naming convention, while also being prefixed with a number k to indicate that the top set W has size k.

$$\begin{aligned}
k\text{-}\textsc{indiff-indiff} &= \bigvee_{W \in \mathcal{P}_k(C)} \left(\textsc{top}_W \wedge \textsc{indiff}_W \wedge \textsc{indiff}_{C\setminus W}\right)\\
k\text{-}\textsc{indiff-incomp} &= \bigvee_{W \in \mathcal{P}_k(C)} \left(\textsc{top}_W \wedge \textsc{indiff}_W \wedge \textsc{incomp}_{C\setminus W}\right)\\
k\text{-}\textsc{incomp-incomp} &= \bigvee_{W \in \mathcal{P}_k(C)} \left(\textsc{top}_W \wedge \textsc{incomp}_W \wedge \textsc{incomp}_{C\setminus W}\right)
\end{aligned}$$

Our omission of k-INCOMP-INDIFF is due to similar reasons as mentioned above for the INCOMP-INDIFF constraint.

6.3.2 Extracting Election Winners

To simulate a multiwinner voting rule, the voters' preferences are turned into judgments that satisfy a suitable rationality constraint Γ. In the case of ordinal preferences, this is RANKING. In the case of approval-based preferences, we are going to use INDIFF-INCOMP, i.e., we are going to assume that a voter i who approves of the set A_i does not share any views regarding the relative desirability of the candidates she does not approve of (rather than to declare indifference between them). We consider this the most natural interpretation of an approval ballot (and this interpretation will turn out to be technically convenient as well).

We can then apply a JA rule F_{JA} to the preferences thus encoded, obtaining a collective judgment. In case that collective judgment satisfies the constraint $\textsc{top}_W$ for some set $W \subseteq C$, we can declare the candidates in W to be the winners of the original election instance. With this in mind, we are now ready to present our central definition that relates multiwinner voting rules and JA rules. This is similar to the definition of single-winner voting rules by Endriss (2018) .

Definition 6.6 (Simulation). *Given a set of candidates C, a JA rule F_{JA} for the preference agenda $\Phi^C_{\succcurlyeq}$, and a multiwinner voting rule F for C, we let:*

(*i*) $\Gamma = \textsc{ranking}$ *in case F uses ordinal preferences, and*

(*ii*) $\Gamma = \textsc{indiff-incomp}$ *in case it uses approval-based preferences.*

Then we say that F_{JA} simulates F if, for every preference profile $(\succcurlyeq_1, \ldots, \succcurlyeq_n) \in \text{Mod}(\Gamma)^n$ and corresponding judgment profile $\boldsymbol{J} = (J_1, \ldots, J_n)$, we have that:

$$F(\succcurlyeq_1, \ldots, \succcurlyeq_n) = \bigcup_{J \in F_{\text{JA}}(\boldsymbol{J})} \{W \subseteq C \mid J \models \textsc{top}_W\}.$$[1]

Observe that F_{JA} can simulate F only if F_{JA} satisfies a feasibility constraint Γ' that ensures that all outcomes $J \in F_{\text{JA}}(\boldsymbol{J})$ satisfy $\textsc{top}_W$ for some $W \subseteq C$. Furthermore, for F_{JA} to simulate a rule that returns committees of size k, the constraint Γ' needs to ensure that these sets W indeed always have size k. So constraints such as k-INDIFF-INDIFF, k-INDIFF-INCOMP and k-INCOMP-INCOMP are natural candidates for Γ' in such scenarios.

6.3.3 Simulation Results

With all the relevant definitions in place, we now present our simulation results for specific multiwinner voting rules. We start with results for two ordinal-based rules that may be regarded as the multiwinner counterparts of corresponding results by Endriss (2018) for single-winner voting rules. The first concerns the k-Borda rule.

Theorem 6.1. *When restricted to* RANKING*-rational profiles, we have that the* max-sum$(\cdot, k\text{-INDIFF-INCOMP})$ *rule simulates k-Borda.*

Proof. For any given set $W \subseteq C$, let J^W be defined as the unique judgment such that, for any given $x, y \in C$, we have $J^W(p_{x \succcurlyeq y}) = 1$ if and only if $x \in W$. Observe that the judgments J^W for sets W with $|W| = k$, are precisely the judgments that satisfy k-INDIFF-INCOMP:

$$\text{Mod}(k\text{-INDIFF-INCOMP}) = \{\, J^W \mid W \in \mathcal{P}_k(C)\}.$$

[1]Note that in general, multiple committees returned by a multiwinner rule F may correspond to a single judgment of F_{JA}. However, this does not apply to the feasibility constraints that we deal with as it is not possible for a judgment J to exist such that $J \models \textsc{top}_W$ and $J \models \textsc{top}_{W'}$ for committees $W \neq W' \subseteq C$.

Now consider any profile $\boldsymbol{J} \in \text{Mod}(\textsc{ranking})^n$. The max-sum rule induced by k-INDIFF-INCOMP, acting on the profile $\boldsymbol{J}$, returns:

$$\underset{J \in \text{Mod}(k\text{-}\textsc{indiff-incomp})}{\text{argmax}} \sum_{i \in N} |\text{Agr}(J, J_i)|.$$

This is equivalent to the following:

$$\underset{J^W \text{ s.t. } W \in \mathcal{P}_k(C)}{\text{argmax}} \sum_{i \in N} |\{\varphi \in \Phi^C_{\succcurlyeq} \mid J^W(\varphi) = J_i(\varphi)\}|.$$

Now, to determine the outcome, we need to compute a score that is obtained by summing over the elements φ of the preference agenda $\Phi^X_{\succcurlyeq}$. For any given $W \subseteq C$, we can separate this agenda into the following five disjoint parts:

1. $\{p_{x \succcurlyeq y} \mid x = y\}$
2. $\{p_{x \succcurlyeq y} \mid x \neq y \text{ and } x, y \notin W\}$
3. $\{p_{x \succcurlyeq y} \mid x \neq y \text{ and } x, y \in W\}$
4. $\{p_{x \succcurlyeq y} \mid x \in W,\, y \notin W\}$
5. $\{p_{x \succcurlyeq y} \mid x \notin W,\, y \in W\}$

We need to count the propositions φ in $\Phi^X_{\succcurlyeq}$ on which the two judgments, namely J^W and $J_i \in \text{Mod}(\textsc{ranking})$ representing voter i's judgment, agree.

First, for propositions in $\{p_{x \succcurlyeq y} \mid x = y\}$, the judgments agree when $x \in W$. This number remains the same regardless of the selection of W so it is omitted from the final count.

Second, J^W rejects all propositions in $\{p_{x \succcurlyeq y} \mid x \neq y \text{ and } x, y \notin W\}$, while J_i, by virtue of denoting a voter's strict ranking of the candidates, accepts exactly half of them (independently of the specific preference reported by voter i). Since this number also does not depend on the choice of the set W, we omit it going forward.

The same is true for $\{p_{x \succcurlyeq y} \mid x \neq y \text{ and } x, y \in W\}$: J^W again accepts all propositions, while J_i accepts exactly half of them, namely those propositions $p_{x \succcurlyeq y}$ for which $x \succ_i y$. However, as will become clear shortly, in this case it will be convenient to explicitly include this number in our count of the agreements.

For the remaining two parts of the agenda, that is $\{p_{x \succcurlyeq y} \mid x \in W,\, y \notin W\}$ and $\{p_{x \succcurlyeq y} \mid x \notin W,\, y \in W\}$, the judgment J^W accepts all propositions in the former and rejects all those in the latter. If J_i accepts a proposition $p_{x \succcurlyeq y}$ in the former, it rejects $p_{y \succcurlyeq x}$ in the latter. So J_i is in agreement with J^W, in both parts of the agenda, for exactly those pairs (x, y) for which $x \succ_i y$. So these judgments agree on the same number for both parts of the agenda, meaning that we need to consider only one.

To summarise, omitting the terms we can ignore, we obtain that the max-sum rule induced by k-INDIFF-INCOMP maps any given strict preference profile $(\succ_1, \ldots, \succ_n)$ to the outcome:

$$\underset{J^W \text{ s.t. } W \in \mathcal{P}_k(C)}{\text{argmax}} \sum_{i \in N} |\{(x,y) \mid x \succ_i y \,\text{and}\, x \in W,\ y \notin W\}| \\ +|\{(x,y) \mid x \succ_i y \,\text{and}\, x, y \in W\}|.$$

The latter can be further simplified to:

$$\underset{J^W \text{ s.t. } W \in \mathcal{P}_k(C)}{\text{argmax}} \sum_{i \in N} |\{y \in C \mid x \succ_i y \,\text{and}\, x \in W\}|.$$

Hence, the elected k-sized committee/s clearly maximise/s the Borda scores of the winning candidates; so this is the k-Borda rule. □

As a feasibility constraint, we used k-INDIFF-INCOMP to capture the intuition of collective incomparability within the losing set. However, there is also another intuitive route, namely the one where we assume indifference between the non-winners. Indeed, when counting agreements, as voters provide strict relations between candidate pairs, one can freely choose between indifference and incomparability amongst candidates grouped together in the outcome. We demonstrate this with the following result using k-INDIFF-INDIFF.

Theorem 6.2. *When restricted to* RANKING*-rational profiles, we have that the* max-sum$(\cdot, k$-INDIFF-INDIFF$)$ *rule simulates k-Borda.*

Proof (sketch). We define the judgment J^W as accepting a proposition $p_{x \succcurlyeq y}$ for $x, y \in C$, if and only if one of the following three conditions are satisfied: (i) $x, y \in W$, (ii) $x \in W$ but $y \notin W$, or (iii) $x, y \notin W$. This ensures that Mod(k-INDIFF-INDIFF) $= \{ J^W \mid W \in \mathcal{P}_k(C) \}$. The proof now proceeds along the same lines as that of Theorem 6.1.

Regarding agreements, both judgments accept all of $\{p_{x \succcurlyeq y} \mid x = y\}$ and agree on half of $\{p_{x \succcurlyeq y} \mid x \neq y \,\text{and}\, x, y \notin W\}$ with J^W accepting all them. The other agreements are as in the proof of Theorem 6.1. The same holds for the final count. □

The same can be seen for the third size-k feasibility constraint that we have highlighted, namely k-INCOMP-INCOMP even though we now assume incomparability amongst those candidates in the top set W.

Theorem 6.3. *When restricted to* RANKING*-rational profiles, we have that the* max-sum$(\cdot, k$-INCOMP-INCOMP$)$ *rule simulates k-Borda.*

Proof (sketch). We define the judgment J^W as accepting a proposition $p_{x \succcurlyeq y}$ for $x, y \in C$, if and only if $x \in W$ but $y \notin W$. This ensures that $\text{Mod}(k\text{-INCOMP-INCOMP}) = \{ J^W \mid W \in \mathcal{P}_k(C) \}$. The remainder of the proof is analogous to the proofs of the other k-Borda simulations.

Both judgments agree on those propositions in $\{p_{x \succcurlyeq y} \mid x = y\}$ when it holds that $x \in W$, the judgments then agree on half of $\{p_{x \succcurlyeq y} \mid x \neq y \text{ and } x, y \notin W\}$ and exactly half of $\{p_{x \succcurlyeq y} \mid x \neq y \text{ and } x, y \in W\}$, with J^W rejecting all of the propositions in both sets. The other agreements, for the remaining two parts of the agenda, are as in the proof of Theorem 6.1. The same holds for the final count. □

So for max-sum, we have observed that using any of our proposed feasibility constraints yields a simulation of k-Borda.

Now, we return to k-INDIFF-INCOMP as we apply this feasibility constraint to max-num, instead of max-sum, as we transition to simulating the multiwinner, Copeland rules. We begin with simulations where we assume that the following holds: $\text{Maj}(\boldsymbol{J}) \in \text{Mod}(\text{COMPLETE}_C \wedge \text{ANTI-SYM}_C)$. Intuitively, this is the assumption that for every candidate pair $x, y \in C$, we have that exactly one of $p_{x \succcurlyeq y}$ and $p_{y \succcurlyeq x}$ is set to true by the majority judgment $\text{Maj}(\boldsymbol{J})$. This is always the case, for example, when n is odd.

Proposition 6.4. *When we restrict ourselves to* RANKING*-rational judgment profiles $\boldsymbol{J}$ where* $\text{Maj}(\boldsymbol{J}) \in \text{Mod}(\text{COMPLETE}_C \wedge \text{ANTI-SYM}_C)$*, we have that the* max-num$(\cdot, k$-INDIFF-INCOMP$)$ *rule simulates k-Copeland.*

Proof (sketch). Take the same judgment J^W and agenda decomposition from Theorem 6.1. We proceed to assess the agreements between J^W and the $\text{Maj}(\boldsymbol{J})$ judgment.

Notice that when checking for the agreements, given the assumption that $\text{Maj}(\boldsymbol{J}) \in \text{Mod}(\text{COMPLETE}_C \wedge \text{ANTI-SYM}_C)$, then $\text{Maj}(\boldsymbol{J})$ sets, for any pair of distinct candidates $x, y \in C$, exactly one of $p_{x \succcurlyeq y}$ and $p_{y \succcurlyeq x}$ to true. This is much like the considered judgment J_i from Theorem 6.1. Hence, it is clear that the rule is equivalent to:

$$\underset{J^W \text{ s.t. } W \in \mathcal{P}_k(C)}{\text{argmax}} |\{p_{x \succcurlyeq y} \mid \text{Maj}(\boldsymbol{J})(p_{x \succcurlyeq y}) = 1 \text{ and } x \in W\}|.$$

The returned k-sized committee/s maximise the pairwise majority wins of the committee members and thus, this is k-Copeland. □

As with the k-Borda rule, k-Copeland has alternative simulations using either k-INDIFF-INDIFF or k-INCOMP-INCOMP as feasibility constraints. The (omitted) proofs of these facts are analogous to that of Theorems 6.2 and 6.3.

Proposition 6.5. *When we restrict ourselves to* RANKING*-rational judgment profiles* $\boldsymbol{J}$ *where* $\mathrm{Maj}(\boldsymbol{J}) \in \mathrm{Mod}(\text{COMPLETE}_C \wedge \text{ANTI-SYM}_C)$*, we have that both the* max-num$(\cdot, k\text{-INDIFF-INDIFF})$ *and the* max-num$(\cdot, k\text{-INCOMP-INCOMP})$ *rule simulate k-Copeland.*

We now assess what can be obtained using max-num when we drop this assumption on the majority judgment $\mathrm{Maj}(\boldsymbol{J})$. We present the following results where k-INDIFF-INCOMP and k-INDIFF-INDIFF are taken to be the feasibility constraints. We now look at the k-Copeland^1 and k-Copeland^0 separately.

Theorem 6.6. *When restricted to* RANKING*-rational profiles, we have that the* max-num$(\cdot, k\text{-INDIFF-INCOMP})$ *rule simulates k-Copeland^1.*

Proof (sketch). Consider, once again, the same judgment J^W and the agenda decomposition from Theorem 6.1. Let us count agreements between J^W and $\mathrm{Maj}(\boldsymbol{J})$.

$$\underset{J^W \text{ s.t. } W \in \mathcal{P}_k(C)}{\mathrm{argmax}} |\{p_{x \succcurlyeq y} \mid \mathrm{Maj}(\boldsymbol{J})(p_{x \succcurlyeq y}) = 1 \text{ and } x \in W\}| \\ + |\{p_{x \succcurlyeq y} \mid \mathrm{Maj}(\boldsymbol{J})(p_{x \succcurlyeq y}) = 0,\ x \neq y \text{ and } x, y \notin W\}|.$$

Observe that the latter maximises the number of pairs of candidates in $C \setminus W$ that contain a strict majority loser in the pairwise majority contest between the two candidates. So in order to maximise this value, the rule must minimise the number of candidate pairs in $C \setminus W$ that are tied according to the majority judgment. And to minimise that value, the rule must maximise the number of candidates $x \in W$ such that $\mathrm{Maj}(\boldsymbol{J})(p_{x \succcurlyeq y}) = 1$ for some $y \in C$. So, in the end, the rule becomes:

$$\underset{J^W \text{ s.t. } W \in \mathcal{P}_k(C)}{\mathrm{argmax}} |\{p_{x \succcurlyeq y} \mid \mathrm{Maj}(\boldsymbol{J})(p_{x \succcurlyeq y}) = 1 \text{ and } x \in W\}|.$$

And as we dropped the assumption that $\mathrm{Maj}(\boldsymbol{J}) \in \mathrm{Mod}(\text{COMPLETE}_C \wedge \text{ANTI-SYM}_C)$, this is exactly k-Copeland^1 (but not k-Copeland^0). □

Moving on to k-INDIFF-INCOMP as the feasibility constraint, we find that max-num no longer simulates k-Copeland^1, but rather, simulates k-Copeland^0.

Theorem 6.7. *When restricted to* RANKING*-rational profiles, we have that the* max-num$(\cdot, k\text{-INDIFF-INDIFF})$ *rule simulates k-Copeland^0.*

Proof (sketch). Take the judgment J^W and the agenda decomposition that was considered in Proposition 6.2. In counting the agreements as usual, we find that the rule is the following:

$$\underset{J^W \text{ s.t. } W \in \mathcal{P}_k(C)}{\operatorname{argmax}} |\{p_{x \succcurlyeq y} \mid \text{Maj}(\boldsymbol{J})(p_{x \succcurlyeq y}) = 1 \text{ and } x \in W\}| + \\ |\{p_{x \succcurlyeq y} \mid \text{Maj}(\boldsymbol{J})(p_{x \succcurlyeq y}) = 1,\ x \neq y \text{ and } x, y \notin W\}|.$$

To maximise this count, the rule must minimise the number of candidate pairs (x, y) with $\text{Maj}(\boldsymbol{J})(p_{x \succ y}) = 1$ such that either (i) $x, y \in W$, (ii) $x, y \notin W$, or (iii) $y \in W, x \notin W$. This is equivalent to:

$$\underset{J^W \text{ s.t. } W \in \mathcal{P}_k(C)}{\operatorname{argmax}} |\{p_{x \succcurlyeq y} \mid \text{Maj}(\boldsymbol{J})(p_{x \succ y}) = 1 \text{ and } x \in W\}|.$$

The rule then returns the k candidates that have the highest Copeland scores with $\tau = 0$ and thus, the rule simulates k-Copeland^0. □

We have simulated k-Borda, k-Copeland^1 and k-Copeland^0 with varying feasibility constraints but we have yet to ask what occurs when we use k-INCOMP-INCOMP as a feasibility constraint, adopt max-num as the JA rule, *and* drop the assumption that $\text{Maj}(\boldsymbol{J}) \in \text{Mod}(\text{COMPLETE}_C \wedge \text{ANTI-SYM}_C)$. In answering this, we veer towards the simulation of rules that exhibit different qualities to the aforementioned rules.

In the multiwinner voting literature, k-Borda and the two k-Copeland variants have been proposed as suitable candidates to perform tasks such as shortlisting, as they satisfy the axiom of *committee monotonicity*. This property ensures that winning candidates in a k-sized committee will remain winners if the target committee size is increased (Elkind et al., 2017; Barberà and Coelho, 2008; Aziz et al., 2017b; Faliszewski et al., 2017). Barberà and Coelho (2008) showed this property to be incompatible with another well-studied property, namely the *Condorcet*-related notion of *weak Gehrlein stability* (Gehrlein, 1985; Barberà and Coelho, 2008; Ratcliff, 2003; Aziz et al., 2017b).[2] We recall its definition below.

Definition 6.7 (Weak Gehrlein stability)**.** *Take a set of candidates C, a target committee size k, and a set of voters N with each $i \in N$ providing a strict ranking $\succ_i$. A committee $W \in \mathcal{P}_k(C)$ is weakly Gehrlein-stable if for any $x \in W$ and $y \in C \setminus W$, it is the case that $|\{i \in N \mid x \succ_i y\}| \geqslant |\{i \in N \mid y \succ_i x\}|$.*

Since we work with weak Gehrlein stability, we can allow for evenly-sized N. And note that for some C, N, k, and strict preference profiles, a weakly Gehrlein-stable committee may not exist.

[2] Aziz et al. (2017b) showed this incompatibility does not occur for *strict Gehrlein stability* with k-Copeland^0 being both committee monotone and strictly Gehrlein-stable.

We now show that, when certain members of the class of *additive majority rules* (AMRs) are induced by k-INCOMP-INCOMP on RANKING-rational profiles, the resultant JA rule returns judgments that correspond to k-sized, weakly Gehrlein-stable committees given such committees exist for the given profiles. Now, we give a definition of these additive majority rules (Botan and Endriss, 2020; Nehring and Pivato, 2019).

Definition 6.8 (Additive Majority Rule, AMR). *A JA rule F_{JA} is an additive majority rule if there exists a non-decreasing gain function $g : [0, n] \to \mathbb{R}$ such that $g(t) < g(t')$ for $t < \frac{n}{2} \leqslant t'$, and for every feasibility constraint Γ' and JA profile $\boldsymbol{J}$, it holds that:*

$$F_{\text{JA}}(\boldsymbol{J}, \Gamma') = \operatorname*{argmax}_{J \in \text{Mod}(\Gamma')} \sum_{\varphi \in J} g(n_{(\boldsymbol{J},\varphi)}).$$

We obtain the max-sum rule for $g(t) = t$, while max-num has $g(t) = 1$ when $t \geqslant \frac{n}{2}$ and $g(t) = 0$ otherwise. Now, we get the following simulation result regarding weakly Gehrlein-stable rules:

Theorem 6.8. *When restricted to* RANKING*-rational profiles, every JA outcome $J \in F_{JA}(\cdot, k\text{-INCOMP-INCOMP})$ for an AMR F_{JA} that is based on a gain function g such that $g(t) = g(t')$ for any two $t, t' \geqslant n/2$, corresponds to a weakly Gehrlein-stable committee, provided such a weakly Gehrlein stable committee exists at all.*

Proof. Take an m-sized candidate set C and suppose that a weakly Gehrlein-stable committee $S \in \mathcal{P}_k(C)$ exists. Moreover, to derive a contradiction, suppose that there is some judgment $J^W \in F_{\text{JA}}(\boldsymbol{J}, k\text{-INCOMP-INCOMP})$ that corresponds to a committee $W \in \mathcal{P}_k(C)$ that is not weakly Gehrlein-stable.

Since W is not weakly Gehrlein-stable, there must be some $x \in W$ and some $y \in C \setminus W$ such that $|\{i \in N \mid x \succ_i y\}| < |\{i \in N \mid y \succ_i x\}|$. Then the score $\sum_{\varphi \in J^W} g(n_{(\boldsymbol{J},\varphi)})$ achieved by J^W is strictly less than $k(m-k)g_{\max}$, where $g_{\max} = g(\lceil n/2 \rceil) = \cdots = g(n)$.

However, the judgment J^S that corresponds to the committee S achieves the score $\sum_{\varphi \in J^S} g(n_{(\boldsymbol{J},\varphi)}) = k(m-k)g_{\max}$, and thus achieves a strictly higher score than J^W. This is a contradiction with our assumption that $J^W \in F_{\text{JA}}(\boldsymbol{J}, k\text{-INCOMP-INCOMP})$. Therefore we can conclude that all judgments in $F_{\text{JA}}(\boldsymbol{J}, k\text{-INCOMP-INCOMP})$ correspond to weakly Gehrlein-stable committees. □

This result makes a large selection from the AMR class available to those interested in committee stability with tools to easily define novel Gehrlein-stable rules. In fact, the subclass of AMRs to which Theorem 6.8 applies includes some rules that correspond to multiwinner voting rules that have been studied in the literature.

For example, consider the AMR based on the gain function g with $g(t) = 1$ when $t \geqslant \frac{n}{2}$ and $g(t) = 0$ otherwise—which coincides with the max-num rule. When restricted to RANKING-rational profiles, using k-INCOMP-INCOMP as the feasibility constraint, this rule simulates the following multiwinner voting rule known as *Number of External Defeats (NED)* (Barberà and Coelho, 2008; Aziz et al., 2017b; Faliszewski et al., 2017). Note that this is without the assumption on the structure of the majority judgment as seen in the k-Copeland simulations.

Definition 6.9 (Number of External Defeats, NED). *Recall that $\succcurlyeq_M$ is the weak majority relation. The NED rule elects the k-sized committee/s with the highest NED scores, defined for a committee W as:*

$$\mathit{NED}(W) = |\{(x, y) \in W \times C \setminus W \mid x \succcurlyeq_M y\}|.$$

Proposition 6.9. *When restricted to* RANKING*-rational profiles, we have that the* max-num$(\cdot, k$-INCOMP-INCOMP$)$ *rule simulates NED.*

Proof. We proceed as in the previously seen simulation proofs. Take the same judgment J^W and agenda decomposition from Theorem 6.3. As we look at the max-num rule, we proceed to assess the agreements between J^W and the $\mathrm{Maj}(\boldsymbol{J})$ judgment. Using similar reasoning as we have seen above, we find that the rule is:

$$\begin{aligned} \underset{J^W \text{ s.t. } W \in \mathcal{P}_k(C)}{\operatorname{argmax}} & |\{p_{x \succcurlyeq y} \mid \mathrm{Maj}(\boldsymbol{J})(p_{x \succcurlyeq y}) = 1 \,\text{and}\, x \in W, y \notin W\}| \\ & + |\{p_{x \succcurlyeq y} \mid \mathrm{Maj}(\boldsymbol{J})(p_{x \succcurlyeq y}) = 0,\; x \notin W, y \in W\}| \\ & + |\{p_{x \succcurlyeq y} \mid \mathrm{Maj}(\boldsymbol{J})(p_{x \succcurlyeq y}) = 0,\; x \neq y \,\text{and}\, x, y \in W\}| \\ & + |\{p_{x \succcurlyeq y} \mid \mathrm{Maj}(\boldsymbol{J})(p_{x \succcurlyeq y}) = 0,\; x \neq y \,\text{and}\, x, y \notin W\}|. \end{aligned}$$

Observe that these agreements are maximised when agreements between the majority judgment $\mathrm{Maj}(\boldsymbol{J})(p_{x \succcurlyeq y})$ and the set of propositions $\{p_{x \succcurlyeq y} \mid x \in W, y \notin W\}$ is maximised which leads us to the following:

$$\begin{aligned} & \underset{J^W \text{ s.t. } W \in \mathcal{P}_k(C)}{\operatorname{argmax}} |\{p_{x \succcurlyeq y} \mid \mathrm{Maj}(\boldsymbol{J})(p_{x \succcurlyeq y}) = 1 \,\text{and}\, x \in W, y \notin W\}| \\ = & \underset{J^W \text{ s.t. } W \in \mathcal{P}_k(C)}{\operatorname{argmax}} |(x, y)) \mid \mathrm{Maj}(\boldsymbol{J})(p_{x \succcurlyeq y}) = 1 \,\text{and}\, x \in W, y \notin W\}|. \end{aligned}$$

With this, we find that the rule is indeed equivalent to the NED rule. □

We can do similarly for the following Gehrlein-stable, multiwinner voting rule called *k-Kemeny* (Ratcliff, 2003; Barberà and Coelho, 2008).

Definition 6.10 (k-Kemeny). *The k-Kemeny rule elects the k-sized committee/s with the highest Kemeny scores, defined for a committee W as:*

$$K(W) = \sum_{x \in W, y \in C \setminus W} \max\{0, |\{i \in N \mid y \succ_i x\}| - |\{i \in N \mid x \succ_i y\}|\}.$$

We find that the k-Kemeny rule can be simulated using the AMR based on the following gain function g: $g(t) = 0$ when $t \geqslant \frac{n}{2}$ and $g(t) = 2t - n$ otherwise, or put differently, $g(t) = \max\{0, 2t - n\}$. Specifically, this simulation of k-Kemeny goes through when this particular AMR is restricted to RANKING-rational profiles, using k-INCOMP-INCOMP as constraint.

Proposition 6.10. *When restricted to* RANKING*-rational profiles, then for every AMR F_{JA} that is based on a gain function g with $g(t) = \max\{0, 2t - n\}$, the $F_{\text{JA}}(\cdot, k\text{-INCOMP-INCOMP})$ rule simulates k-Kemeny.*

Proof (sketch). For this proof sketch, we first note that with similar reasoning as used in Proposition 6.9, we find that the score of the AMR is maximised when the agreement for propositions in $\{p_{x \succcurlyeq y} \mid x \in W, y \notin W\}$ is maximised.

Then we note that this correspondence revolves around the fact that for each candidate pair (x, y), if candidate x is selected in the outcome W and candidate y is not, i.e., the AMR sets $p_{x \succcurlyeq y} = 1$ and $p_{y \succcurlyeq x} = 0$, a score of $\max\{0, |\{i \in N \mid y \succ_i x\}| - |\{i \in N \mid x \succ_i y\}|\}$ is added for $p_{x \succcurlyeq y}$ to the total score (as is done by k-Kemeny). □

	max-sum	max-num	AMR with $g(t) = \max\{0, 2t - n\}$
k-INDIFF-INDIFF	k-Borda	k-Copeland k-Copeland^0	-
k-INDIFF-INCOMP	k-Borda	k-Copeland k-Copeland^1	-
k-INCOMP-INCOMP	k-Borda	k-Copeland NED	k-Kemeny

Table 6.1: Simulations of fixed-size multiwinner voting rules on RANKING-rational profiles. Those simulations highlighted in blue are those that hold only when we have that $\text{Maj}(\boldsymbol{J}) \in \text{Mod}(\text{COMPLETE}_C \land \text{ANTI-SYM}_C)$.

We now transition to approval-based multiwinner voting rules, with a natural starting point being the simple AV rule. As previously mentioned, the rationality constraint for these rules will be INDIFF-INCOMP, which allows voters to have approval ballots of arbitrary size.

Proposition 6.11. *When restricted to* INDIFF-INCOMP*-rational profiles, we have that the* max-sum$(\cdot, k$-INDIFF-INCOMP$)$ *rule simulates AV.*

Proof. Take again, along with the usual agenda decomposition, the judgment J^W from Theorem 6.1 that accepts a proposition $p_{x \succcurlyeq y}$ for $x, y \in C$ if and only if $x \in W$. With INDIFF-INCOMP-rational profiles, each voter sets indifference between her most-preferred candidates. We fix an approval set A_i for voter i such that $J_i(p_{x \succcurlyeq y}) = 1$ if and only if $x \in A_i$. Note that this captures a voter specifying incompatibility amongst those candidates she does not approve of.

It is easy to verify through counting relevant agreements that the max-sum rule induced by k-INDIFF-INCOMP on INDIFF-INCOMP-rational profiles is:

$$\underset{J^W \text{ s.t. } W \in \mathcal{P}_k(C)}{\text{argmax}} \sum_{i \in N} |\{x \in W \mid J_i(p_{x \succcurlyeq y}) = 1\}|$$

Thus, in the end, we get that this rule is equivalent to:

$$\underset{J^W \text{ s.t. } W \in \mathcal{P}_k(C)}{\text{argmax}} \sum_{i \in N} |\{x \in W \cap A_i\}|$$

The elected committee/s of size k maximise/s the approvals of the committee members, which gives us a simulation of AV. □

We now extend the AV simulation to other Thiele rules with the use of the following refinement of the max-sum rule:

Definition 6.11 (u-max-sum)**.** *Given a judgment profile* $\boldsymbol{J}$*, a feasibility constraint* Γ' *and some non-increasing scoring function* $u : \mathbb{Z}_{\geqslant 0} \to \mathbb{R}_{\geqslant 0}$*, the* u-max-sum *rule is defined as:*

$$u\text{-max-sum}(\boldsymbol{J}, \Gamma') = \underset{J \in \text{Mod}(\Gamma')}{\text{argmax}} \sum_{i \in N} \sum_{t=1}^{|\text{Agr}(J, J_i)|} u(t).$$

That the u-max-sum rule facilitates the simulation of PAV and α-CC is immediate from its definition and our proof sketch for Proposition 6.11. However, we present the next results in some detail.

Proposition 6.12. *When restricted to* INDIFF-INCOMP*-rational profiles, we have that the* u-max-sum *rule induced by* k-INDIFF-INCOMP *simulates PAV when using the scoring function where* $u(t) = {}^1\!/_t$ *for all* t*.*

Proof. Once again, we consider the agenda decomposition, judgment J^W and fixing of voters' approval sets from Theorem 6.1. Now, we have that the general u-max-sum rule is defined as follows:

$$\underset{J^W \text{ s.t. } W \in \mathcal{P}_k(C)}{\operatorname{argmax}} \sum_{i \in N} \sum_{t=1}^{|\mathrm{Agr}(J,J_i)|} u(t)$$

Consider the specific scoring function u where $u(t) = 1/t$ for all t. By counting agreements, we see that the u-max-sum rule induced by k-INDIFF-INCOMP on INDIFF-INCOMP-rational profiles is equivalent to:

$$\underset{J^W \text{ s.t. } W \in \mathcal{P}_k(C)}{\operatorname{argmax}} \sum_{i \in N} \sum_{t=1}^{|\{(x,y) | x \succcurlyeq_i y \text{ and } x \in W, y \in C\}|} u(t)$$

Now, to see that this rule with $u(t) = 1/t$ for all t, simulates PAV, observe that for every voter i, voter i's approval of a candidate in committee W results in exactly $m-1$ agreements. So a voter i that gave a committee W a PAV score of $\sum_{t=1}^{|A_i \cap W|} u(t)$ will give said committee's corresponding judgment a score of $\sum_{t=1}^{|A_i \cap W| \cdot (m-1)} u(t)$ for this particular u-max-sum rule. And so, a judgment J^W representing a committee W that returned by PAV will still yield a higher score via the following rule than those judgments that correspond to non-winning committees of PAV:

$$\underset{J^W \text{ s.t. } W \in \mathcal{P}_k(C)}{\operatorname{argmax}} \sum_{i \in N} \sum_{t=1}^{|\{(x,y) | x \succcurlyeq_i y \text{ and } x \in W, y \in C\}|} u(t)$$

$$= \underset{J^W \text{ s.t. } W \in \mathcal{P}_k(C)}{\operatorname{argmax}} \sum_{i \in N} \sum_{t=1}^{|\{x \in W | J_i(p_{x \succcurlyeq y})=1\}|} u(t).$$

□

Proposition 6.13. *When restricted to* INDIFF-INCOMP*-rational profiles, the* u-max-sum *rule induced by* k-INDIFF-INCOMP *simulates* α*-CC when using the scoring function where* $u(1) = 1$ *and* $u(t) = 0$ *for all* $t > 1$.

Proof (sketch). By applying the same reasoning as in the proof of Proposition 6.12, but for the scoring function where $u(1) = 1$ and $u(t) = 0$ for all $t > 1$, we find that u-max-sum rule simulates the α-CC rule. □

6.4 Computational Considerations

Worst-case intractability has been shown for many JA rules. Specifically, it has been shown that computing outcomes under max-sum and max-num is Θ_2^{p}-hard (Endriss et al., 2020). Thus, when simulating multiwinner rules in the JA model, we encounter the paradox that ordinarily easy-to-compute rules, such as k-Borda and AV (Aziz et al., 2015; Elkind et al., 2017), now seem computationally difficult to implement. To address this mismatch, we build on the approach for identifying tractable fragments of JA developed by De Haan (2018). Of particular interest to us, De Haan (2018) showed that certain JA rules can be used efficiently when the integrity constraint (that is the feasibility constraint in the JA model we study) is represented as a circuit in *decomposable negation normal form*, or a DNNF circuit for short. Recently, this technique has been used to obtain tractable embeddings of participatory budgeting into JA (Rey et al., 2020, 2023). We begin with the definition of this type of circuit, as given by Darwiche and Marquis (2002).

Definition 6.12 (DNNF circuits)**.** *A Boolean circuit in negation normal form (NNF) is a directed acyclic graph with a single root where each internal node is labelled with $\vee$ or $\wedge$, and every leaf is labelled with $\top$, $\bot$, x or $\neg x$ for a propositional variable x. A* DNNF circuit *is an NNF circuit that satisfies decomposability: for each conjunction in the circuit, no two conjuncts share a propositional variable.*

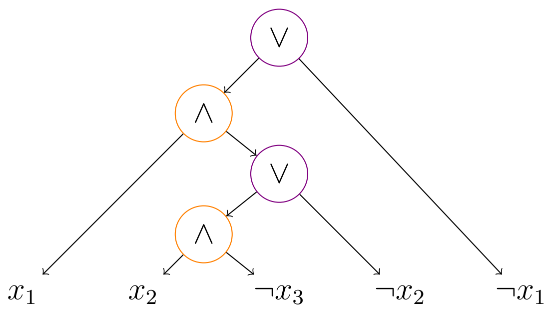

Figure 6.1: An example of a DNNF circuit.

The results of De Haan (2018) cover certain members in the class of scoring rules (Dietrich, 2014), including max-sum and max-num. Scoring rules select those constraint-satisfying JA outcomes that maximise the score of an associated scoring function. Such a function attaches a score to each issue with respect to an voter's judgment. Before restating the relevant result in Theorem 6.14, we define the *outcome determination problem* $\textsc{Outcome}(F_{\mathrm{JA}})$ for a given JA rule F_{JA}.

OUTCOME(F_{JA})	
Given:	A judgment profile $\boldsymbol{J}$ for an agenda Φ, an integrity constraint Γ', and a partial judgment d on Φ.
Question:	Is there a $J \in F_{\mathrm{JA}}(\boldsymbol{J}, \Gamma')$ that agrees with d?

This decision problem is formulated with the presence of a partial judgment d in order to mirror the problem used by De Haan (2018), who mention that this problem is natural when one desires collective outcomes with a particular structure or property, e.g., a judgment that agrees with the partial judgment d.

Theorem 6.14 (De Haan, 2018)**.** *When the integrity constraint Γ' is represented as a DNNF circuit, then* OUTCOME(F_{JA}) *is polynomial-time solvable for $F_{\mathrm{JA}} \in \{\text{max-sum}, \text{max-num}\}$.*

This result was extended to a general class of *additive rules*, which includes both the scoring rules and the AMRs (Rey et al., 2020, 2023).

We now show that the k-INDIFF-INCOMP constraint can be represented as a DNNF circuit. Recall that this constraint sets indifference amongst the top candidates and incomparability between those in the bottom set. Observe that for any candidate x in the top set, the proposition $p_{x \succcurlyeq y}$ is true for all $y \in C$. On the other hand, if x is in the bottom set, the proposition $p_{x \succcurlyeq y}$ is false for all candidates $y \in C$.

Proposition 6.15. *Given a finite set C of candidates and a corresponding preference agenda $\Phi^C_{\succcurlyeq}$, the k-INDIFF-INCOMP constraint can be encoded into a DNNF circuit in polynomial time.*

Proof. Given the set of candidates $C = \{x_1, \dots, x_m\}$, we construct the circuit according to the (arbitrary) ordering $x_1, \dots, x_m$ of C. An ordering of propositions such as $p_{x \succcurlyeq x}$ for each $x \in C$ is then set. We say x_i is the i-th candidate in C while $p_{x_i \succcurlyeq x_i}$ is the corresponding proposition in $\Phi^C_{\succcurlyeq}$. We also use the counting variables i and j during the circuit construction, both starting at 0.

The circuit contains nodes $N_{i,j}$, each of which denoting that we have assessed the propositions in the sequence up to (and including) index $i-1$ with j the current size of the winning set. We set $N_{0,0}$ to be the root of the circuit. If $i = |X| + 1$ and $j = k$, then $N_{i,j} = \top$. If $i = |X| + 1$ and $j \neq k$, then $N_{i,j} = \bot$. Now if $i < |X| + 1$, we either have (i) $p_{x_i \succcurlyeq x_i}$ is true or (ii) $p_{x_i \succcurlyeq x_i}$ is false. We set the node $N_{i,j}$ to be the disjunction $\alpha \vee \beta$, where we have both of the following:

- $\alpha = (N(i+1,\, j+1) \wedge \bigwedge_{y \in C} p_{x_i \succcurlyeq y})$.

- $\beta = (N(i+1, j) \wedge \bigwedge_{y \in C} \neg p_{x_i \succcurlyeq y})$.

We have that every leaf is either $\top$, $\bot$ or $p_{x \succcurlyeq y}$ for some x and y. Thus, we have a NNF circuit. Each proposition appears exactly once in the circuit so we also have that it is decomposable. The circuit is only satisfied by a preference agenda $\Phi^C_{\succcurlyeq}$ if the agenda has a k-sized top set of candidates with indifference between them while the bottom set's candidates are incomparable. So we have that the circuit corresponds to our constraint. This circuit can also be constructed in polynomial time as the process terminates once each candidate in C has been assessed exactly once. □

Note that the proof of the following result, regarding the circuit encoding for INDIFF-INCOMP, is omitted as it works almost identically to that in Proposition 6.15, except there is no tracking of the winning set's size.

Proposition 6.16. *Given a finite set C of candidates and a corresponding preference agenda $\Phi^C_{\succcurlyeq}$, the* INDIFF-INCOMP *constraint can be encoded into a DNNF circuit in polynomial time.*

These results ensure that for the max-sum and max-num rules, when using k-INDIFF-INCOMP as the feasibility constraint, such as in our simulations, computing the outcomes can still be done in polynomial time. We continue with this approach to analyse k-INDIFF-INDIFF.

Recall that k-INDIFF-INDIFF sets indifference within both the k-sized top set and the bottom set. We now show that this constraint cannot be constructed as a DNNF circuit in polynomial time. The claim is that, given a set of candidates C, computing the max-sum rule induced by k-INDIFF-INDIFF with the rationality constraint $\top$, i.e., when the ballots are unconstrained, is a computationally difficult problem. This implies that we cannot construct a DNNF circuit representing k-INDIFF-INDIFF in polynomial time (assuming that $\mathsf{P} \neq \mathsf{NP}$). We show that this problem is NP-hard by giving a reduction from the following problem.

NAE-3SAT	
Given:	A formula φ in 3CNF.
Question:	Is there a truth assignment satisfying φ that falsifies at least one literal in each clause of φ?

Theorem 6.17. *Given a finite set C of candidates and a corresponding preference agenda $\Phi^C_{\succcurlyeq}$, computing any outcome of the* max-sum *rule induced by the* k-INDIFF-INDIFF *constraint on $\top$-restricted ballots is NP-hard.*

Proof. We reduce from NAE-3SAT. Let φ be an arbitrary instance of NAE-3SAT, where $x_1, \ldots, x_m$ are the variables in φ and $c_1, \ldots, c_u$ are the clauses in φ. For each variable x_i in φ, we create a candidate a_{x_i} for each literal of the variable x_i. This produces $2m$ candidates. The profile on $\Phi^C_{\succcurlyeq}$, is as follows: for each variable x_i, we add $10u$ voters, each of which has a judgment set that sets every preference issue to true except for $p_{a_{x_i} \succcurlyeq a_{\neg x_i}}$ and $p_{a_{\neg x_i} \succcurlyeq a_{x_i}}$. For each clause c_j, and each pair of literals (l_1, l_2) within c_j, we create a voter that has the judgment set that sets every issue in the agenda to true except for $p_{a_{l_1} \succcurlyeq a_{l_2}}$ and $p_{a_{l_2} \succcurlyeq a_{l_1}}$. We claim that there exists a truth assignment that both satisfies and falsifies at least one literal in each clause of φ if and only if the translated outcome, given by max-sum induced by m-INDIFF-INDIFF, has a score of at least $(10mu + 3u) \cdot (4\binom{m}{2} + m^2) - 10mu - 4u$.

($\Longrightarrow$) Assume we have a truth assignment that satisfies and falsifies at least one literal in each clause of φ. Those variables that are set to true in this assignment correspond to those candidates in the top set of the outcome. A truth assignment will set one of variables x_i and $\neg x_i$ to true and correspondingly, we construct the top set to contain one of a_{x_i} and $a_{\neg x_i}$. This gives us m candidates, as is required by the feasibility constraint m-INDIFF-INDIFF. Now to check if this assignment gives the correct score. We have $10u$ voters for each top-set member and three voters for each clause who agree on the following: $2\binom{m}{2}$ propositions in each of the top and bottom sets, which represents $4\binom{m}{2}$ in the score, as well as, for each member of the top set, they agree on a proposition for every candidate in the bottom set they are preferred to. This gives m^2 propositions. This constitutes the $(10mu + 3u) \cdot (4\binom{m}{2} + m^2)$ agreements. Of propositions such as $p_{a_{x_i} \succcurlyeq a_{\neg x_i}}$ and $p_{a_{\neg x_i} \succcurlyeq a_{x_i}}$, exactly one is true. This accounts for $10mu$ disagreements. In each clause, four of the six propositions associated with clause voters are set to true. This accounts for the $4u$ disagreements. The difference between agreements and disagreements gives the score $(10mu + 3u) \cdot (4\binom{m}{2} + m^2) - 10mu - 4u$.

($\Longleftarrow$) Assume we have a max-sum outcome with at least a score of $(10mu + 3u) \cdot (4\binom{m}{2} + m^2) - 10mu - 4u$ which we refer to as $s^\star$ for the remainder of this proof. There are m candidates in the top set due to the feasibility constraint. We must check that they give the correct assignment of φ. The m candidates cannot contain conflicting pairs, such as a_{x_i} and $\neg a_{x_i}$, as this leads to a score of at most $(10mu + 3u) \cdot (4\binom{m}{2} + m^2) - 20mu - 4u$ which is less than score $s^\star$. This gives a truth assignment for φ. Now to check if the clauses are satisfied by this assignment. From the score, we know that all literals in a clause cannot have same truth value, as that would give a score of at most $(10mu + 3u) \cdot (4\binom{m}{2} + m^2) - 10mu - 6u$. This is also smaller than the score $s^\star$ so we have a valid assignment for φ which also satisfies all clauses of φ in the correct manner. □

So in general, it is not possible to efficiently use the max-sum rule induced by k-INDIFF-INDIFF (and also, INDIFF-INDIFF). However, when used on RANKING-restricted profiles, it simulates k-Borda (Theorem 6.2). So in this particular case, we obtain computational efficiency. Moreover, we know that we can simulate NED using the max-num rule that is induced by the feasibility constraint k-INCOMP-INCOMP, and since computing outcomes for the NED rule is NP-hard (Aziz et al., 2017b), it follows that k-INCOMP-INCOMP cannot be encoded into a DNNF circuit in polynomial time (assuming that $\mathsf{P} \neq \mathsf{NP}$). However, in a similar fashion as we have done with k-INDIFF-INDIFF, we have shown in Theorem 6.3 that a polynomial-time computable multiwinner voting rule, namely k-Borda, can be simulated using a max-sum rule induced by k-INCOMP-INCOMP on RANKING-rational profiles. These observations lead to the following:

Theorem 6.18. *Let $\mathcal{F}$ be a class of feasibility constraints. Suppose computing outcomes for a rule $F_{JA} \in \{\text{max-sum}, \text{max-num}\}$ induced by a feasibility constraint $\Gamma' \in \mathcal{F}$ on $\top$-restricted ballots is* NP*-hard. It then holds that no $\Gamma' \in \mathcal{F}$ can be encoded into a DNNF circuit in polynomial time, unless $\mathsf{P} = \mathsf{NP}$.*

These cases highlight the care that is required in selecting the constraints used in this JA model as certain constraint-rule combinations have these computational drawbacks.

6.5 Chapter Summary

In this chapter, we illustrated how the JA model with rationality and feasibility constraints enables us to simulate important multiwinner voting rules in judgment aggregation. We subsequently showed, by encoding the constraints as DNNF circuits, that some of these simulations retain the computational efficiency of their multiwinner counterparts. On the other hand, for specific feasibility constraints, namely k-INDIFF-INDIFF and k-INCOMP-INCOMP, this efficiency cannot be retained in general. Also, we demonstrated how the class of JA rules called AMRs can help produce multiwinner voting rules that satisfy the property of weak Gehrlein-stability.

Future Work. For future research, the JA simulation of more sophisticated multiwinner voting rules, such as sequential rules or other rules that are tailored towards proportional representation, should be explored. Also, in line with our result related to weak Gehrlein-stability, we mention that JA rules could aid the enrichment of multiwinner voting with new rules satisfying multiwinner notions other than committee stability.

Part Two

Constraining the Feasible Committees

Chapter 7

Constrained Public Decisions

This chapter gets us back on the track of providing proportional representation within complex domains. And with this, we study the public decisions model (detailed in Chapter 3) as a base. Recall that this is the model where a group of voters is presented with a set of issues for which they are expected to make a binary choice: typically, deciding to either *accept* or *reject* each issue. It is of particular interest due to the real-world scenarios captured by it. Notable examples include the following (Lang and Xia, 2016):

1. *Multiple Referenda:* The scenarios where the public vote directly on the resolution of political issues.

2. *Group Activity Planning:* A group of individuals are to choose, as a collective, the activities that the entire group shall partake in.

3. *Committee Elections*: A set of candidates are in the running for multiple positions on a committee and a group of decision-makers must select the committee members.

Given the collective nature of the problem, one of the natural desiderata is that the outcome represents a *fair* compromise for the participating voters. And in keeping with the approach of the thesis, we focus on proportional representation among the numerous possible interpretations of fairness. And indeed, as shown in Chapter 3, when zooming in on the model of public decisions, the goal of producing collective outcomes that proportionally reflect the opinions of the voter population has been drawing increasing attention in recent years (Freeman at al., 2020; Masarík et al., 2023; Skowron and Górecki, 2022). However, a component that has so far not received much attention in this growing literature on proportionality is the presence of *constraints* that restrict the possible outcomes that can be returned (with the exception of (Masarík et al., 2023)).

In this chapter, we focus on answering the question of what one may do when outcomes that would satisfy classical proportionality axioms—and thus be considered fair outcomes—are no longer feasible due to the presence of constraints. When examining real-world examples of the public-decision model, there are many scenarios where enriching the model with constraints fits naturally:

1. *Multiple Referenda:* Take an instance of multiple referenda where a constraint sets the acceptance of a political issue to be conditioned on the acceptance (or rejection) of another issue.

2. *Committee Elections*: Consider diversity constraints being imposed on a committee election instance that require that the number of candidates with certain characteristics that may be selected must be a precise number (or must fall within a specific range of numbers).

3. *Group Activity Planning:* Suppose we have instances of group activity planning where the feasible combinations of chosen activities are constrained by factors such as the activities' costs, the distances between the activity locations, the activities only being available in certain timeslots, and so on.

In tackling our task, we build on existing notions of proportionality that have been posed for less rich models and tailor them to the challenges that comes with the existence of constraints. Our work contributes to the ongoing research on this topic in the following three respects:

- **Axioms.** We provide (what we argue to be) natural adaptations of existing proportionality axioms—that are based on varying public-decision interpretations of justified representation—for this setting with constraints. Note that the axioms that we adapt are discussed in Chapter 3. Also, we import to our setting the priceability notion from multiwinner voting which provides another promising route to introducing proportionality for public decisions under constraints. Our adaptations require the introduction of the notion of feasible group deviations within this setting.

- **Constraint Restrictions.** We identify a class of constraints that provides promising results when looking to ensure proportional representation in this constrained setting. Specifically, for each of our axioms, we show that although it is challenging to satisfy these properties in general constrained instances, when one hones in on a restricted—yet highly expressive—class of constraints, we can achieve strong proportionality guarantees that represent approximations of the desirable justified-representation axioms.

- **Decision Rules.** We define novel adaptations of recently studied decision rules to our public-decision setting with constraints, such as MES (for public decisions) and the local-search variant of PAV.

Additional Related Work. We begin by noting that our constrained public-decision model closely resembles that of *judgment aggregation* and it also naturally fits into the area of collective decisions in *combinatorial domains* (see the book chapters by Endriss (2016) and Lang and Xia (2016) for general introductions to these two topics, respectively).[1]

Most relevant to this chapter's work is the research conducted on proportionality in the context of public decisions without constraints that we discussed in Chapter 3. Conitzer et al. (2017) focused on individually proportional outcomes, thus, our work more closely aligns with those that look to ensure adequate representation to groups of voters (Freeman at al., 2020; Masarík et al., 2023; Skowron and Górecki, 2022). Note that with the general social-choice model of Masarík et al. (2023), one can model our setting of interest. But by focusing on this setting, we explore properties that are specifically made for it, which in turn allows us to define, and subsequently conduct an analysis of, constrained public-decision rules that are not touched upon by Masarík et al. (2023). Thus, this chapter's results complement their work by showing further possibilities, and also limitations, for proportionality within this constrained public-decision model.

Also, the works on proportionality in models of sequential decision-making (that are discussed in Chapter 3) are relevant to our own as they can be seen as generalisations of the public-decision model without constraints (Bulteau et al., 2021; Chandak et al., 2024; Lackner, 2020). Amongst these sequential decision-making papers, those of Bulteau et al. (2021) and Chandak et al. (2024) relate to our work the most as they also implement justified-representation notions.

In related fields, previous work studied proportionality in various models that differ from the constrained public-decision model but feature collective choices on interconnected propositions: the belief merging setting (Haret et al., 2020), interdependent binary issues via conditional ballots (Brill at al., 2023c), and approval-based shortlisting with constraints (presented in a model of judgment aggregation) (Chingoma et al., 2022). The former two were mentioned in Chapter 3 while the latter is explored in the later Chapter 8.

Chapter Outline. We begin by detailing the constrained public-decision model in Section 7.1. We continue with Section 7.2 where we discuss two known ways in which justified representation is formalised for public decisions, and also present our notion of voter groups having feasible deviations. Then each of Sections 7.3 and 7.4 deal with a particular public-decision interpretation of justified representation. Finally, we deal with our constrained version of the priceability axiom in Section 7.5 before concluding in Section 7.6.

[1]Note that the judgment aggregation framework is discussed in more detail in Chapter 6.

7.1 The Model

A finite set of n voters $N = \{1, \ldots, n\}$ has to take a collective decision on a finite set of m issues $\mathcal{I} = \{a_1, \ldots, a_m\}$. It is typical in the public decisions setting to consider there only being two available decisions per issue but we instead adopt the following, more general setup. Each issue $a_t \in \mathcal{I}$ is associated with a finite set of *alternatives* called a *domain* $D_t = \{d_t^1, d_t^2, \ldots\}$ where $|D_t| \geqslant 2$ holds for all $t \in [m]$. The design decision of going beyond binary issues is motivated by the wider real-life applicability of this model when more than two alternatives are possible for each issue.

Each voter $i \in N$ submits a *ballot* $\boldsymbol{b}_i = (\boldsymbol{b}_i^1, \ldots, \boldsymbol{b}_i^m) \in D_1 \times \ldots \times D_m$ where $\boldsymbol{b}_i^t = d_t^c$ indicates that voter i chooses the decision d_t^c for the issue issue a_t. A *profile* $\boldsymbol{B} = (\boldsymbol{b}_1, \ldots, \boldsymbol{b}_n) \in (D_1 \times \ldots \times D_m)^n$ is a vector of the n voters' ballots. An *outcome* $\boldsymbol{w} = (w_1, \ldots, w_m) \in D_1 \times \ldots \times D_m$ is then a vector providing a decision for every issue at stake.

We focus on situations where some *constraints* limit the set of possible collective outcomes: we denote by $\mathcal{C} \subseteq D_1 \times \ldots \times D_m$ the set of *feasible* outcomes. We write $(\boldsymbol{B}, \mathcal{C})$ to denote an *election instance*. By a slight abuse of notation we also refer to $\mathcal{C}$ as the *constraint*, and thus, we refer to elections instances where $\mathcal{C} = D_1 \times \ldots \times D_m$ as *unconstrained* election instances.[2] Note that some of the results hinge on the way in which we chose to represent the constraint $\mathcal{C}$ and may be affected if a different constraint representation were to be used.

Note that voter ballots need not be consistent with the constraints, i.e., for an election instance $(\boldsymbol{B}, \mathcal{C})$, we do not require that $\boldsymbol{b}_i \in \mathcal{C}$ for all voters $i \in N$.[3]

Remark 7.1. *While not common in work done in the related judgment aggregation model, our assumption that voters ballots need not correspond to feasible outcomes is common in other settings of social choice. In multiwinner voting, voters can approve more candidates than the committee target size while in participatory budgeting, the sum of the costs of a voter's approved projects may exceed the instance's budget. For our setting, we argue that this approach helps capture real-world, constrained decision-making scenarios where either the constraint is* uncertain *when voters submit their ballots, or possibly, the voting process becomes more burdensome for voters as they attempt to create ballots with respect to a (possibly difficult to understand) constraint. For example, consider a group of friends deciding on the travel destinations of their shared holiday across the*

[2]Note that while we work formally with the constraint being an enumeration of all feasible outcomes, in practice, it is often possible to represent the set of feasible outcomes in more concise forms—via the use of formulas of *propositional logic*, for example—to help with parsing said constraint and/or speed up computation by exploiting the constraint's representation structure.

[3]This assumption takes our model closer to the particular model of judgment aggregation where the constraints on the output may differ from the constraints imposed on the the voters' input judgments (Endriss, 2018; Chingoma et al., 2022). The earlier Chapter 6 gives a more detailed outline of this model.

world, visiting one country in each continent. On a booking platform, there are a certain number of locations that can be selected per continent such as: Amsterdam, Paris and Vienna in Europe; Mexico City and Toronto in North America; Cairo, Nairobi and Cape Town in Africa; and so on. Each friend has a preferred combination of cities and their collective itinerary is subject to factors such as their travel budget or the available flight connections between cities. However, as flight costs and connections may change significantly on a day-to-day basis, it may be unclear which combination of cities are affordable. Therefore, it is not reasonable to impose the requirement, by default that is, that voter ballots are constraint-consistent. ◁

If needed, we explicitly state when we pivot from this assumption and require that voter ballots be constraint-consistent. At times, we shall restrict ourselves to election instances where $D_t = \{0, 1\}$ holds for every issue a_t. We refer to such cases as *binary* election instances. When necessary, we explicitly state whether any result hinges on the restriction to binary instances. Given an outcome $\boldsymbol{w}$ for a binary instance, the vector $\bar{\boldsymbol{w}} = (\bar{w}_1, \ldots, \bar{w}_m)$ is such that $\bar{w}_t = 1 - w_t$ for all issues $a_t \in \mathcal{I}$.

Now, consider an outcome $\boldsymbol{w}$, a set of issues $S \subseteq \mathcal{I}$ and some vector $\boldsymbol{v} = (v_1, \ldots, v_m) \in D_1 \times \ldots \times D_m$ (that can be interpreted as either an outcome or voter's ballot). We write $\boldsymbol{w}[S \leftarrow \boldsymbol{v}] = (w'_1, \ldots, w'_m)$ where $w'_t = w_t$ for all issues $a_t \in \mathcal{I} \setminus S$ and $w'_t = v_t$ for all issues $a_j \in S$. In other words, $\boldsymbol{w}[S \leftarrow \boldsymbol{v}]$ is the resultant vector of updating outcome $\boldsymbol{w}$'s decisions on the issues in S by fixing them to those of vector $\boldsymbol{v}$. For a given issue $a_t \in \mathcal{I}$ and a decision $d \in D_t$, we use $N(a_t, d) = \{i \in N \mid \boldsymbol{b}_i^t = d\}$ to denote the set of voters that agree with decision d on issue a_t. Given two vectors $\boldsymbol{v}, \boldsymbol{v}' \in D_1 \times \ldots \times D_m$, we denote the *agreement* between them by $\text{Agr}(\boldsymbol{v}, \boldsymbol{v}') = \{a_t \in \mathcal{I} \mid v_t = v'_t\}$. Then, the *satisfaction* that a voter i obtains from an outcome $\boldsymbol{w}$ corresponds to $u_i(\boldsymbol{w}) = |\text{Agr}(\boldsymbol{b}_i, \boldsymbol{w})|$, i.e., the number of decisions on which the voter i is in agreement with outcome $\boldsymbol{w}$.

Example 7.1. Suppose that a group of three voters must vote on the design of a logo and there are three relevant aspects that go into its design, each with their own set of available options: the outer shape of the logo can either be a circle or pentagon; the pattern within the logo's shape can either be horizontal lines, a checkerboard or a crosshatch; and the colour of this inner pattern can either be blue, red, yellow or green.

Formally, we can model this with three issues $\mathcal{I} = \{\text{Shape}, \text{Pattern}, \text{Colour}\}$, and here are the domains of each issue:

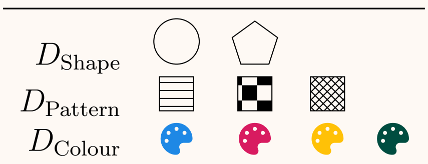

Next, suppose that the five voters cast the following ballots (where each voter's desired logo may, or may not, correspond to a feasible logo as determined by the constraint $\mathcal{C}$).

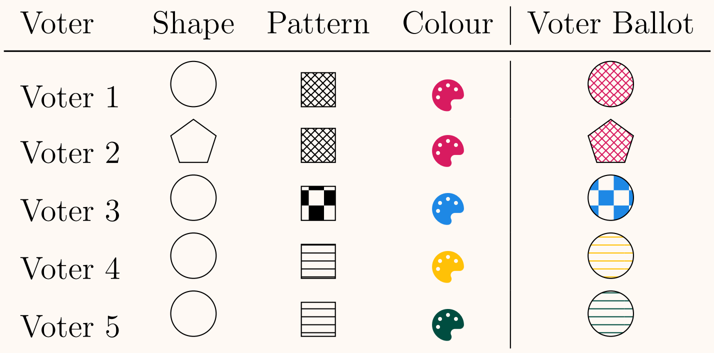

Now, we present an example of a constraint $\mathcal{C}$ that indicates the feasible logo designs, i.e., the feasible combinations of decisions on the three logo aspects.

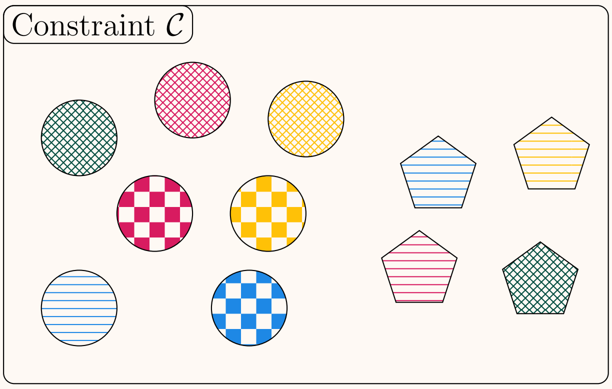

Note that with this constraint, we can see that only voters 1 and 3 submit feasible logos as their ballots.

7.2 Proportionality via Justified Representation

This section starts with the observation that classical notions of proportionality fall short when considering interconnected decisions (in the upcoming Example 7.2), and then follows with our proposed generalisations of such axioms that deal with constraints.

Ideally, when looking to make a proportional collective choice, we would like to meet the following criteria: a group of similarly-minded voters that is an α fraction of the population should have their opinions reflected in an α fraction of the m

issues. We wish to define an axiom for our model that captures this idea within our richer framework. In the setting of multiwinner voting, this is formally captured with the justified representation axioms such as EJR (detailed in Chapter 2). In the setting of public decisions, there are two different adaptations that have been studied and both were discussed in Chapter 3. These two notions are agreement-EJR and cohesiveness-EJR. In this chapter, we shall look at both. Recall that with agreement-EJR, the aim is to guarantee that 'a group of voters that agree on a set of issues T and represent an α fraction of the voter population, should control an $\alpha \cdot |T|$ number of the total issues in $\mathcal{I}$'. This approach differs from cohesiveness-EJR which states that 'a group of voters that agree on an α fraction of the issues and represent an α fraction of the total voters, should control $\alpha \cdot m$ of the issues in $\mathcal{I}$' (Chandak et al., 2024; Freeman at al., 2020).

Meeting the ideal outlined by both of these notions is not easy in our setting as the constraint $\mathcal{C}$ could rule out a seemingly fair outcome from the onset.

Example 7.2. Take the following binary election instance version of the logo design task from Example 7.1. Here, we have two issues $\mathcal{I} = \{\text{Patterned Shape}, \text{Colour}\}$ with each issue having a domain of size two.

$D_{\text{Patterned Shape}}$	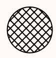	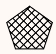
D_{Colour}		

Now, suppose that we have the following constraint with two feasible outcomes.

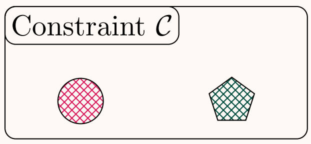

Then suppose that the two voters submitting the following ballots.

Voter	Patterned Shape	Colour	Voter Ballot
Voter 1	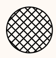		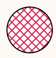
Voter 2	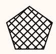		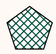

Note that voters 1 and 2 are both, on their own, cohesive groups. Now, both aforementioned EJR interpretations require each voter to obtain at least 1 in satisfaction, i.e., deciding half of the two issues. However, there exists no feasible outcome that provides agreement-EJR or cohesiveness-EJR as one voter $i \in \{1, 2\}$ will have satisfaction $u_i(\boldsymbol{w}) = 0$ for any outcome $\boldsymbol{w} \in \mathcal{C}$.

Example 7.2 makes clear an issue that we must take into account when defining proportionality properties when there are constraints. That is, a voter group that is an α fraction of the population may lay claim to deciding an α fraction of the issues, but in doing so, they may be resolving, or influencing the decision on, a larger portion of the issues than they are entitled to (due to the constraint).

In doing so, we look for a meaningful way to identify those voter groups that can justifiably complain at the selection of some outcome $\boldsymbol{w}$. This is formalised by the following definition where we identify the voter groups whose displeasure is justified as those groups that can propose an alternative outcome $\boldsymbol{w}^*$ that is feasible and yields greater satisfaction for each group member.

Definition 7.1 (($S, \boldsymbol{w}$)-deviation)**.** *Given election instance* $(\boldsymbol{B}, \mathcal{C})$ *and outcome* $\boldsymbol{w} \in \mathcal{C}$*, a set of voters* $N' \subseteq N$ *has an* $(S, \boldsymbol{w})$-deviation *if* $\emptyset \neq S \subseteq \mathcal{I}$ *is a set of issues such that all of the following hold:*

- $S \subseteq \mathrm{Agr}(\boldsymbol{b}_i, \boldsymbol{b}_j)$ *for all* $i, j \in N'$ *(the voters agree on the decisions on all issues in* S*).*
- $S \subseteq \mathcal{I} \setminus \mathrm{Agr}(\boldsymbol{b}_i, \boldsymbol{w})$ *for all* $i \in N'$ *(the voters disagree with outcome* $\boldsymbol{w}$*'s decisions on all issues in* S*).*
- $\boldsymbol{w}[S \leftarrow \boldsymbol{b}_i] \in \mathcal{C}$ *for all* $i \in N'$ *(fixing outcome* $\boldsymbol{w}$*'s decisions on issues in* S*, so as to agree with the voters in* N'*, induces a feasible outcome).*

Intuitively, given an outcome $\boldsymbol{w}$, a voter group having an $(S, \boldsymbol{w})$-deviation indicates the presence of another feasible outcome $\boldsymbol{w}^* \neq \boldsymbol{w}$ where every group member would be better off. Thus, our goal in providing a fair outcome reduces to finding an outcome where every group of voters that has an $(S, \boldsymbol{w})$-deviation is sufficiently represented. We shall use this $(S, \boldsymbol{w})$-deviation notion to convert proportionality axioms from unconstrained settings to axioms that deal with constraints. But first, we look at the following computational question associated with $(S, \boldsymbol{w})$-deviations: given an election instance $(\boldsymbol{B}, \mathcal{C})$ and an outcome $\boldsymbol{w} \in \mathcal{C}$, the problem is to find all groups of voters with an $(S, \boldsymbol{w})$-deviation.

Proposition 7.1. *Given an election instance* $(\boldsymbol{B}, \mathcal{C})$ *and an outcome* $\boldsymbol{w} \in \mathcal{C}$*, there exists an algorithm that finds all groups of voters* N' *such that there exists an* $S \subseteq \mathcal{I}$ *with* N' *having an* $(S, \boldsymbol{w})$*-deviation, that runs in* $O(|\mathcal{C}|^2 mn)$ *time.*

> Proof. Take $(\boldsymbol{B}, \mathcal{C})$ and outcome $\boldsymbol{w} \in \mathcal{C}$. Consider the following algorithm that operates in $|\mathcal{C}|$ rounds, assessing an outcome $\boldsymbol{w} \in \mathcal{C}$ in each round (with each outcome assessed once throughout): at each round for an outcome $\boldsymbol{w} \in \mathcal{C}$, iterate through all other outcomes $\boldsymbol{w}^* \neq \boldsymbol{w} \in \mathcal{C}$; fix S to be the issues that $\boldsymbol{w}$ and $\boldsymbol{w}^*$ disagree on; in at most mn steps, it can be checked if there is a

set of voters that agree with $\boldsymbol{w}^*$ on all issues in S which verifies the existence of a voter group N' with an $(S, \boldsymbol{w})$-deviation; keep track of all such groups N'; if all outcomes have been assessed, terminate, otherwise, move to the next outcome. This algorithm takes $O(|\mathcal{C}|^2 mn)$ time to complete in the worst case, which is polynomial in the input size given our assumptions. □

We offer the following remark in regards to the nature of Proposition 7.1.

Remark 7.2. *Proposition 7.1 can be seen as positive whenever the constraint $\mathcal{C}$ under consideration is 'not too large'. Such an assumption is reasonable for many real-life examples. Consider the quite general, collective task of selecting the features of some product. Our running example of the logo design is an instance of this. Other applicable scenarios include choosing the technical features of a shared computer or the items to be placed in an organisation's common area. In many cases, factors such as a limited budget (or limited space in the case of the common area) may result in very few feature combinations being feasible for said product. These are natural scenarios where we may encounter a 'small' constraint (according to our definition) with respect to the number of issues at hand and the size of their domains.* ◁

Our goal is to answer the following question: how much representation can we guarantee from some outcome $\boldsymbol{w}$, to a group of voters that has an $(S, \boldsymbol{w})$-deviation and that qualifies as underrepresented?

7.3 Justified Representation with Cohesiveness

We now propose the following adaptations of cohesiveness-EJR to public decisions with constraints. To adapt cohesiveness-EJR, we adapt ℓ-cohesiveness from multiwinner voting in a similar manner as done by Freeman at al. (2020). First, we define the notion of *T-agreeing* groups.

Definition 7.2 (T-agreeing)**.** *For a set of issues $T \subseteq \mathcal{I}$, we say that a group of voters is* T-agreeing *if $T \subseteq \mathrm{Agr}(\boldsymbol{b}_i, \boldsymbol{b}_j)$ holds for all voters $i, j \in N'$.*

We can then define our cohesiveness notion as the following:

Definition 7.3 (T-cohesiveness)**.** *For a set of issues $T \subseteq \mathcal{I}$, we say that a set of voters $N' \subseteq N$ is* T-cohesive *if N' is T-agreeing and it holds that $|N'| \geqslant |T| \cdot {}^{n}/_{m}$.*

Using T-cohesiveness, we can adapt the formulation of EJR from (Freeman at al., 2020) to our public-decision model with constraints.

Definition 7.4 (cohEJR$_{\mathcal{C}}$)**.** *Given an election instance* $(\boldsymbol{B}, \mathcal{C})$*, an outcome* $\boldsymbol{w}$ *provides* cohEJR$_{\mathcal{C}}$ *if for every* T*-cohesive group of voters* $N' \subseteq N$ *for some* $T \subseteq \mathcal{I}$ *with an* $(S, \boldsymbol{w})$*-deviation for some* $S \subseteq T$*, there exists a voter* $i \in N'$ *such that* $u_i(\boldsymbol{w}) \geqslant |T|$*.*

Intuitively, cohEJR$_{\mathcal{C}}$ deems an outcome to be unfair if there exists a T-cohesive voter group with (i) none of its group members having at least $|T|$ in satisfaction, and (ii) them fixing outcome $\boldsymbol{w}$'s decisions to match their own on some of the issues in T leads to some other feasible outcome.

We have the following result that can be interpreted as positive when the size of $\mathcal{C}$ is 'not too large'.

Proposition 7.2. *Given an election instance* $(\boldsymbol{B}, \mathcal{C})$ *and an outcome* $\boldsymbol{w} \in \mathcal{C}$*, there exists an algorithm that decides in* $O((\max_{t\in[m]} |D_t|)^m |\mathcal{C}|^3 mn)$ *time whether outcome* $\boldsymbol{w}$ *provides cohEJR*$_{\mathcal{C}}$*.*

> Proof. From Proposition 7.1 we know that, given an outcome $\boldsymbol{w}$, we can find all groups with some $(S, \boldsymbol{w})$-deviation for some $S \subseteq \mathcal{I}$ in $O(|\mathcal{C}|^2 mn)$ time. There can be at most $(\max_{t\in[m]} |D_t|)^m(|\mathcal{C}| - 1)$ such groups (recall that $\max_{t\in[m]} |D_t|$ is the maximal size of any issue's domain). Then, for each group N' with an $(S, \boldsymbol{w})$-deviation, we can check their size in polynomial time and thus verify whether they are a T-cohesive with $S \subseteq T$, and if so, we can check if there exists any voter $i \in N'$ with $u_i(\boldsymbol{w}) > |T|$. □

Now, Chandak et al. (2024) have already shown that, in general, cohesiveness-EJR is not always satisfiable in their sequential decisions model. This negative result carries over to the unconstrained public-decision setting. Although we shall later analyse the extent to which we can achieve positive results with cohesiveness-EJR in our constrained setting, this negative result motivates the study of the following weaker axiom that is an adaptation of the multiwinner JR axiom and can always be satisfied in the public-decision setting without constraints (Bulteau et al., 2021; Chandak et al., 2024; Freeman at al., 2020).

Definition 7.5 (cohJR$_{\mathcal{C}}$)**.** *Given an election instance* $(\boldsymbol{B}, \mathcal{C})$*, an outcome* $\boldsymbol{w}$ *provides* cohJR$_{\mathcal{C}}$ *if for every* T*-cohesive group of voters* $N' \subseteq N$ *for some* $T \subseteq \mathcal{I}$ *with an* $(S, \boldsymbol{w})$*-deviation for some* $S \subseteq T$ *where* $|S| = |T| = 1$*, there exists a voter* $i \in N'$ *such that* $u_i(\boldsymbol{w}) \geqslant 1$*.*

Unfortunately, when considering arbitrary constraints, even cohJR$_{\mathcal{C}}$ cannot always be achieved. Note that this even holds for binary election instances.

Proposition 7.3. *There exists an election instance where no outcome provides cohJR*$_{\mathcal{C}}$*.*

Proof. Consider the following binary election instance where we have two issues $\mathcal{I} = \{\text{Patterned Shape}, \text{Colour}\}$ with each issue having a domain with two alternatives.

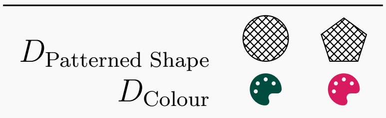

Then suppose we have the following constraint with two feasible outcomes where we see that the patterned shape must be a circle for the logo to be feasible.

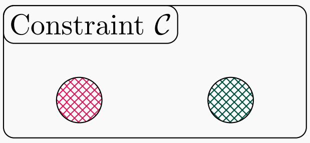

Now, suppose that there are two voters that submit the following ballots with both voters choosing the pentagon as the patterned shape and thus, neither of their ballots are feasible.

Voter	Patterned Shape	Colour	Voter Ballot
Voter 1	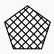		
Voter 2	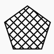		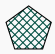

Note that for both outcomes $\boldsymbol{w} \in \mathcal{C}$, one voter will have satisfaction of 0 while being a T-cohesive group with an $(S, \boldsymbol{w})$-deviation for $|S| = |T| = 1$. As each voter is half of the population, they may 'flip' issue Colour to deviate towards the alternative feasible outcome, which provides them greater satisfaction than the current one. □

Let us now restrict the constraints that we consider. To do so, we introduce notation for the *fixed decisions* for a set of outcomes $C \subseteq \mathcal{C}$, which are the issues in $\mathcal{I}$ whose decisions are equivalent across all the outcomes in C. For a set of outcomes $C \subseteq \mathcal{C}$, we represent this as:

$$\mathcal{I}_{\text{fix}}(C) = \{a_t \in \mathcal{I} \mid \text{there exists some } d \in D_t \text{ such that } w_t = d \text{ for all } \boldsymbol{w} \in C\}.$$

Definition 7.6 (No Fixed Decisions (NFD) property)**.** *We say a constraint $\mathcal{C}$ has the* NFD *property if $\mathcal{I}_{fix}(\mathcal{C}) = \emptyset$ holds for $\mathcal{C}$.*

Remark 7.3. *At first glance, this NFD property seems more than a reasonable requirement but rather a property that should be assumed to hold by default. We* *argue however, that by doing so, we will neglect election instances where decisions that are fixed from the get-go may contribute to the satisfaction of voters and, specifically for our goal, these fixed decisions may aid in giving the voters their fair, proportional representation. It is this reason, why we did not restrict ourselves to election instances where the NFD property holds.* ◁

Now, we show that with the NFD property, the cohEJR$_\mathcal{C}$ axiom can always be satisfied, albeit only for 'small' election instances. We begin with cases where the number of feasible outcomes is limited to two.

Proposition 7.4. *For election instances* $(\boldsymbol{B}, \mathcal{C})$ *with* $|\mathcal{C}| = 2$ *where* $\mathcal{C}$ *has the NFD property, cohEJR*$_\mathcal{C}$ *can always be satisfied.*

> Proof. Take some feasible outcome $\boldsymbol{w} \in \mathcal{C}$. Observe that when $|\mathcal{C}| = 2$, if property NFD holds, then the two feasible outcomes differ on the decisions of all issues. Thus, it is only possible for T-cohesive groups with an $(S, \boldsymbol{w})$-deviation for $|S| \leqslant |T| = m$ to have an allowable deviation from $\boldsymbol{w}$ to the only other feasible outcome. This means only the entire voter population have the potential to deviate. And if such deviation to $\boldsymbol{w}'$ exists, then outcome $\boldsymbol{w}'$ sufficiently represents the entire voter population. □

Now we ask the following: can we guarantee cohEJR$_\mathcal{C}$ when $m \leqslant 3$? We answer in the positive when we restrict ourselves to binary election instances.

Proposition 7.5. *For binary election instances* $(\boldsymbol{B}, \mathcal{C})$ *with* $m \leqslant 3$ *where the constraint* $\mathcal{C}$ *has the NFD property, cohEJR*$_\mathcal{C}$ *can always be provided.*

> Proof. The case for $m = 1$ is trivially satisfied so we present the proof as two separate cases where the number of issues is either $m = 2$ or $m = 3$.
>
> **Case $\boldsymbol{m = 2}$:** Take a constraint $\mathcal{C}$ and a feasible outcome $\boldsymbol{w} = (d_x, d_y) \in \mathcal{C}$ where $d_x, d_y \in \{0, 1\}$. Let us now consider voter groups with an $(S, \boldsymbol{w})$-deviation over some set of issues $S \subseteq T$ who are witness to a violation of cohEJR$_\mathcal{C}$. As $m = 2$, the agreement among voters and the deviation may concern at most two issues, i.e., $|S|, |T| \in \{1, 2\}$.
>
> First, consider $|T| = 1$. Since $|S| \leqslant |T|$ and $S \neq \emptyset$, we have $|S| = 1$ for any T-cohesive group (which is thus of size $|N'| \geqslant n/2$) wishing to perform a $(S, \boldsymbol{w})$-deviation from $\boldsymbol{w}$ to some other feasible outcome $\boldsymbol{w}' \in \mathcal{C}$. If there is a voter $i \in N'$ such that $u_i(\boldsymbol{w}) \geqslant 1$, group N' would be sufficiently satisfied—therefore, cohEJR$_\mathcal{C}$ is satisfied. Otherwise, we have that all $i \in N'$ are unanimous and that $u_i(\boldsymbol{w}) = 0$; hence, $\boldsymbol{b}_i = (1 - d_x, 1 - d_y)$ for all

$i \in N'$. There are two possible outcomes that differ from $\boldsymbol{w}$ in only one issue-decision. If neither outcome is in $\mathcal{C}$, then no $(S, \boldsymbol{w})$-deviation is possible for N' and cohEJR$_\mathcal{C}$ is satisfied. Otherwise, assume without loss of generality that $\boldsymbol{w}' = (1 - d_x, d_y) \in \mathcal{C}$. Now, if there is a voter $i \in N \setminus N'$ such that $u_i(\boldsymbol{w}) \geqslant 1$, then we are done (as the group $N \setminus N'$ would be sufficiently satisfied if it were T-cohesive for $|T| = 1$). If that is not the case, then all voters $j \in N \setminus N'$ are unanimous on ballot $\boldsymbol{b}_j = (d_x, 1 - d_y)$. But then, since $\mathcal{C}$ satisfies the NFD property, there exists some outcome $\boldsymbol{w}'' \in \mathcal{C}$ such that $\boldsymbol{w}'' = 1 - d_y$. Then, $u_i(\boldsymbol{w}'') \geqslant 1$ for all $i \in N$ and no deviation is possible.

Finally, consider $|T| = m = 2$. In order for a group N' that is T-cohesive to have a $(S, \boldsymbol{w})$-deviation for $|S| \leqslant |T|$, it must be the case that $N' = N$, and $u_i(\boldsymbol{w}) = 0$ for all $i \in N$. By property NFD, there must be some outcome $\boldsymbol{w}' \neq \boldsymbol{w} \in \mathcal{C}$, and thus $u_i(\boldsymbol{w}') \geqslant 1$ for all $i \in N$.

Case $\boldsymbol{m = 3}$: Let $(\boldsymbol{B}, \mathcal{C})$ be an election instance satisfying the conditions in the proposition statement. We now reason on the existence of possible T-cohesive groups that are a witness to the violation of cohEJR$_\mathcal{C}$, for each possible size $1 \leqslant |T| \leqslant 3$.

For $|T| = 1$, suppose by contradiction that for all $\boldsymbol{w} \in \mathcal{C}$, there is some voter group N' such that $|N'| \geqslant n/3$ and each voter in N' has satisfaction of 0. Thus, for all voters $i \in N'$ we have $\boldsymbol{b}_i = \bar{\boldsymbol{w}}$. Moreover, for a T-cohesive group with an $(S, \boldsymbol{w})$-deviation for $|S| = |T| = 1$ to be possible, there has to exist a $\boldsymbol{w}' \in \mathcal{C}$ whose decisions differ from $\boldsymbol{w}$ in exactly one issue (i.e., $\text{Agr}(\boldsymbol{w}, \boldsymbol{w}') = 2$). To fit all these disjoint T-cohesive groups for $|T| = 1$, one for each outcome in $\mathcal{C}$, it must be that $n \geqslant |\mathcal{C}| \cdot n/3$, hence $|\mathcal{C}| \leqslant 3$ must hold. If $|\mathcal{C}| = 1$, the NFD property cannot be met. If $|\mathcal{C}| = 2$, then the two feasible outcomes cannot differ in the decision of only one issue while also satisfying the NFD property. For $|\mathcal{C}| = 3$, to get a T-cohesive voter group with an $(S, \boldsymbol{w})$-deviation for $|S| = |T| = 1$ at every $\boldsymbol{w} \in \mathcal{C}$, the three outcomes must differ by at most one decision, contradicting the NFD property.

For $|T| = 2$, we only consider $(S, \boldsymbol{w})$-deviations from a T-cohesive group N' with $|S| \in \{1, 2\}$. Consider the case of $|S| = 1$. Without loss of generality, assume that $\boldsymbol{w} = (0, 0, 0)$ and that there exists a T-cohesive group N' (where $|N'| \geqslant n \cdot 2/3$) with every voter having satisfaction $u_i(\boldsymbol{w}) < 2$, with an $(S, \boldsymbol{w})$-deviation towards some outcome, e.g., $\boldsymbol{w}' = (1, 0, 0)$. So $u_i(\boldsymbol{w}') \geqslant 1$ holds for at least one voter $i \in N'$. If N' has no $(S, \boldsymbol{w}')$-deviation, then there are not a witness to a violation of cohEJR$_\mathcal{C}$. Otherwise, suppose no voter in N' has satisfaction of 2 (so all voters in N' have a ballot $(1, 1, 1)$) and the group N' has an $(S, \boldsymbol{w}')$-deviation, to outcome $\boldsymbol{w}'' = (1, 1, 0)$. Then all voters in N' have satisfaction of at least 2 with outcome $\boldsymbol{w}''$. Suppose that the remaining $n/3$ of the voters $N \setminus N'$ have an $(S, \boldsymbol{w}'')$-deviation (so each of these voters derives zero satisfaction from outcome $\boldsymbol{w}''$ and all disagree with group N' on the first two issues a_1 and a_2). This must be a deviation of size $|S| = 1$ for the $n/3$ of the voters $N \setminus N'$ to demand it. This deviation can only be to outcome $(1,1,1)$, $(0,1,0)$ or $(1,0,0)$. By the NFD property, we know that one of the outcomes $\{(0,0,1),(0,1,1),(1,0,1),(1,1,1)\}$ must be in $\mathcal{C}$. See that for each of the outcomes $(1,1,1)$, $(1,0,1)$ and $(0,1,1)$, some voter in N' gets at least satisfaction of 2 while all voters in $N \setminus N'$ get at least 1 in satisfaction, thus, cohEJR$_\mathcal{C}$ is satisfied. Now, out of these outcomes $\{(0,0,1),(0,1,1),(1,0,1),(1,1,1)\}$, if only outcome $\boldsymbol{w}''' = (0,0,1)$ is in $\mathcal{C}$, then group N' has no $(S, \boldsymbol{w}''')$-deviations of size $|S| = 1$ as this is only possible to one of $(0,1,1)$ or $(1,0,1)$. So in this case too, cohEJR$_\mathcal{C}$ is satisfied.

Now we look at the case for $|S| = 2$. Without loss of generality, consider the outcome $\boldsymbol{w} = (0,0,0)$ and assume that there exists a T-cohesive group N' (where $|N'| \geqslant n \cdot 2/3$) with an $(S, \boldsymbol{w})$-deviation towards some outcome, e.g., $\boldsymbol{w}' = (1,1,0)$. Thus, there is some voter i in N' with satisfaction $u_i(\boldsymbol{w}') \geqslant 2$. At this point, the only possible further $(S, \boldsymbol{w})$-deviation could arise for $|S| = 1$ in case there are $n/3$ voters in $N \setminus N'$ each have a satisfaction of 0 for $\boldsymbol{w}'$, i.e., each has the ballot $(0,0,1)$ and either one of the outcomes in $\{(1,0,0),(0,1,0),(1,1,1)\}$ is in $\mathcal{C}$. Now take instead that $u_i(\boldsymbol{w}') = 2$ and consider two cases where either voter i agrees or disagrees with the voters in $N \setminus N'$ on the decision of issue a_3. First, assume that voter $i \in N'$ *agrees* with the voters in $N \setminus N'$ on issue a_3 (so voter i had the ballot $\boldsymbol{b}_i = (1,1,1)$). Then if either $(0,1,1) \in \mathcal{C}$ or $(1,1,1) \in \mathcal{C}$ holds, we have that cohEJR$_\mathcal{C}$ is provided. And if $(0,0,1) \in \mathcal{C}$ holds, then voters in $N \setminus N'$ are entirely satisfied and the voters in N' may only have an $(S, \boldsymbol{w})$-deviation for $|S| \leqslant |T| = 2$ if either $(0,1,1) \in \mathcal{C}$ or $(1,1,1) \in \mathcal{C}$ holds (as they only 'flip' issues they disagree with), which means that cohEJR$_\mathcal{C}$ is provided. In the second case, assume that voter $i \in N'$ *disagrees* with the voters in $N \setminus N'$ on issue a_3 and so, voter i had the ballot $\boldsymbol{b}_i = (1,1,0)$. This means that $u_i(\boldsymbol{w}') = 3$ holds, hence, any outcome that the voters in $N \setminus N'$ propose given they have an $(S, \boldsymbol{w})$-deviation for $|S| = 1$, would be one that provides cohEJR$_\mathcal{C}$.

Finally, a T-cohesive group for $|T| = 3$ implies a unanimous profile; if there exists an allowable $(S, \boldsymbol{w})$-deviation for $|S| \leqslant |T| = 3$, then the outcome in $\mathcal{C}$ maximising the sum of agreement with the profile provides cohEJR$_\mathcal{C}$. □

We leave it open whether the above result holds if we do not restrict our view to binary election instances. Unfortunately, the good news ends there as we provide an example showing that cohJR$_\mathcal{C}$ cannot be guaranteed when do not have $m \leqslant 3$ (even for binary election instances).

Proposition 7.6. *There exists an election instance $(\boldsymbol{B}, \mathcal{C})$ where $m > 3$ and the constraint $\mathcal{C}$ has the NFD property but no cohJR$_\mathcal{C}$ outcome exists.*

Proof. Consider the binary election instance where voters are to choose the pattern and colour of each of the four quadrants of a circular logo. We can represent this with eight issues that have the following domains.

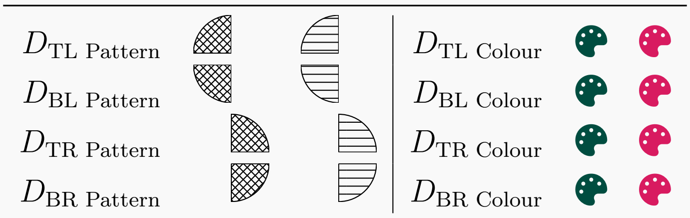

Now, take the constraint $\mathcal{C} = \{\boldsymbol{w}_1, \boldsymbol{w}_2, \boldsymbol{w}_3, \boldsymbol{w}_4\}$ having four outcomes being feasible where $\boldsymbol{w}_1 = (1, 1, 1, \ldots, 1)$, $\boldsymbol{w}_2 = (0, 0, 1, \ldots, 1)$, $\boldsymbol{w}_3 = (1, 1, 0, \ldots, 0)$ and $\boldsymbol{w}_4 = (0, 0, 0, \ldots, 0)$.

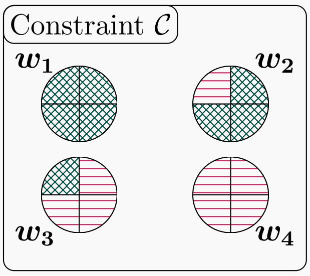

Then consider now a profile of four voters $\{1, 2, 3, 4\}$ with $\boldsymbol{b_i} = \boldsymbol{w_i}$, i.e., each of the four feasible outcomes is submitted by one of the voters. Given that $m = 8$, note that for every outcome $\boldsymbol{w} \in \mathcal{C}$, there exists some voter that deserves 2 in satisfaction by being T-cohesive for $|T| = 2$ with an $(S, \boldsymbol{w})$-deviation (by changing the pattern and colour of the top left quadrant) but with zero in satisfaction. And by cohJR$_{\mathcal{C}}$, such a voter would be entitled to at least 1 in satisfaction, so there is no outcome that provides cohJR$_{\mathcal{C}}$. □

We now turn our attention towards a weakening of cohEJR$_{\mathcal{C}}$ that takes inspiration from EJR-1 studied in the context of participatory budgeting (discussed in Chapter 3).

Definition 7.7 (cohEJR$_{\mathcal{C}}$-1). *Given an election* $(\boldsymbol{B}, \mathcal{C})$*, an outcome* $\boldsymbol{w}$ *provides* cohEJR$_{\mathcal{C}}$-1 *if for every* T*-cohesive group of voters* $N' \subseteq N$ *for some* $T \subseteq \mathcal{I}$ *with an* $(S, \boldsymbol{w})$*-deviation for some* $S \subseteq T$*, there exists a voter* $i \in N'$ *such that* $u_i(\boldsymbol{w}) \geqslant |T| - 1$.

As cohEJR$_{\mathcal{C}}$ implies cohEJR$_{\mathcal{C}}$-1, the results of Propositions 7.4 and 7.5 immediately apply to cohEJR$_{\mathcal{C}}$-1.

Corollary 7.7. *For binary election instances* $(\boldsymbol{B}, \mathcal{C})$ *with* $|\mathcal{C}| = 2$ *where the constraint* $\mathcal{C}$ *has the NFD property, cohEJR*$_{\mathcal{C}}$*-1 can always be satisfied.*

Corollary 7.8. *For binary election instances* $(\boldsymbol{B}, \mathcal{C})$ *with* $m \leqslant 3$ *where the constraint* $\mathcal{C}$ *has the NFD property, cohEJR*$_{\mathcal{C}}$*-1 can always be satisfied.*

Note that for the computational result for cohEJR$_{\mathcal{C}}$ in Proposition 7.2, a simple alteration of the proof given for Proposition 7.2 (replacing the value $|T|$ with $|T| - 1$ in the final satisfaction check) yields a corresponding computational result for cohEJR$_{\mathcal{C}}$-1.

Proposition 7.9. *Given an election instance* $(\boldsymbol{B}, \mathcal{C})$ *and an outcome* $\boldsymbol{w} \in \mathcal{C}$*, there exists an algorithm that decides in* $O((\max_{t \in [m]} |D_t|)^m |\mathcal{C}|^3 mn)$ *time whether outcome* $\boldsymbol{w}$ *satisfies cohEJR*$_{\mathcal{C}}$*-1.*

For the result of stating that cohEJR$_{\mathcal{C}}$ can be provided when $m = 2$ given that NFD holds (see Proposition 7.5), we can show something stronger for cohEJR$_{\mathcal{C}}$-1 by dropping the assumption that the NFD property holds.

Proposition 7.10. *For election instances* $(\boldsymbol{B}, \mathcal{C})$ *with* $m = 2$*, cohEJR*$_{\mathcal{C}}$*-1 can always be satisfied.*

Proof. Consider an election over two issues, where a T-cohesive group of voters has an $(S, \boldsymbol{w})$-deviation for some outcome $\boldsymbol{w}$. Observe that, when $m = 2$, $(S, \boldsymbol{w})$-deviation are only possible for $|S| \in \{1, 2\}$. Take a T-cohesive group N' for $|T| = 1$ with an $(S, \boldsymbol{w})$-deviation from $\boldsymbol{w}$ to some other feasible outcome $\boldsymbol{w}' \in \mathcal{C}$. Even if $u_i(\boldsymbol{w}) = 0$ for every voter $i \in N'$, we have $u_i(\boldsymbol{w}) \geqslant |T| - 1 = 1 - 1 = 0$, and thus cohEJR$_{\mathcal{C}}$-1 is satisfied. Take now a T-cohesive group N' for $|T| = 2$: for them to deviate, it must be the case that $N' = N$, and $u_i(\boldsymbol{w}) = 0$ for all $i \in N$. If they have an $(S, \boldsymbol{w})$-deviation for $|S| = |T| = 2$, the outcome $\boldsymbol{w'}$ they wish to deviate to must increase the satisfaction of each voter by at least 1, which thus satisfies $u_i(\boldsymbol{w}) \geqslant |T| - 1 = 2 - 1 = 1$, and thus cohEJR$_{\mathcal{C}}$-1. □

Can we show that an outcome providing cohEJR$_{\mathcal{C}}$-1 always exists when there are more than three issues, unlike for cohEJR$_{\mathcal{C}}$? Unfortunately, this is not the case, even assuming property NFD, as the same counterexample used to prove Proposition 7.6 yields the following (so also for binary election instances).

Proposition 7.11. *There exists an election instance* $(\boldsymbol{B}, \mathcal{C})$ *where* $m > 3$ *and the constraint* $\mathcal{C}$ *has the NFD property but there exists no outcome that satisfies cohEJR*$_{\mathcal{C}}$*-1.*

We demonstrate that the challenge of satisfying cohEJR$_\mathcal{C}$-1 lies in the constraints. To do so, we show that in the setting without constraints, it is always possible to find an outcome that provides cohEJR$_\mathcal{C}$-1. To do so, we define the constrained version of MES that has been studied for the public-decision setting without constraints (mentioned in Chapter 3). Our adaptation allows for the prices associated with fixing the outcome's decisions on issues to vary. This contrasts with the unconstrained MES that fixes the prices of every issue's decision to n from the onset. And this pricing is determined by a particular pricing type λ.

Definition 7.8 (MES$_\mathcal{C}$)**.** *The rule runs for at most m rounds. Each voter has a budget of m. In every round, for every undecided issue a_t in a partial outcome $\boldsymbol{w}^*$, we identify those issue-decision pairs (a_t, d) where fixing some decision $d \in D_t$ on issue a_t allows for a feasible outcome to be returned in future rounds. If no such issue-decision pair exists, then the rule stops. Otherwise, for every such pair (a_t, d), we calculate the minimum value for $\rho_{(a_t,d)}$ such that if each voter in $N(a_t, d)$ were to pay either $\rho_{(a_t,d)}$ or the remainder of their budget, then these voters could afford to pay the price $\lambda(a_t, d)$ (determined by the pricing type λ). If there exists no such value for $\rho_{(a_t,d)}$, then we say that the issue-decision pair (a_t, d) is* not affordable *in round, and if in a round, there are no affordable issue-decision pairs, the rule stops. Otherwise, we update $\boldsymbol{w}^*$ by setting decision d on issue a_t for the pair (a_t, d) with a minimal value $\rho_{(a_t,d)}$ (breaking ties arbitrarily, if necessary) and have each voter in $N(a_t, d)$ either paying $\rho_{(a_t,d)}$, or the rest of their budget. Note that MES$_\mathcal{C}$ may terminate with not all issues being decided and we assume that all undecided issues are decided arbitrarily.*

A natural candidate for a pricing type is the standard pricing of unconstrained MES where the price for every issue-decision pair (a_t, d) is set to $\lambda(a_t, d) = n$. And with this pricing, that we refer to as *unit* pricing λ_{unit}, we can show that MES$_\mathcal{C}$ satisfies cohEJR$_\mathcal{C}$-1 for unconstrained, binary elections.

Proposition 7.12. *For binary election instances, when $\mathcal{C} = \{0,1\}^m$, MES$_\mathcal{C}$ with unit pricing λ_{unit} satisfies cohEJR$_\mathcal{C}$-1.*

Proof. Take an outcome $\boldsymbol{w}$ returned by MES$_\mathcal{C}$ with unit pricing λ_{unit} and consider a T-cohesive group of voters N'. Let us assume that for every voter $i \in N'$, it holds that $u_i(\boldsymbol{w}) < |T| - 1$ and then set $\ell = |T| - 1$. So to conclude the run of MES$_\mathcal{C}$, each voter in N' paid for at most $\ell - 1 = |T| - 2$ issues.

Now, assume that the voters in N' paid at most $m/(\ell+1)$ for any decision on an issue. We know that each voter has at least the following funds remaining at that moment:

$$m - (\ell - 1)\frac{m}{\ell + 1} = \frac{2m}{\ell + 1} = \frac{2m}{|T|} \geqslant \frac{2n}{|N'|}.$$

The last step follows from the group N' being . So now we know that the voters in N' hold at least $2n$ in funds when some at the end of $\text{MES}_\mathcal{C}$'s run. Thus, we know that at least two issues have not been funded and for at least one of these two issues, at least half of N' agree on the decision of this issue (as the election instance is binary) and they hold enough funds to pay for it, hence, we have a contradiction to $\text{MES}_\mathcal{C}$ terminating.

Now, assume that some voter i in N' paid more than $m/(\ell+1)$ for a decision on an issue. Since we know that at the end of $\text{MES}_\mathcal{C}$'s execution, each voter in N' paid for at most $\ell - 1 = |T| - 2$ issues, then at the round r that voter i paid more than $m/(\ell+1)$ for an issue's decision, the voters in N' collectively held at least $2n$ in funds. Since at least two issues in were not funded, there exists some issue that could have been paid for in round r, where voters each pay $m/(\ell+1)$, contradicting the fact that voter i paid more than $m/(\ell+1)$ in round r. So, we have that this group of voters N' cannot exist and that $\text{MES}_\mathcal{C}$ satisfies $\text{cohEJR}_\mathcal{C}$-1. □

This result provides us with an axiom 'close to' EJR that we know is always satisfiable when the issues have size-two domains and there are no constraints.

7.4 Justified Representation with Agreement

Given the mostly negative results regarding the cohesiveness-EJR notion, we move on to justified representation based on agreement. We justify this move as the notion based on agreement is weaker and yields more positive results in the unconstrained setting. Thus, by assessing it here, we are able to establish a baseline of what can be achieved in terms of EJR-like proportionality guarantees in our constrained model. First, we formalise agreement-based EJR with the following axiom.

Definition 7.9 ($\text{agrEJR}_\mathcal{C}$)**.** *Given an election $(\boldsymbol{B}, \mathcal{C})$, an outcome $\boldsymbol{w}$ provides* $\text{agrEJR}_\mathcal{C}$ *if for every T-agreeing group of voters $N' \subseteq N$ for some $T \subseteq \mathcal{I}$ with an $(S, \boldsymbol{w})$-deviation for some $S \subseteq T$ with $|S| \leqslant |T| \cdot |N'|/n$, there exists a voter $i \in N'$ such that $u_i(\boldsymbol{w}) \geqslant |N'|/n \cdot |T|$.*

Now, in more unfortunate news, we find that $\text{agrEJR}_\mathcal{C}$ is not always satisfiable in general. In fact, the counterexample of Proposition 7.6 suffices to show this as each voter requires at least 1 in satisfaction for to $\text{agrEJR}_\mathcal{C}$ to be satisfied.

Proposition 7.13. *There exists an election instance where no outcome provides agrEJR$_\mathcal{C}$ (even when the NFD property holds for $\mathcal{C}$).*

We now focus on a particular class of constraints as we import agreement-EJR into our setting. Specifically, we consider a class that allows us to talk about how restrictive, and thus how costly, the fixing of a particular issue-decision pair is.

Akin to work by Rey et al. (2020, 2023), we consider constraints $\mathcal{C}$ that can be equivalently expressed as a set of implications $Imp_{\mathcal{C}}$, where each implication in $Imp_{\mathcal{C}}$ is a propositional formula with the following form: $\ell_{(a_x,d_x)} \rightarrow \ell_{(a_y,d_y)}$. This class of constraints allows us, for instance, to express simple dependencies and conflicts such as 'selecting x means that we must select y' and 'selecting x means that y cannot be selected', respectively. These constraints correspond to *propositional logic formulas* in *2CNF*.

Example 7.3. Take a set of issues $\mathcal{I} = \{a_1, a_2, a_3, a_4, a_5\}$ for a binary election instance. Here is an example of an implication set:

$$Imp_{\mathcal{C}} = \{(a_1, 1) \rightarrow (a_2, 1), (a_3, 1) \rightarrow (a_5, 0), (a_4, 1) \rightarrow (a_5, 0)\}.$$

Here, accepting a_1 means that a_2 must also be accepted while accepting either a_3 or a_4 requires the rejection of a_5.

Given a (possibly partial) outcome $\boldsymbol{w} \in \mathcal{C}$ and the set $Imp_{\mathcal{C}}$, we construct a directed *outcome implication graph* $G_{\boldsymbol{w}} = \langle H, E\rangle$ where $H = \bigcup_{a_t \in \mathcal{I}} \{(a_t, d) \mid d \in D_t\}$ as follows:

1. Add the edge $((a_x, d_x), (a_y, d_y))$ to E if $\ell_{(a_x,d_x)} \rightarrow \ell_{(a_y,d_y)} \in Imp_{\mathcal{C}}$ and $w_y \neq d_y$;

2. Add the edge $((a_y, d_y^*), (a_x, d_x^*))$ for all $d_y^* \neq d_y \in D_y, d_x^* \neq d_x \in D_x$ to E if $\ell_{(a_x,d_x)} \rightarrow \ell_{(a_y,d_y)} \in Imp_{\mathcal{C}}$ and $w_x = d_x$.

Given such a graph $G_{\boldsymbol{w}}$ for an outcome $\boldsymbol{w}$, we use $G_{\boldsymbol{w}}(a_x, d_x)$ to denote the set of all vertices that belong to some path in $G_{\boldsymbol{w}}$ having vertex (a_x, d_x) as the source (note that $G_{\boldsymbol{w}}(a_x, d_x)$ excludes (a_x, d_x)).

Example 7.4. Consider a binary election instance and take a set of issues $\mathcal{I} = \{a_1, a_2, a_3, a_4\}$ and the implication set $Imp_{\mathcal{C}} = \{(a_1, 1) \rightarrow (a_2, 1), (a_1, 1) \rightarrow (a_3, 1), (a_2, 1) \rightarrow (a_4, 1)\}$ of some constraint $\mathcal{C}$. Consider the outcome implication graph for $\boldsymbol{w}_1 = (0, 0, 0, 0)$ (vertices with no adjacent edges are omitted for readability):

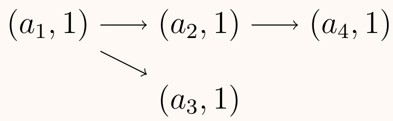

Then, we have that $G_{\boldsymbol{w}}(a_1, 1) = \{(a_2, 1), (a_3, 1), (a_4, 1)\}$ holds and therefore, we also get that $|G_{\boldsymbol{w}}(a_1, 1)| = 3$.

Thus, for an issue-decision pair (a_x, d_x), we can count the number of affected issues in setting a decision d_x for the issue a_x. This leads us to the following class of constraints.

Definition 7.10 (k-restrictive constraints)**.** *Take some constraint $\mathcal{C}$ expressible as a set of implications $Imp_{\mathcal{C}}$. For some positive integer $k \geqslant 2$, we say that $\mathcal{C}$ is* k-restrictive *if for every outcome $\boldsymbol{w} \in \mathcal{C}$, it holds that:*

$$\max \left\{ |G_{\boldsymbol{w}}(a_x, d_x)| \;\middle|\; (a_x, d_x) \in \bigcup_{a_t \in \mathcal{I}} \{(a_t, d) \mid d \in D_t\} \right\} = k - 1,$$

where $G_{\boldsymbol{w}}$ is the outcome implication graph constructed for outcome $\boldsymbol{w}$ and the implication set $Imp_{\mathcal{C}}$.

Intuitively, with a k-restrictive constraint, if one were to fix/change an outcome $\boldsymbol{w}$'s decision for one issue, this would require fixing/changing $\boldsymbol{w}$'s decisions on at most $k-1$ other issues. So intuitively, when dealing with k-restrictive constraints, we can quantify (at least loosely speaking) how 'difficult' it is to satisfy a constraint via the use of this value k. Thus, we can use this value k to account for the constraint's difficulty when designing proportionality axioms.

Before assessing how k-restrictive constraints affect our goal of providing proportionality, we touch on the computational complexity of checking, for some constraint $\mathcal{C}$, whether there exists a set of implications $Imp_{\mathcal{C}}$ that is equivalent to $\mathcal{C}$. For the case of binary elections, this problem been studied under the name of *Inverse Satisfiablility* and it has been shown that for formulas in 2CNF, the problem is in P (Kavvadias and Sideri, 1998). So in the remainder of the chapter, when we refer to a k-restrictive constraint $\mathcal{C}$, we thus assume that $\mathcal{C}$ is expressible using an implication set $Imp_{\mathcal{C}}$.

We now import the agreement-EJR notion and an approximate variant into our framework with constraints.

Definition 7.11 (α-agrEJR$_{\mathcal{C}}$-β)**.** *Given an election $(\boldsymbol{B}, \mathcal{C})$, some $\alpha \in (0, 1]$ and some positive integer β, an outcome $\boldsymbol{w}$ provides α-agrEJR$_{\mathcal{C}}$-β if for every T-agreeing group of voters $N' \subseteq N$ for some $T \subseteq \mathcal{I}$ with an $(S, \boldsymbol{w})$-deviation for some $S \subseteq T$ with $|S| \leqslant |T| \cdot {}^{|N'|}/_{n}$, there exists a voter $i \in N'$ such that $u_i(\boldsymbol{w}) \geqslant \alpha \cdot {}^{|N'|}/_{n} \cdot |T| - \beta$.*

With this axiom, we formalise agreement-EJR to our constrained public-decision model with the presence of the multiplicative and additive factors allowing us to measure how well rules satisfy this notion even if they fall short providing the ideal representation.[4] Note that for the sake of readability, when we have either $\alpha = 1$ or $\beta = 0$, we omit them from the notation when referring to α-agrEJR$_{\mathcal{C}}$-β.

Example 7.5. Suppose we have the binary election instance where there are four issues $\mathcal{I} = \{\text{TL Colour}, \text{TR Colour}, \text{BL Colour}, \text{BR Colour}\}$.

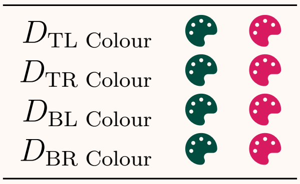

Here is the constraint $\mathcal{C} = \{(1,1,0,0),(1,1,1,0)\}$ that has two feasible outcomes which only differ in a single quadrant colour.

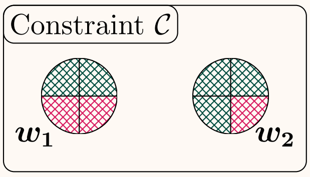

Next, suppose that there are two voters with ballots $\boldsymbol{b}_1 = (1,1,1,1)$ and $\boldsymbol{b}_2 = (0,0,0,0)$ that are represented as the following.

Voter	TL Colour	BL Colour	TR Colour	BR Colour	Voter Ballot
Voter 1					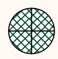
Voter 2					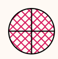

See that each voter deserves at least 2 in satisfaction according to agreement-EJR. See that outcome $\boldsymbol{w_1} = (1,1,0,0)$ provides agrEJR$_{\mathcal{C}}$ while the outcome $\boldsymbol{w_2} = (1,1,1,0)$ only provides $^1/_2$-agrEJR$_{\mathcal{C}}$ as voter 2 only obtains 1 in satisfaction whilst having a sufficiently small $(S, \boldsymbol{w})$-deviation for the issue BL Colour (deviating to outcome $\boldsymbol{w_1}$).

[4] Observe that we include the axiom's size requirement on the set S such that a group has an $(S, \boldsymbol{w})$-deviation in order to prohibit considering cases such as a single voter only having an $(S, \boldsymbol{w})$-deviation for $S = \mathcal{I}$ while not intuitively being entitled to that much representation.

Next, we analyse $\text{MES}_{\mathcal{C}}$ with respect to this axiom for k-restrictive constraints. We say that for $\text{MES}_{\mathcal{C}}$, the price for an issue-decision pair (a_x, d) given a partial outcome $\boldsymbol{w}^*$ is $\lambda(a_x, d) = n \cdot (|G_{\boldsymbol{w}^*}(a_x, d)| + 1)$ and we refer to this as a *fixed* pricing λ_{fix}. Then we can show the following for binary election instances.

Theorem 7.14. *For binary election instances $(\boldsymbol{B}, \mathcal{C})$ where $\mathcal{C}$ is k-restrictive for some k, $MES_{\mathcal{C}}$ with fixed pricing λ_{fix} satisfies $1/k$-agrEJR$_{\mathcal{C}}$-1.*

Proof. For a binary election instance $(\boldsymbol{B}, \mathcal{C})$ where $\mathcal{C}$ is k-restrictive, take an outcome $\boldsymbol{w}$ returned by $\text{MES}_{\mathcal{C}}$ with fixed pricing λ_{fix}. Consider a T-agreeing voter group N'. Let us assume that for every $i \in N'$, it holds that $u_i(\boldsymbol{w}) < {|N'|}/{nk} \cdot |T| - 1$ and then set $\ell = {|N'|}/{nk} \cdot |T| - 1$. So to conclude $\text{MES}_{\mathcal{C}}$, each voter $i \in N'$ paid for at most $\ell - 1 = {|N'|}/{nk} \cdot |T| - 2$ issues. Note that for a k-restrictive constraint $\mathcal{C}$, the maximum price $\text{MES}_{\mathcal{C}}$ with fixed pricing λ_{fix} sets for any issue-decision pair is nk (as at most k issues are fixed for a $\text{MES}_{\mathcal{C}}$ purchase). Now, assume that the voters in N' paid at most $m/(\ell+1)$ for any decision on an issue. We know that each voter has at least the following funds remaining at that moment:

$$m - (\ell - 1)\frac{m}{\ell + 1} = \frac{2m}{\ell + 1} = \frac{2m}{{|N'|}/{kn} \cdot |T|} = \frac{2mnk}{|N'||T|} \geqslant \frac{2nk}{|N'|}.$$

We now have that voter group N' holds at least $2nk$ in funds at the rule's end. Thus, we know that at least k issues have not been funded and for at least one of these k issues, at least half of N' agree on the decision for it (as the election is a binary instance) while having enough funds to pay for it. Hence, we have a contradiction to $\text{MES}_{\mathcal{C}}$ terminating.

Now, assume that some voter $i \in N'$ paid more than $m/(\ell+1)$ for fixing an issue's decision. Since we know that at the end of $\text{MES}_{\mathcal{C}}$'s run, each voter in N' paid for at most $\ell - 1$ issues, then at the round r that voter i paid more than $m/(\ell+1)$, the voters group N' collectively held at least $2nk$ in funds. Since at least k issues in were not funded, there exists some issue that could have been paid for in round r, where voters each pay $m/(\ell+1)$. This contradicts the fact that voter i paid more than $m/(\ell+1)$ in round r. So, we have that this group of voters N' cannot exist, which concludes the proof. □

Towards an even more positive result, and one where we are not limited to binary election instances, we now provide an adaptation of the MeCorA rule (mentioned in Chapter 3). In the unconstrained public-decision model, MeCorA is presented by Skowron and Górecki (2022) as an auction-style variant of MES that allows voter groups to change the decision of an issue all while increasing the price for any further change to this issue's decision. In our constrained model, groups are allowed to pay for changes to the decisions on sets of issues, as long as these changes represent a feasible deviation.

Definition 7.12 (MeCorA$_\mathcal{C}$)**.** *Take some constant $\epsilon > 0$. Start by setting $\lambda_t = 0$ as the current price of every issue $a_t \in \mathcal{I}$, endow each voter $i \in N$ with a personal budget of m and take some arbitrary, feasible outcome $\boldsymbol{w} \in \mathcal{C}$ as the current outcome. A groups of voters can 'update' the current outcome $\boldsymbol{w}$'s decisions on some issues $S \subseteq \mathcal{I}$ if the group:*

(i) can propose, for each issue $a_t \in S$, a new price $\lambda_t^* \geqslant \lambda_t + \epsilon$,

(ii) can afford the sum of new prices for issues in S, and

(iii) has an $(S, \boldsymbol{w})$-deviation.

The rule then works as follows. Given a current outcome $\boldsymbol{w}$, it computes, for every non-empty $S \subseteq \mathcal{I}$, the smallest possible value $\rho_{(t,S)}$ for each issue $a_t \in S$ such that for some N', if voters in N' each pay $\rho_S = \sum_{a_t \in S} \rho_{(t,S)}$ (or their remaining budget), then N' is able to 'update' the decisions on every $a_t \in S$ as per conditions $(i) - (iii)$. If there exists no such voter group for issues S then it sets $\rho_S = \infty$.

If $\rho_S = \infty$ for every $S \subseteq \mathcal{I}$, the process terminates and returns the current outcome $\boldsymbol{w}$. Otherwise, it selects the set S with the lowest value ρ_S (any ties are broken arbitrarily) and does the following:

1. *updates the current outcome $\boldsymbol{w}$'s decisions on issues in S to the decisions agreed upon by the voters with the associated $(S, \boldsymbol{w})$-deviation,*
2. *updates the current price of every issue $a_t \in S$ to λ_t^*,*
3. *returns all previously spent funds to all voters who paid for the now-changed decisions on issues in S,*
4. *and finally, for each voter in N', deduct $\sum_{a_t \in S} \rho_{(t,S)}$ from their personal budget (or the rest of their budget).*

Next, we show the representation guarantees can be achieved on instances with k-restrictive constraints via the use of modified version of MeCorA$_\mathcal{C}$. Moreover, we can drop the restriction to binary election instances that was key for the result of Theorem 7.14. In this MeCorA$_\mathcal{C}$ variant, we first partition the voter population into groups where members of each group agree on some set of issues. Then, for each group, its members may only pay to change some decisions as a collective and only on those issues that they agree on. In contrast to MeCorA$_\mathcal{C}$, voter groups cannot pay to change some decisions if this leads to the group's members gaining 'too much' satisfaction from the altered outcome (i.e., a voter group exceeding their proportional share of their agreed-upon issues, up to some additive factor q that parameterises the rule).

Definition 7.13 (Greedy MeCorA$_\mathcal{C}$-q)**.** *The set of the voters N is partitioned into p disjoints sets $N(T_1), \ldots, N(T_p)$ such that:*

(i) for every $x \in \{1, \ldots, p\}$, a voter group $N(T_x) \subseteq N$ is T_x-agreeing for some $T_x \subseteq \mathcal{I}$, and

(ii) for all $x \in \{1, \ldots, p-1\}$, it holds that $|N(T_x)| \cdot |T_x| \geqslant |N(T_{x+1})| \cdot |T_{x+1}|$

As with MeCorA$_\mathcal{C}$, voter groups shall pay to change the decisions of some issues during the rule's execution. However, given the initial partition, during the run of Greedy MeCorA$_\mathcal{C}$-q, the voters in $N(T_x)$ may only change decisions for the issues in T_x.

Moreover, if a voter group $N(T_x)$ for some $x \in \{1, \ldots, p\}$ wishes to change some decisions at any moment during the process, this change does not lead to any voter in $N(T_x)$ having satisfaction greater than ${}^{|N(T_x)|}/_{n} \cdot |T_x| - q$ with the updated outcome. Besides these two differences, the rule works exactly as MeCorA$_\mathcal{C}$.

Now, we can show the following for Greedy MeCorA$_\mathcal{C}$-q working on a k-restrictive constraint. For this result, we require the additional assumption that voter ballots represent feasible outcomes in $\mathcal{C}$.

Theorem 7.15. *For election instances $(\boldsymbol{B}, \mathcal{C})$ where voters' ballots are consistent with the constraint $\mathcal{C}$ and $\mathcal{C}$ is k-restrictive for some $k \geqslant 2$, Greedy MeCorA$_\mathcal{C}$-$(k-1)$ satisfies agrEJR$_\mathcal{C}$-$(k-1)$.*

Proof. Take an outcome $\boldsymbol{w}$ returned by Greedy MeCorA$_\mathcal{C}$-$(k-1)$. Assume that $\boldsymbol{w}$ does not provide agrEJR$_\mathcal{C}$-$(k-1)$. Thus, there is a T-agreeing group N' such that $u_i(\boldsymbol{w}) < {}^{|N'|}/_{n} \cdot |T| - k + 1 = \ell$ holds for every $i \in N'$. Now, consider the partition of voters $N(T_1), \ldots, N(T_p)$ constructed by Greedy MeCorA$_\mathcal{C}$-$(k-1)$ to begin its run. Assume first that there is some $x \in \{1, \ldots, p\}$ such that $N' = N(T_x)$, i.e., voters N' appear in their entirety in said partition. We then have $T = T_x$. Moreover, voters in N' each contribute to at most ℓ decisions at any moment of the run of Greedy MeCorA$_\mathcal{C}$-$(k-1)$, as this is the limit the rule imposes on their total satisfaction. We now consider two cases. Assume that the voters in N' contributed at most ${}^{m}/_{(\ell+k-1)}$ to change some decisions during the rule's execution. It follows that each voter has at least the following funds remaining: $m - (\ell-1) \cdot {}^{m}/_{(\ell+k-1)} \geqslant {}^{nmk}/_{|N'||T|}$.

In this case, the voters in N' would have at least ${}^{nmk}/_{|T|}$ in collective funds, so it follows that each distinct $(S, \boldsymbol{w})$-deviation available to N' must cost at least ${}^{nmk}/_{|T|}$. As N' is T-agreeing, it must be that N' has at least a ${}^{(|T|-\ell+1)}/_{k}$ many $(S, \boldsymbol{w})$-deviations due to $\mathcal{C}$ being k-restrictive and as the voters' ballots are consistent with $\mathcal{C}$.

Now, consider the case where some voter in N' contributed more than $m/(\ell+k-1)$ to change some decisions. The first time that this occurred, the change of decisions did not lead to any voter in N' obtaining a satisfaction greater than $\ell = |N'|/n \cdot |T| - k + 1$ (otherwise the rule would not allow these voters to pay for the changes). Thus, each voter in N' must have contributed to at most $\ell - 1$ issues before this moment. From the reasoning above, it must hold that in this moment, each voter held at least $nmk/|N'||T|$ in funds with there being at least $(|T|-\ell+1)/k$ feasible deviations available to N' and each such deviation costing at least $nmk/|T|$. So in both cases, for the $(S, \boldsymbol{w})$-deviations that are present in T that voters in N' wish to make, outcome $\boldsymbol{w}$'s decisions must have been paid for by voters within the remaining voter population $N \setminus N'$. And so, these decisions must have cost the voters in $N \setminus N'$ at least:

$$\frac{nmk}{|T|} \cdot \left(\frac{|T| - \ell + 1}{k} \right) = \frac{nm}{|T|} \cdot \left(|T| - \frac{|N'|}{n} \cdot |T| + k \right)$$

$$> \frac{nm}{|T|} \cdot \left(\frac{n|T| - |N'||T|}{n} \right) = m(n - |N'|).$$

However, voters $N \setminus N'$ have at most $m(n - |N'|)$ in budget. Thus, the rule cannot have terminated with the voter group N' existing.

Now, assume that the group N' did not appear in their entirety within the partition $N(T_1), \ldots, N(T_p)$ made by Greedy MeCorA$_{\mathcal{C}}$-$(k-1)$. This means that some voter $i \in N'$ is part of another voter group $N(T_x)$ that is T_x-agreeing such that $|N(T_x)|/n \cdot |T_x| \geqslant N'/n \cdot |T|$. Now, recall that for each voter group $N(T)$ in the partition, the voters in $N(T)$ have the same satisfaction to end the rule's execution (as they only pay to flip decisions as a collective). Thus, from the arguments above, it holds for this voter $i \in N' \cap N(T_x)$ that $u_i(\boldsymbol{w}) \geqslant |N(T_x)|/n \cdot |T_x| - k + 1 \geqslant |N'|/n \cdot |T| - k + 1$, which contradicts the assumption that every voter in N' has satisfaction less than $|N'|/n \cdot |T| - k + 1$.

□

Now, we offer yet another way towards producing proportional outcomes when using k-restrictive constraints. It is a constrained adaptation of the LS-PAV (detailed in Chapter 2). Recall that LS-PAV is polynomial-time computable and satisfies EJR. In the multiwinner voting setting, the rule begins with an arbitrary committee of some fixed size k and in iterations, searches for any swaps between committee members and non-selected candidates that brings about an increase of the PAV score by at least n/k^2. To translate LS-PAV to our model, we must define the PAV score of some feasible outcome $\boldsymbol{w} \in \mathcal{C}$:[5]

[5]Note the overload of terms when referencing the PAV score of feasible outcomes in this setting and PAV score of committees in the multiwinner setting (see Definition 2.16).

Definition 7.14 (PAV score of an outcome)**.** *Given a constrained election instance* $(\boldsymbol{B},\mathcal{C})$*, the PAV score of an outcome* $\boldsymbol{w} \in \mathcal{C}$ *is defined to be:*

$$PAV\text{-}score_{\mathcal{C}}(\boldsymbol{w}) = \sum_{i \in N} \sum_{t=1}^{u_i(\boldsymbol{w})} \frac{1}{t}.$$

With this PAV score notion in place, we can now lift LS-PAV to our setting with constraints in the following manner.

Definition 7.15 (LS-PAV$_{\mathcal{C}}$)**.** *Beginning with an arbitrary outcome* $\boldsymbol{w} \in \mathcal{C}$ *as the current winning outcome, the rule looks for all possible deviations. If there exists an* $(S,\boldsymbol{w})$*-deviation for some voter group to some outcome* $\boldsymbol{w}' \in \mathcal{C}$ *such that* $PAV\text{-}score_{\mathcal{C}}(\boldsymbol{w}') - PAV\text{-}score_{\mathcal{C}}(\boldsymbol{w}) \geqslant n/m^2$*, i.e., the new outcome* $\boldsymbol{w}'$ *yields a PAV score that is at least* n/m^2 *higher than that of* $\boldsymbol{w}$*, then the rule sets* $\boldsymbol{w}'$ *as the current winning outcome. The rule terminates once there exists no deviation that improves on the PAV score of the current winning outcome by at least* n/m^2*.*

As there is a maximum obtainable PAV score, LS-PAV$_{\mathcal{C}}$ is guaranteed to terminate. The question is how long this rule takes to return an outcome when we have to take k-restrictive constraints into account.

Proposition 7.16. *For elections instances where* $\mathcal{C}$ *is* k*-restrictive where* k *is a fixed constant and* $k \geqslant 2$*, LS-PAV*$_{\mathcal{C}}$ *terminates in polynomial time.*

> Proof. We show that given an outcome $\boldsymbol{w}$, finding all possible deviations can be done in polynomial time for a k-restrictive constraint $\mathcal{C}$. This can be done by exploiting the presence of the implication set $Imp_{\mathcal{C}}$. Note that the size of the implication set $Imp_{\mathcal{C}}$ is polynomial in the number of issues. So we can construct the outcome implication graph of $Imp_{\mathcal{C}}$ and the outcome $\boldsymbol{w}$ in polynomial time. Then for each issue $a_t \in \mathcal{I}$, we can find the set $G_{\boldsymbol{w}}(a_t, d)$ for some $d \neq w_t \in D_t$ in polynomial time and the issue-decision pairs represent the required additional decisions to be fixed in order to make a deviation from outcome $\boldsymbol{w}$ by changing the $\boldsymbol{w}$'s decision on issue a_t to d. Doing this for each issue a_t allows us to find a deviation that can improve the PAV score, if such a deviation exists. With similar reasoning used in other settings (Aziz et al., 2017a; Chandak et al., 2024), we end by noting that since there is a maximum possible PAV score for an outcome, and each improving deviation increases the PAV score by at least n/m^2, the number of improving deviations that LS-PAV$_{\mathcal{C}}$ makes is polynomial in the number of issues m. □

Off the back of this positive result, our attention immediately turns to LS-PAV$_{\mathcal{C}}$'s performance with respect to constrained agreement-EJR. We find the following holds under similar conditions of the result of Theorem 7.15.

Theorem 7.17. *For election instances $(\boldsymbol{B}, \mathcal{C})$ where the voters' ballots are consistent with the constraint $\mathcal{C}$ and $\mathcal{C}$ is k-restrictive for some $k \geqslant 2$, LS-PAV$_\mathcal{C}$ satisfies $^2/_{(k+1)}$-agrEJR$_\mathcal{C}$-$(k-1)$.*

Proof. For an election instance $(\boldsymbol{B}, \mathcal{C})$ where $\mathcal{C}$ is k-restrictive for $k \geqslant 2$, take an outcome $\boldsymbol{w}$ returned by LS-PAV$_\mathcal{C}$ and consider a group of voters N' that agree on some set of issues T. Let us assume that for every voter $i \in N'$, it holds that $u_i(\boldsymbol{w}) < {}^2/_{k+1} \cdot {}^{|N'|}/_n \cdot |T| - k + 1$ and then set $\ell = {}^2/_{k+1} \cdot {}^{|N'|}/_n \cdot |T| - k + 1$. We use r_i to denote the number of outcome $\boldsymbol{w}$'s decisions that a voter $i \in N$ agrees with.

For each voter $i \in N \setminus N'$, we calculate the maximal reduction in PAV score that may occur from a possible deviations by LS-PAV$_\mathcal{C}$ when $\mathcal{C}$ is k-restrictive. This happens when for each of at most $^{r_i}/_k$ deviations, we decrease their satisfaction by k and remove $\sum_{t=0}^{k-1} {}^1/_{(r_i - t)}$ in PAV score. So for these voters in $N \setminus N'$, we deduct at most the following:

$$\sum_{N \setminus N'} \frac{r_i}{k} \cdot \left(\sum_{t=0}^{k-1} \frac{1}{r_i - t} \right) \leqslant \sum_{N \setminus N'} \frac{r_i}{k} \cdot \left(\frac{\sum_{t=1}^{k} t}{r_i} \right) = \frac{k+1}{2} \cdot (n - |N'|).$$

Now, so there are $|T| - (\ell - 1) = |T| - \ell + 1$ issues that all voters in N' agree on but they disagree with outcome $\boldsymbol{w}$'s decisions on these issues. Since we assume the constraint is k-restrictive, then for each of these $|T| - \ell + 1$ issues, they fix at most $k-1$ other issues and thus, there are at least $^{(|T| - \ell + 1)}/_k$ feasible deviations that can be made by LS-PAV$_\mathcal{C}$ amongst these issues. For the voters in N', we now consider the minimal increase in PAV score that may occur from these possible deviations by LS-PAV$_\mathcal{C}$. For each such deviation, we increase their satisfaction by at least k and thus, for a voter $i \in N'$, we increase the PAV score by $\sum_{t=1}^{k} {}^1/_{(r_i + t)}$. Since for each voter $i \in N'$ we have $r_i \leqslant \ell - 1$, and as there are at least $^{(|T| - \ell + 1)}/_k$ feasible deviations in T, it follows that we add at least the following to the PAV score:

$$\frac{|T| - \ell + 1}{k} \cdot \left(\sum_{i \in N'} \sum_{t=1}^{k} \frac{1}{r_i + t} \right) \geqslant \frac{|T| - \ell + 1}{k} \cdot \left(\sum_{i \in N'} \sum_{t=1}^{k} \frac{1}{\ell + t - 1} \right).$$

Taking into account that $k \geqslant 2$ and $\ell = {}^{2|N'||T|}/_{(n(k+1))} - k + 1$, then with further simplification, we find that at least the following is added to the PAV score:

$$> \frac{n(k+1)}{2} - |N'| + \frac{n(k+1)}{|T|} \geqslant \frac{k+1}{2} \cdot (n - |N'|) + \frac{n(k+1)}{|T|}.$$

So the total addition to the PAV score due to satisfying voters in N' is strictly greater than the PAV score removed for the added dissatisfaction of voters in $N \setminus N'$ (which is at most $^{(k+1)(n-|N'|)}/_2$). And specifically, this change in score is at least $^{n(k+1)}/_{|T|} > {}^{n}/_{|T|}$ and thus, at least one of the $^{(|T|-\ell+1)}/_k$ many deviations must increase the PAV score by more than:

$$\frac{k}{|T|-\ell+1} \cdot \frac{n}{|T|} \geqslant \frac{1}{|T|} \cdot \frac{n}{|T|} \geqslant \frac{n}{|T|^2} \geqslant \frac{n}{m^2}.$$

Thus, LS-PAV$_{\mathcal{C}}$ would not terminate but would instead make this deviation in order to increase the total PAV score. Thus, contradicting that such a group N' cannot exist. □

With this result, we have a rule that when focused on k-restrictive constraints, is both polynomial-time computable and provides substantial proportional representation guarantees (assuming voter ballots are constraint consistent).

7.5 Proportionality via Priceability

With this section, we offer an alternative to the justified-representation-like interpretation of proportional representation, and this is through the notion of priceability (see Definition 2.12). Recent work has shown the promise of this market-based approach for a general social choice model (Masarík et al., 2023) and the sequential choice model (Chandak et al., 2024). We look to employ it for constrained public decisions (albeit looking at a weaker priceability axiom than the axiom that Masarík et al. (2023) studied).

Definition 7.16 (Priceability). *Each voter has a personal budget of m and they have to collectively fund the decisions on some issues, with each decision coming with some price. A* price system $\boldsymbol{ps} = (\{p_i\}_{i\in N}, \{\pi_{(a_t,d)}\}_{(a_t,d)\in H})$ *where* $H = \bigcup_{a_t\in\mathcal{I}}\{(a_t,d) \mid d \in D_t\}$ *is a pair consisting of (i) a collection of* payment functions $p_i : \mathcal{I} \times \{0,1\} \to [0,b]$, *one for each voter* $i \in N$, *and (ii) a collection of* prices $\pi_{(a_t,d)} \in \mathbb{R}_{\geqslant 0}$, *one for each decision pair* (a_t,d) *for* $a_t \in \mathcal{I}$ *and* $d \in D_t$. *We consider priceability with respect to outcomes* $\boldsymbol{w} \in \mathcal{C}$ *where decisions are made on all issues. We say that an outcome* $\boldsymbol{w} = (w_1,\dots,w_m)$ *is* priceable *if there exists a price system* $\boldsymbol{ps}$ *such that:*

(P1) : *For all* $a_t \in \mathcal{I}$ *and* $d \in D_t$, *it holds that if* $d \neq \boldsymbol{b}_i^t$ *we have* $p_i(a_t,d) = 0$, *for every* $i \in N$.

(P2) : $\sum_{(a_t,d)\in H} p_i(a_t,d) \leqslant m$ *for every* $i \in N$ *where it holds that* $H = \bigcup_{a_t\in\mathcal{I}}\{(a_t,d) \mid d \in D_t\}$.

(P3) : $\sum_{i\in N} p_i(a_t,d) = \pi_{(a_t,w_t)}$ *for every* $a_t \in \mathcal{I}$.

(P4) : $\sum_{i \in N} p_i(a_t, d) = 0$ *for every* $a_t \in \mathcal{I}$ *and every* $d \neq w_t \in D_t$.

(P5) : *There exists no group of voters* N' *with an* $(S, \boldsymbol{w})$*-deviation for some* $S \subseteq \mathcal{I}$*, such that for each* $a_t \in S$*:*

$$\sum_{i \in N'} \left(m - \sum_{(a'_t, d') \in H} p_i(a'_t, d') \right) > \pi_{(a_t, w_t)}$$

where $H = \bigcup_{a_t \in \mathcal{I}} \{(a_t, d) \mid d \in D_t\}$.

Condition (P1) states that each voter only pays for decisions that she agrees with; (P2) states that a voter does not spend more than her budget m; (P3) states that for every decision in the outcome, the sum of payments for this decision is equal to its price; (P4) states that no payments are made for any decision not in the outcome; and, finally, (P5) states that for every set of issues S, there is no group of voters N' agreeing on all decisions for issues in S, that collectively hold more in unspent budget to 'update' outcome $\boldsymbol{w}$'s decision on every issue $a_t \in S$ to a decision that they all agree with (where 'updating' these issues leads to a feasible outcome). We illustrate priceability in our setting with the following example of a binary election instance.

Example 7.6. Take a binary election instance with four issues $\mathcal{I} = \{\text{TL Colour}, \text{TR Colour}, \text{BL Colour}, \text{BR Colour}\}$.

$D_{\text{TL Colour}}$	
$D_{\text{TR Colour}}$	
$D_{\text{BL Colour}}$	
$D_{\text{BR Colour}}$	

Here is the constraint $\mathcal{C} = \{\boldsymbol{w}_1, \boldsymbol{w}_2\}$ where the two feasible outcomes are $\boldsymbol{w}_1 = (1, 1, 0, 0)$ and $\boldsymbol{w}_2 = (1, 1, 1, 1)\}$.

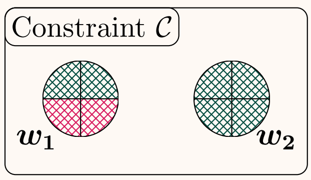

Now, take two voters with the ballots $\boldsymbol{b}_1 = (1, 1, 1, 1)$ and $\boldsymbol{b}_2 = (0, 0, 0, 0)$.

Voter	TL Colour	BL Colour	TR Colour	BR Colour	Voter Ballot
Voter 1					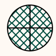
Voter 2					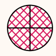

Note that outcome $\boldsymbol{w_1} = (1,1,1,1)$ is not priceable as any price system where voter 2 does not exceed her budget would have voter 2 having enough in leftover budget to cause a violation of condition (P5) (with her entire budget being leftover, she can afford more than price of the $(S, \boldsymbol{w})$-deviation to outcome $\boldsymbol{w}$). On the other hand, $\boldsymbol{w_2} = (1,1,0,0)$ is priceable where we set the price of this outcome's decisions to 1.

The following result gives some general representation guarantees whenever we have priceable outcomes.

Proposition 7.18. *Consider a priceable outcome* $\boldsymbol{w}$ *with price system* $\boldsymbol{ps} = (\{p_i\}_{i \in N}, \{\pi_{(a_t,d)}\}_{(a_t,d)\in H})$ *where* $H = \bigcup_{a_t \in \mathcal{I}} \{(a_t, d) \mid d \in D_t\}$*. Then, for every* T*-cohesive group of voters* $N' \subseteq N$ *for some* $T \subseteq \mathcal{I}$ *with an* $(S, \boldsymbol{w})$*-deviation for some* $S \subseteq T$*, it holds that:*

$$\sum_{i \in N'} u_i(\boldsymbol{w}) \geqslant \frac{n}{q} \cdot |T| - |S|$$

where $q = \max\{\pi_{(a_t,w_t)}\}_{a_t \in S}$.

Proof. Take a priceable outcome $\boldsymbol{w}$ and consider a T-cohesive group of voters N'. Suppose that $\sum_{i \in N'} u_i(\boldsymbol{w}) < {}^n\!/_q \cdot |T| - |S|$ where $q = \max\{\pi_{(a_t,w_t)}\}_{a_t \in S}$. As a group, the voters N' have a budget of $m|N'|$. Now, the voters in N' collectively contributed to at most ${}^n\!/_q \cdot |T| - |S| - 1$ decisions in outcome $\boldsymbol{w}$, and for each decision, the price was at most q (as q is the the price system's maximal price). So, we have that voter group N' has at least the following in leftover budget:

$$m|N'| - q \cdot \left(\frac{n}{q} \cdot |T| - |S| - 1\right) \geqslant m \cdot \frac{n|T|}{m} - n|T| + q|S| + q = q \cdot (|S| + 1).$$

Note we made use of the fact that N' is T-cohesive. Thus, we know that N' has strictly more than $q|S|$ in funds and for each issue in $a_t \in S$, holds more than in funds than $q \geqslant \pi_{(a_t,w_t)}$. This presents a violation of condition P5 of priceability. Hence, voter group N' cannot exist. □

However, we now must ascertain whether priceable outcomes always exist, regardless of the nature of the constraint. We see that this is possible thanks to the rule we have already defined, namely MeCorA$_\mathcal{C}$. The next result shows that MeCorA$_\mathcal{C}$ captures the notion of priceability.

Proposition 7.19. *MeCorA$_\mathcal{C}$ always returns priceable outcomes.*

Proof. Let $\boldsymbol{w} = (w_1, \ldots, w_m)$ be the outcome returned by MeCorA$_\mathcal{C}$. We define the following price system **ps**: For each issue $a_t \in \mathcal{I}$, fix the prices $\pi_{(a_t,w_t)} = \pi_{(a_t,d)} = \lambda_t$ for all $d \neq w_t \in D_t$ where λ_t is issue a_t's last MeCorA$_\mathcal{C}$ price (before being set to ∞) prior to the rule's termination. Fix the payment functions p_i for each voter to the money they spent to end the execution of MeCorA$_\mathcal{C}$. Observe that the priceability conditions (P1)-(P4) clearly hold: since we have that, to end MeCorA$_\mathcal{C}$'s run, voters do not pay for decisions that (*i*) they do not agree with (condition (P1)) and (*ii*) are not made by outcome $\boldsymbol{w}$ (condition (P4)); MeCorA$_\mathcal{C}$ limits each voter a budget of m (condition (P2)) (P2); and the sum of payments for decisions made by outcome $\boldsymbol{w}$ will equal exactly $\pi_{(a_t,w_t)} = \lambda_t$ (condition (P3)). Now, for condition (P5), note that if such a group of voters N' existed for some set of issues S, then MeCorA$_\mathcal{C}$ would not have terminated as this group of voters could have changed the decisions of these issues in S while increasing each issues' prices. □

This is a positive result that, combined with that of Proposition 7.18, gives us a rule that always returns us priceable outcomes for any election instance.

7.6 Chapter Summary

In this chapter, we first considered two different interpretations of justified representation from multiwinner voting and adapted them to a public-decision model with constraints. In analysing the feasibility of the axioms, we devised restricted classes of constraints (the NFD property, simple implications and k-restrictiveness). While we could show mostly negative results for the satisfaction of cohesiveness-EJR under constraints, we were able to adapt successfully three known voting rules (MES, MeCorA and LS-PAV) to yield positive proportional guarantees that meet, in an approximate sense, the requirements of agreement-EJR. Additionally, we defined a suitable notion of priceability and showed that our adaptation of MeCorA always returns priceable outcomes.

Future Work. Our work opens up a variety of paths for future research. First, assessing a class of constraints that are more expressive than the simple implications seems a natural starting point in extending our work. Then, on a more technical level, it would be interesting to check if the representation guarantees that are offered by $\text{MES}_{\mathcal{C}}$, Greedy $\text{MeCorA}_{\mathcal{C}}$-$(k-1)$ and $\text{LS-PAV}_{\mathcal{C}}$ still hold for a wider range of election instances. Also, it is imperative that our axioms based on the cohesiveness notion such as $\text{cohEJR}_{\mathcal{C}}$ are explored when restricted to k-restrictive constraints and other restricted classes of constraints that we did not consider. Regarding our adaptation of priceability, the question is open as to whether there are more constrained public-decision rules that always produce complete priceable outcomes. Given that we opted to represent the constraints as an enumeration of all feasible outcomes, it is natural to ask what occurs to results such as Propositions 7.1 and 7.2, when we consider the constraint takes a particular form of representation, e.g., $\mathcal{C}$ is represented as a Boolean formula of propositional logic. We also note some lingering computational questions such as the computational complexity of (i) computing outcomes for rules such as $\text{MES}_{\mathcal{C}}$ and Greedy $\text{MeCorA}_{\mathcal{C}}$-$(k-1)$ for general constraints, and (ii) of checking whether a given feasible outcome is priceable. Finally, the list of proportionality notions to be tested on the constraints test-bed is not exhausted, with the proportionality degree (Lackner and Skowron, 2023) most notably still to be considered.

Chapter 8

Constrained Shortlisting Using Approvals

Chapter 7 saw us study a model of public decisions that is enriched with constraints. This chapter sees us investigate a quite similar model. Specifically, we shall be using constraints to enrich the model of multiwinner voting where there is not a fixed, committee target size. Instead, the size of the committee is variable and is determined by the decision-making method. We refer to such a setting as *approval-based shortlisting* (Lackner and Maly, 2021a). The similarity of the public-decision and approval-based shortlisting models can be seen by also looking at the former (with binary election instances specifically) as an instance of variable-sized, approval-based multiwinner voting as the number of accepted issues is not fixed beforehand but may vary. This however, also reveals the difference between these settings. This difference lies in the way in which voters attain satisfaction from a committee. In the model of public decisions, when voters are presented a winning committee, they care about both the number of committee members that they approve of as well as the number of candidates that they disapprove of, that are not in the winning committee. This contrasts with the approval-based shortlisting model where voters only care about the former—the size of the intersection between the committee and their approval ballots—as is most common in the multiwinner voting literature. The similarity is such that the model of approval-based shortlisting is one that also admits a substantial number of real-world applications and is thus worthy of our attention.

Now that the intuition of this shortlisting model has been given, we move on to outline the goal of the chapter's analysis, and given what this thesis has discussed before this point, it is quite a familiar one: how can the proportionality notions of multiwinner voting be extended to this model of approval-based shortlisting with the presence of constraints? To be precise, we adapt the justified-representation notion to this model and study voting rules that we designed with the specific task of providing proportionality in mind. Our study of these rules takes the shape of testing them against some of our adapted proportionality axioms and also, assessing the computational aspects of computing these rules' outcomes. These novel rules take inspiration from two sources. The first is the class of Thiele methods (with us focusing on two well-known members of it), and the second is a class of rules introduced for the task of variable-sized multiwinner voting (Kilgour, 2016; Faliszewski et al., 2020).

Additional Related Work. We are looking at a model of multiwinner voting that does not have a target committee size that is fixed. There has been some interest in such settings. One example is that of public decisions, but in that work, the voter satisfaction is treated differently. In the public-decision model, voter satisfaction depends on those approved candidates that are in the committee as well as those disapproved candidates that are not in the committee. We take a different route to voter satisfaction: only consider voter approval of candidates elected to the committee. With this, we have the model of *variable-sized multiwinner voting* (Duddy et al., 2016; Kilgour, 2016; Faliszewski et al., 2020) or *approval-based shortlisting* (Lackner and Maly, 2021a). The model we study can be captured by the aforementioned general social-choice model of Masarík et al. (2023). Also, one may consider settings such as participatory budgeting to be related as in that setting, there is no limit on the number of projects that are to be implemented but rather, the only restriction on the selection of projects is imposed by the budget limits constraints. This can be done in our constrained shortlisting case when the considered constraint corresponds to a budget constraint.

Chapter Outline. We present the model of this chapter's study in Section 8.1. Section 8.2 is where we present the proportionality notion that we aim for and then we formalise it in an axiom called EJR_{sh} along with weakenings of EJR_{sh}. We then develop some novel voting rules and test them against the EJR_{sh} weakenings in Section 8.3. Section 8.3 also sees us assessing the computational complexity of these rules. We then conclude the chapter with Section 8.4.

8.1 The Model

This model is only a slight variation on the approval-based multiwinner voting model that we have seen previously. We take a set of voters $N = \{1, \dots, n\}$ and a set of m *candidates* $C = \{a, b, c, \dots\}$. In the usual fashion, each voter $i \in N$ then submits a set of candidates $A_i \subseteq C$ that indicate those candidates that voter i approves of. Then an *approval profile* is then a vector $\boldsymbol{A} = (A_1, \dots, A_n)$ of approval ballots, one for each voter.

This constrained shortlisting model then differs from that of standard approval-based multiwinner elections in the presence of a constraint $\mathcal{C}$ that denotes the set of feasible committees that may be of varying size, i.e., $\mathcal{C} \subseteq \mathcal{P}(C)$. The goal is to select, as an outcome, a feasible committee $W \in \mathcal{C}$ (we sometimes refer to such an outcome as a shortlist). A *constrained election instance* is a pair $(\boldsymbol{A}, \mathcal{C})$ and an (irresolute) *constrained shortlisting rule* F_{sh} takes a constrained election instance $(\boldsymbol{A}, \mathcal{C})$ as input, and maps it to a non-empty set of winning, feasible committees denoted as $F_{sh}(\boldsymbol{A}, \mathcal{C})$.

8.2 Defining Proportionality Axioms

In defining an axiom to aim for, we use the same approach as in our paper (where we did so for PJR instead of EJR) (Chingoma et al., 2022). Specifically, we aim for an output committee $W \in \mathcal{C}$ to provide EJR (in the sense of approval-based multiwinner voting (see Definition 2.9)) where the committee target size is set to $k = |W|$. Formally, we have the following definition:

Definition 8.1 (EJR$_{sh}$)**.** *Given an election instance $(\boldsymbol{A}, \mathcal{C})$, we say a committee $W \in \mathcal{C}$ provides* EJR$_{sh}$ *if W provides EJR for the multiwinner election instance $(\boldsymbol{A}, k)$ where $k = |W|$.*

To the best of our knowledge, with this PJR-variant of this axiom, we initiated the study of justified representation in this setting, however, the axiom was too narrow. Subsequently, more general, but different, notions have been studied. Specifically, we must touch on the work by Masarík et al. (2023). We note that our EJR$_{sh}$ axiom is not implied by the axioms that they developed independently (namely, their EJR and BEJR axioms that we mentioned in Chapter 3). Intuitively, see that our axiom effectively ignores the constraint in its definition whereas theirs do not. Formally, the difference becomes clear with the following bad news for our setting: it is not always possible to satisfy EJR$_{sh}$. This is not the case for their BEJR axiom, for example, which they show is always satisfiable.

Proposition 8.1. *There exist constrained election instances where EJR$_{sh}$ cannot be provided.*

> Proof (sketch). We only provide the intuition behind this result. The result is clear as there may not exist a committee $W \in \mathcal{C}$ that provides multiwinner EJR for size $|W|$ as the constraint $\mathcal{C}$ could exclude, for every $k \in [m]$, the committees W that provide EJR for the target committee size $k = |W|$. □

Now, we do have a way to determine if it is possible to satisfy this criterion given some election instance: for all possible committee sizes $k \in [m]$, apply a multiwinner voting rule that is known to satisfy EJR, e.g., MES or PAV, to see if at least one of committees returned by the rule is feasible. The existence of one such committee means that it is possible to provide EJR$_{sh}$. If the rule does

not return a feasible committee for any of the target sizes $k \in [m]$, then we know that EJR_{sh} cannot be provided. This same approach can be used to produce committees for our constrained shortlisting setting. For example, we could apply MES[arbitrary] for all possible committee sizes $k \in [m]$, in increasing order, until the rule produces a size k committee W that is feasible. This means that we can satisfy our axiom in polynomial-time (at least for election instances where such a committee exists).
The downside of this approach is its ad-hoc nature and we wish to look at rules that are more tailored for the task of variable-sized committee selection.

To conclude this section on a more positive side however, observe that if we have a constraint $\mathcal{C}$ where all committees W of some size k are feasible, i.e., $\mathcal{P}_k(C) \subseteq \mathcal{C}$, then EJR_{sh} can always be satisfied. This follows as at least one of these size-k committees satisfies multiwinner EJR.

Remark 8.1. *A key question is whether such constraints capture natural, real-world scenarios. We argue that they do so despite seeming to be so limited. For example, consider a conference organiser having to select a set of talks (of varying lengths) to fill the conference day's five timeslots. The constraint may require at least five talks to be selected (to fill each timeslot with at least one talk) with any set of, let's say, five to ten talks being a feasible outcome. However, due to the talks having different lengths, we also have that not all selections of size greater than ten are feasible due to the presence of too many lengthy talks, i.e., some size-*11 *selections may run over time when certain talks are chosen together. This case falls into the identified class of constraints.* ◁

The above results and observations prompt us to assess weakenings of EJR_{sh} that take greater account of the constraint $\mathcal{C}$ in their definitions. So we move forward with weaker axioms based on the notion of a candidate being *independent* within a constraint $\mathcal{C}$. Given a constraint $\mathcal{C}$ and a candidate $c \in C$, we say that candidate c is *$\mathcal{C}$-independent* if it holds that (i) for every committee $W \in \mathcal{C}$ with $c \notin W$, we have that $W \cup \{c\} \in \mathcal{C}$ and (ii) for every committee $W \in \mathcal{C}$ with $c \in W$, we have that $W \setminus \{c\} \in \mathcal{C}$. This leads us to this modified version of the ℓ-cohesiveness notion where we take into account the variable outcome size, denoted by t.

Definition 8.2 ((ℓ,t)-ind-cohesiveness)**.** *For some constraint $\mathcal{C}$, some positive integer t and an integer $\ell \in \{1, \dots, t\}$, we say that group of voters $N' \subseteq N$ is* (ℓ,t)-ind-cohesive *if both of the following conditions hold:*

- $|N'| \geqslant \ell \cdot n/t$.
- $|\{c \in \bigcap_{i \in N'} A_i \mid c \text{ is } \mathcal{C}\text{-independent}\}| \geqslant \ell$.

Observe that with the use of the $\mathcal{C}$-independence notion in defining cohesiveness, we take the constraint into account in a natural manner. We then use this variant of cohesiveness to define the following justified-representation axioms.

Definition 8.3 (Independent EJR, ℓ-ind-EJR$_{sh}$)**.** *When we are given an election instance* $(\boldsymbol{A}, \mathcal{C})$*, we say a committee* $W \in \mathcal{C}$ *provides* ℓ-ind-EJR$_{sh}$ *if for every* $(\ell,|W|)$*-ind-cohesive group of voters* N'*, there exists a voter* $i \in N'$ *such that* $|A_i \cap W| \geqslant \ell$*.*

Definition 8.4 (Independent PJR, ℓ-ind-PJR$_{sh}$)**.** *When we are given an election instance* $(\boldsymbol{A}, \mathcal{C})$*, we say a committee* $W \in \mathcal{C}$ *provides* ℓ-ind-PJR$_{sh}$ *if for every* $(\ell,|W|)$*-ind-cohesive group of voters* N'*, it holds that* $|\bigcup_{i\in N'} A_i \cap W| \geqslant \ell$*.*

Definition 8.5 (Independent JR, ind-JR$_{sh}$)**.** *Given an election instance* $(\boldsymbol{A}, \mathcal{C})$*, we say a committee* $W \in \mathcal{C}$ *provides* ℓ-ind-JR$_{sh}$ *if for every* $(1,|W|)$*-ind-cohesive group of voters* N'*, there exists a voter* $i \in N'$ *such that* $|A_i \cap W| \geqslant 1$*.*

With the parameter ℓ, we can account for committees that provide the proportionality property when one only considers $(\ell,|W|)$-ind-cohesive group for certain sizes. We then use *ind-JR$_{sh}$* to refer to ℓ-ind-EJR$_{sh}$ for $\ell = 1$. Observe that when the constraint $\mathcal{C}$ states that the feasible committees are all those committees of some target size k, then these axioms reduce to those in from multiwinner voting (albeit with EJR and PJR being parameterised by ℓ).

Note that this extends our earlier work in (Chingoma et al., 2022)—where we only studied ℓ-ind-PJR$_{sh}$—by looking at the stronger ind-EJR$_{sh}$. These axioms for independent candidates represent weak requirements of a committee that any reasonably proportional rule ought to satisfy in this constrained setting.

8.3 Rules to Provide Proportionality

We are now going to propose constrained shortlisting rules geared towards proportionality. In the variable-sized multiwinner literature, there is a class of rules that take both approvals and disapprovals into account when scoring a committee (Brams and Kilgour, 2014; Faliszewski et al., 2020), namely the class of *Net Approval Voting (NAV)* rules. We adopt this approval-disapproval dynamic and apply it to an outcome's shortlisted candidates.

Given a scoring vector $\boldsymbol{u}^{(m)} = (u_1, \ldots, u_m)$, let $f_{\boldsymbol{u}^{(m)}}(k) = \sum_{i=0}^{k} u_i$ be the sum of the first k elements of $\boldsymbol{u}^{(m)}$. Then we have the following general rule for two scoring vectors $\boldsymbol{a}^{(m)}$ and $\boldsymbol{d}^{(m)}$ for approvals and disapprovals, respectively.

Definition 8.6 ($(\boldsymbol{a}, \boldsymbol{d})$-NAV rule)**.** *Given two scoring vectors* $\boldsymbol{a}^{(m)}$ *and* $\boldsymbol{d}^{(m)}$*, we define an* $(\boldsymbol{a}, \boldsymbol{d})$-NAV *rule to be the following:*

$$(\boldsymbol{a}, \boldsymbol{d})\text{-}NAV(\boldsymbol{A}, \mathcal{C}) = \operatorname*{argmax}_{W\in\mathcal{C}} \sum_{i\in N} f_{\boldsymbol{a}^{(m)}}(|W \cap A_i|) - f_{\boldsymbol{d}^{(m)}}(|W \setminus A_i|).$$

We concretely analyse two instances of this general $(\boldsymbol{a}, \boldsymbol{d})$-NAV rule. The first attaches standard AV scoring to approved candidates in the shortlist, and 'penalises' disapprove shortlist members with an 'inverted' PAV scoring.

Definition 8.7 (c-PAV)**.** *c-PAV is defined as the* $(\boldsymbol{a}, \boldsymbol{d})$*-NAV rule with the scoring vectors* $\boldsymbol{a}^{(m)} = (1, 1, \ldots, 1)$ *and* $\boldsymbol{d}^{(m)} = (1/m, \ldots, 1/2, 1)$*.*

For the second rule, a voter awards points to an outcome as with α-CC scoring, but subtracts a point if the majority threshold, a commonly-used threshold (Alon et al., 2015; Faliszewski et al., 2020; Fishburn and Pekeč, 2017), of disapproved shortlist members is crossed.

Definition 8.8 (c-CC)**.** *c-CC is defined as the* $(\boldsymbol{a}, \boldsymbol{d})$*-NAV rule with the scoring vectors* $\boldsymbol{a}^{(m)} = (1, 0, \ldots, 0)$ *and* $\boldsymbol{d}^{(m)} = (0, \ldots, 0, 1, 0, \ldots, 0)$*. In* $\boldsymbol{d}^{(m)}$*, we set the scoring vector's 'threshold' at position* $\lceil m/2 \rceil + 1$*.*

Next, we test our proposed Thiele method adaptations against these axioms and get the following results. We begin with c-CC.

Theorem 8.2. *Assuming all candidates are* $\mathcal{C}$*-independent, c-CC satisfies ind-JR*$_{sh}$ *but fails* ℓ*-ind-PJR*$_{sh}$ *with* $\ell > 1$*.*

Proof. Our claim is that, for any $\ell > 1$, we can construct a candidate set C and approval profile $\boldsymbol{A}$ such that, for $\mathcal{C} = \mathcal{P}_+(C)$, there is a $W \in$ c-CC$(\boldsymbol{A}, \mathcal{C})$ that does not provide ℓ-ind-PJR$_{sh}$. Consider an arbitrary $\ell > 1$. We choose an agenda with an odd number $m \geqslant 5$ of candidates such that $\ell = \lfloor m/2 \rfloor$. We fix a set of voters N with $|N| = m+1$ and a voter subset $N^\ell \subset N$ such that:

$$|N^\ell| = \frac{|N| \cdot \lfloor m/2 \rfloor}{\lceil m/2 \rceil} = |N| - 2.$$

This ensures the existence of two voters $x, y \notin N^\ell$. Given this voter population, we can define the approval profile $\boldsymbol{A}$. In this profile, the voters in N^ℓ uniformly approve the same ℓ candidates, so we have $|\bigcap_{i \in N^\ell} A_i| = |\bigcup_{i \in N^\ell} A_i| = \ell$. For the approval ballots of voters $x, y \notin N^\ell$, we have two candidates $c_x, c_y \notin \bigcap_{i \in N^\ell} A_i$, i.e., neither c_x nor c_y are accepted by voters in N^ℓ, such that $A_x = \{c_x\}$ and $A_y = \{c_y\}$. This completes the construction of the profile. Note that $|\bigcup_{i \in N} A_i| = \lceil m/2 \rceil + 1$. Now we assess the candidates shortlisted by c-CC for the profile $\boldsymbol{A}$. Once the voter group N^ℓ is represented by $\ell - 1$ candidates, observe that (due to $\ell > 1$) accepting the candidate c_x for voter $x \notin N^\ell$ gives a greater score than accepting an ℓ-th candidate in $\bigcap_{i \in N^\ell} A_i$. The same holds for c_y and voter $y \notin N^\ell$. After c_x and c_y are selected, including an ℓ-th candidate in $\bigcap_{i \in N^\ell} A_i$ decreases the score as the majority threshold is crossed. Notice that c-CC accepts exactly $\lceil m/2 \rceil$ candidates and thus, from the definition of the group N^ℓ's size, it follows that N^ℓ is an ℓ-cohesive group. So we have an outcome where an ℓ-cohesive group is only represented by $\ell - 1$ candidates.

Next, we show that every $W \in$ c-CC$(\boldsymbol{A},\mathcal{C})$ provides ind-JR$_{sh}$ (ℓ-ind-PJR$_{sh}$/ℓ-ind-EJR$_{sh}$ for $\ell = 1$), assuming $\mathcal{C}$-independence throughout the candidate set. The argument is that any non-selected candidate in $\bigcap_{i\in N^*} A_i$ for an unrepresented (1,$|W|$)-ind-cohesive group N^* can be 'swapped' for some chosen candidate in a committee W. This only decreases the c-CC score if every already-selected candidate represents a unique group that is at least as large as N^*. Thus, any such candidate would at least match the contribution to the score of a candidate accepted by N^*. However, this cannot be the case for all $|W|$ candidates as this would imply that at least $|W| \cdot {}^{n}/_{|W|} = n$ voters have been represented thus far, contradicting our assumption of the existence of this unrepresented group N^*. □

With this result, we unsurprisingly find that our novel rule yields similar proportionality guarantees in the multiwinner voting literature (as Theorem 8.2 assumes that there are no constraints) as α-CC does.

Off the back of these results for c-CC, we now turn our attention towards c-PAV and find a more encouraging outcome.

Theorem 8.3. *A committee W returned by c-PAV satisfies ℓ-ind-EJR$_{sh}$ for every value $\ell > {}^{|W|}/_{(m-|W|+2)}$.*

Proof. Take a $W \in$ c-PAV$(\boldsymbol{A},\mathcal{C})$ and assume there is an (ℓ,$|W|$)-ind-cohesive group N^* such that $|W \cap A_i| < \ell$ for all voters $i \in N^*$. We can show that a committee W' that additionally selects a currently-rejected candidate in $\bigcap_{i\in N^*} A_i$, all else being equal, yields a strictly higher c-PAV score than W.

We now detail the change in score which occurs with the acceptance of candidate $c_x \in \bigcap_{i\in N^*} A_i \setminus W$. The group N^* adds at least $\ell \cdot {}^{n}/_{|W|}$ to the score. At least one voter is already satisfied by W and this voter deducts at most ${}^{1}/_{(m-|W|+1)}$ while at most $n - |N^*| - 1$ voters will each deduct at most ${}^{1}/_{(m-|W|)}$. So the score strictly increases when we have:

$$|N^*| \geqslant \ell \cdot \frac{n}{|W|} > \frac{n-|N^*|-1}{m-|W|} + \frac{1}{m-|W|+1}$$

$$> \frac{n-|N^*|-1}{m-|W|+1} + \frac{1}{m-|W|+1} = \frac{n-|N^*|}{m-|W|+1}$$

$$\geqslant \frac{n - {}^{\ell n}/_{|W|}}{m-|W|+1} = \frac{n|W| - n\ell}{|W|(m-|W|+1)} = \frac{n(|W|-\ell)}{|W|(m-|W|+1)}.$$

To conclude, observe that the score is strictly positive when we have:

$$\ell \cdot \frac{n}{|W|} > \frac{n(|W|-\ell)}{|W|(m-|W|+1)} \implies \ell > \frac{|W|}{m-|W|+2}.$$

So we find that c-PAV does provide ℓ-ind-EJR$_{sh}$ when we have that $\ell > {}^{|W|}/_{(m-|W|+2)}$ □

We have now assessed both of our constrained shortlisting rules with respect to the proportionality axioms that we have devised for this setting. With the next section, we look at the computational aspects of both c-CC and c-PAV.

This section continues by conducting a computational complexity analysis of our rules. To be precise, beginning with c-PAV, we consider the following decision problem that is similar to that seen in the previous chapter when assessing judgment aggregation rules (note the overloading of notation with the naming of the decision problem).

Outcome(c-PAV)	
Given:	An approval profile $\boldsymbol{A}$ for a candidate set C, a positive integer t and a subset of candidates $C' \subseteq C$.
Question:	Is there an outcome W of size t that is returned by c-PAV on profile $\boldsymbol{A}$ such that $C' \subseteq W$?

Note that for the Outcome(c-PAV) problem, the constraint $\mathcal{C}$ is not explicitly included in the problem input but rather, it is implicit that the constraint restrains the feasible committees to those of size t.

In order to show that Outcome(c-PAV) problem is NP-hard, we must first show NP-hardness for a variant of the classical problem of Independent Set where given an arbitrary graph, the question is whether the graph contains an independent set of certain size, with an independent set being a subset of the vertices S such that no two vertices in S have an edge connecting them. In the variant we consider, which refer to as Target-Independent Set, the size of the maximum independent set is known in advance, and the question is whether there is a maximum independent set that contains a given vertex

Lemma 8.4. Target-Independent Set *is* NP*-hard.*

Proof. We show NP-hardness by reducing from Independent Set where we are given a graph $G = \langle V, E \rangle$ and a positive integer t, with the question being whether there exists an independent set of size t in G. We reduce to Target-Independent Set as follows: the input for the problem if the graph $G^* = \langle V \cup \{v^*\}, E \rangle$ where is an extra vertex v^* such that no edge is connected to it. We claim that there exists an independent set of size t in G if and only if there exists an independent set of size $t+1$ in G^* that contains the vertex v^*.

($\Longrightarrow$) Assume that we have a positive instance of INDEPENDENT SET so there exists a size t independent set $S \subseteq V$ in G. As the vertex v^* is not connected to any vertex in G^*, it is contained in every independent set of G^* and specifically, it forms an independent set $S \cup \{v^*\}$ of size $t+1$ with S being an independent set in G^* as all is equal between G and G^* besides the presence of this vertex v^*.

($\Longleftarrow$) It is the clear to see that if in the modified graph G^*, there is an independent set $S^* \subseteq V \cup \{v^*\}$ of size $t+1$ that contains vertex v^*, then $S^* \setminus \{v^*\}$ is an independent set in the original graph G that is of size t. □

We are now able to show NP-hardness for the OUTCOME(c-PAV) problem.

Theorem 8.5. OUTCOME(*c-PAV*) *is* NP*-hard.*

Proof. We show NP-hardness by reducing from the following problem. Take a positive integer $t \in \mathbb{N}$ and an approval-based multiwinner election. Moreover, we know that there exists some $b \in \mathbb{N}$ such that: (*i*) the maximum PAV score of any committee of size t equals $b \cdot t$, and this can only be achieved by getting a score of 1 from $b \cdot t$ different voters, (*ii*) there exists at least one such committee of size $t \leqslant m$, and (*iii*) each voter approves of exactly 2 candidates. The problem is to decide, for a given candidate $c^* \in C$, whether it is part of a committee of size t with maximum PAV score.

This problem can straightforwardly be shown to be NP-hard, by adapting the proof of a known result for multiwinner PAV voting (Aziz et al., 2015, Theorem 1), which uses a reduction from the classical problem of INDEPENDENT SET. By instead considering TARGET-INDEPENDENT SET (which is NP-hard by Lemma 8.4), this proof directly yields NP-hardness of the problem that we will reduce from.

Assume that we are restricted to the setting where conditions (*i*)–(*iii*) hold. For the PAV problem that we reduce from, we have a set of voters N, a set of candidates $C = \{c_1, \ldots, c_m\}$ and the approval profile $\boldsymbol{A}$. In reducing to the problem of c-PAV, we use the same set of voters $N' = N$ and then we add m^2 dummy candidates to C such that no voter in N' approves of any of the dummy candidates. So, this means that the new set of candidates is $C' = \{c_1, \ldots, c_m, c_{m+1}, \ldots, c_{m+m^2}\}$ and we have the same approval profile $\boldsymbol{A}' = \boldsymbol{A}$ but now, it is with voters in N' voting over the modified candidate set C'. Finally, we construct a partial outcome $C' = \{c^*\}$. We then claim that the given candidate $c^* \in C$ is part of a committee of size t that with maximum PAV score on the election $(\boldsymbol{A}, t)$ if and only if there exists an outcome W returned by c-PAV$(\boldsymbol{A}', \mathcal{C})$ such that $C' = \{c^*\} \subseteq W$.

To verify correctness in both directions, observe that under the conditions $(i)-(iii)$, the size t committee that yields maximum PAV score is exactly the outcome of size t that obtains the maximum c-PAV score. To see this, recall that the maximum PAV score $b \cdot t$ for some $b \in \mathbb{N}$ occurs from $b \cdot t$ different voters giving the committee a score of 1, i.e., $|A_i \cap W| = 1$ for $b \cdot t$ voters $i \in N$. Let us consider the c-PAV score that this outcome W attains when considering that each voter approves of exactly two candidates. Note that the c-PAV scoring vectors for approvals and disapprovals are $\boldsymbol{a}^{(m)} = (1, \ldots, 1)$ and $\boldsymbol{d}^{(m)} = (1/(m+m^2), \ldots, 1/2, 1)$, respectively. The $b \cdot t$ voters that approve of a single candidate in W also give this outcome a positive score of $b \cdot t$ and each such voter removes $\sum_{j=0}^{t-2} 1/(m+m^2-j)$ from the outcome's c-PAV score. The remaining voters, who approve of no candidates in W, remove $\sum_{j=0}^{t-1} 1/(m+m^2-j)$ from the c-PAV score. Observe that this is the maximal c-PAV score in this instance where conditions $(i)-(iii)$ hold, so we know that this outcome W is one of the outcomes returned by c-PAV. So if it holds that the candidate c^* is in the committee W with a maximum PAV score, then this candidate must be in an outcome returned by c-PAV. □

Now, we find that not only is α-CC is NP-hard in the standard multiwinner voting setting (Lu and Boutilier, 2011; Procaccia et al., 2008), but computing outcomes for the c-CC rule in our constrained shortlisting setting is also hard. The decision problem that we use is the following, where we overload notation once more:

OUTCOME(c-CC)	
Given:	An approval profile $\boldsymbol{A}$ for a candidate set C, a constraint $\mathcal{C}$ and a subset of candidates $C' \subseteq C$.
Question:	Is there an outcome $W \in$ c-CC$(\boldsymbol{A}, \mathcal{C})$ of size t such that $C' \subseteq W$?

For the next result, we must introduce an additional complexity class that is not covered in Chapter 2. The class Θ_2^{p} is the class of decisions problems that can be decided in polynomial time by a Turing machine that (i) has access to an oracle that is complete for NP, and (ii) only makes a number of oracles queries that is logarithmic in the input (Arora and Barak, 2009).

We will now show that OUTCOME(c-CC) is Θ_2^{p}-complete and in our proof of Θ_2^{p}-hardness, we reduce from the following variant of the well-known MAXSAT problem (Papadimitriou, 1994):

TARGET-MAXSAT	
Given:	A set $\mathcal{L}$ of literals, two sets φ_1 and φ_2 of clauses, with clauses in both being of size at most 3, and some variable x^* occurring in φ_2.
Question:	Among the truth assignments that satisfy all clauses in φ_1 and that satisfy a maximum number of clauses in φ_2, is there a truth assignment that sets x^* to true?

It is known that MAXSAT is a Θ_2^{p}-complete problem. One can also prove that TARGET-MAXSAT is Θ_2^{p}-complete with proof of hardness making use of a reduction from the MAX-MODEL problem that is a Θ_2^{p}-complete (Chen and Toda, 1995; Krentel, 1988; Wagner, 1990).[1]

Theorem 8.6. OUTCOME(*c-CC*) *is* Θ_2^{p}*-complete.*

Proof. We describe Θ_2^{p}-hardness by reducing from the TARGET-MAXSAT described above. We omit the straightforward proof of membership in Θ_2^{p}.

We introduce a candidate c_l for each literal l over the variables in the sets φ_1 and φ_2. Hence, we have candidates such as c_{x_1} and $c_{\neg x_1}$ in the candidate set C. For each clause v_i appearing in φ_2, we create a voter i. These voters approve of those candidates which correspond to literals appearing in their associated clause. Selecting a candidate c_l in the outcome corresponds to setting a literal l to true while not selecting c_l represents setting the literal l to false. We construct a constraint $\mathcal{C}$ that expresses that exactly one of each pairs of candidates, c_{x_i} or $c_{\neg x_i}$, is selected in an outcome in such a way that satisfies all clauses in φ_1. Finally, we construct a partial outcome C' that only selects candidate c_{x^*}, i.e., $C' = \{c_{x^*}\}$. We claim that the outcomes of the c-CC rule over the constructed profile correspond to the truth assignments that satisfy all of φ_1 and that maximise the number of satisfied clauses in φ_2. Therefore, there is an outcome $W \in \text{c-CC}(\boldsymbol{A}, \mathcal{C})$ such that $C' = \{c_{x^*}\} \subseteq W$ if and only if the original instance of TARGET-MAXSAT is a yes-instance.

($\Longrightarrow$) Assume we have a truth assignment that satisfies all the clauses in φ_1 and maximises the satisfied clauses in φ_2 with variable x_* set to true. Since the clauses in φ_2 correspond to voters, we have that a maximum number of voters have been satisfied with respect to their approvals, as at least one of their approved issues are set to true. We must confirm that they do not have their disapproval thresholds crossed. Note that the number of issues that each

[1] For MAX-MODEL, we are given a satisfiable propositional formula φ along with a variable x^* that occurs in φ, and the question is whether there exists a satisfying assignment for φ that sets a maximal number of variables in φ to true while setting the variable x^* to true.

voter i disapproves of is $|A_i| + 2 \cdot (m - |A_i|)|$. Thus their majority threshold is set to $\lceil |A_i|/2 + (m - |A_i|) \rceil + 1$. The maximum number of disapprovals that can appear for a voter i, if their clause is satisfied in the Target-MaxSAT instance, is $|A_i| - 1 + (m - |A_i|) = m - 1$. Since each voter can approve at most three candidates, we know that the number of their disapprovals that can be selected in an outcome, given at least one of their approved candidates does appear, cannot reach their majority threshold. Thus, if their clause is satisfied then their disapproval threshold cannot be crossed. If this is not the maximum number of voters under c-CC then this cannot be the maximum number of clauses being satisfied in φ_2 by the truth assignment. So we have confirmed that such a truth assignment indeed satisfies a maximum number of voters under c-CC. And we know that the variable x_* is set to true, then that means that the candidate c_{x^*} is also selected in the corresponding outcome returned by c-CC.

($\Longleftarrow$) Assume we have an outcome $W \in \mathcal{C}$ with $c_{x^*} \in W$ that satisfies a maximum number of voters using the c-CC rule, thus representing a winning outcome. Given our constraint, we know we have a valid truth assignment as exactly one of the conflicting candidates, such as c_{x_1} or $c_{\neg x_1}$, are selected in W. This corresponds to literals being satisfied in Target-MaxSAT. We also have that this truth assignment ensures that all the clauses in φ_1 are satisfied by the construction of $\mathcal{C}$. And since a maximum number of voters are satisfied, we have a maximum number of φ_2 clauses being satisfied as well. And to conclude, as the candidate c_{x^*} is the outcome W, we have that the variable x^* is set to true in φ_2. □

These negative computational results make our $(\boldsymbol{a}, \boldsymbol{d})$-NAV rules less appealing for practical use in general. However, there may be some promise in defining approximate versions of these rules that are more computationally efficient, or possibly, in identifying scenarios where these rules are polynomial-computable, e.g., for certain types of preferences profiles.

8.4 Chapter Summary

This chapter saw us aim towards providing proportionality in the constrained approval-based shortlisting setting. In doing so, we suggested an adaptation of EJR to our setting called EJR_{sh}. Although we find that EJR_{sh} is not always satisfiable, when we assume that it is, we can find a suitable committee satisfying it in polynomial-time. To follow these results, we moved on to weaker axioms based on EJR_{sh}. These axioms made use of the notion of some candidates being independent of the constraint $\mathcal{C}$. This led to us to defining the class of $(\boldsymbol{a}, \boldsymbol{d})$-NAV rules that we deem to be natural extensions of the Thiele methods of multiwinner voting. These rules' performances in satisfying the independence-based weakenings of EJR_{sh} were mixed while the subsequent computational analysis provided more negative results as we found that the rules are hard to compute in general.

Future Work. Our most positive axiomatic results came via the weakening of EJR_{sh} using the notion of candidate independence from the constraint. The study of differing notions of weakening EJR_{sh} would then be a natural follow-up of our work and could prove to be a fruitful research direction. As we only considered constraints in their most general form, it is important to investigate whether a more positive outlook can be found when one restricts oneself to constraints of a certain structure. Also, it would be of interest to further study the class of $(\boldsymbol{a}, \boldsymbol{d})$-NAV rules to obtain greater insight into their axiomatic behaviour in a general sense and not only through the lens of proportional representation.

Coming to a Close

Chapter 9

Conclusion

With the exception of the detour in Chapter 6, there is a clear throughline within the thesis: how to extend proportional representation to richer domains. This has hopefully given readers a (mostly) seamless read up to this point and leaves them with only a single chapter to go. This last chapter comprises two sections. The first offers a summary of the work that the thesis has covered. The second presents thoughts on potential research paths that follow from the work that the thesis has covered.

9.1 Taking It All In

We now recap the what was presented in the various parts of this thesis.

Setting the Stage. We began with the introductory chapter of the thesis in **Chapter 1** and we followed with **Chapter 2** where the formal foundations of the thesis were laid. This included the presentation of axioms and proportional voting rules that are frequently referred to, adapted or used for inspiration in much of the work in the rest of the thesis. We then closed out this first part of the thesis with **Chapter 3** where we presented a 'mini-survey' of work from the literature that holds a similar goal to that of this thesis: providing proportional representation to complex domains.

Part One: Some Seats Have More Value Than Others. The thesis then shifted to the main goal in the next part which consisted of Chapters 4 and 5. In **Chapter 4**, we looked at a generalisation of the apportionment model that introduced weights that are associated with the parliamentary seats. The aim was to lift the proportionality axioms and apportionment methods from standard apportionment to this weighted-seat model. The axioms in focus were the two quota-based axioms, lower quota and upper quota, while we also defined weighted-seat apportionment methods (WSAMs) that generalised the standard methods of

D'Hondt, Adams and LRM. We then found that the most faithful adaptations of these axioms, namely WLQ^o and WUQ^o, were too strong for this setting and thus, yielded mostly negative results, both axiomatically and computationally. Our subsequent study of relaxations of these axioms—which were inspired by 'up-to-one/any' relaxations from the fair division literature—brought us more positive results as our WSAMs were each able to satisfy at least one axiom's relaxation. This upper-quota analysis also saw us investigate our WSAMs with respect to a related notion of envy-freeness and in this case, Adams$_\omega$ and Greedy$_\omega$ both fared well. Our axiomatic study also included the study of a weighted adaptation of house monotonicity which is an axiom that has been deemed desirable in the standard apportionment context. In that analysis, we found that the strongest house monotonicity adaptation was too strong for our WSAMs but our weighted-seat divisor methods, D'Hondt$_\omega$ and Adams$_\omega$, did satisfy the weakening called min-HM. Our analysis of the WSAMs was supplemented with an experimental study where we tested our WSAMs on (*i*) real-world Bundestag committee assignment data, and (*ii*) artificially generated data instances. These experiments yielded greater insight to our WSAMs and ultimately, our WSAMs proved to have differing but appealing qualities. **Chapter 5** saw us take the weight-seat notion of the previous chapter and introduce it the the approval-based multiwinner voting model. In terms of adapting multiwinner voting rules to this model, we defined two classes of rules, those that are seat-based and those that are weight-based. For the latter, we defined weighted generalisations of rules such as MES and seq-Phragmén. These WMWV rules were to then be tested on our weighted-seat version/s of the justified-representation notion. We found that the strongest EJR adaptation, called WEJR, was not always satisfiable. However, we found that the 'up-to' weakenings of WEJR, namely WEJR-1 and WEJR-X-r, could always be satisfied. This analysis highlighted the potential for a multiwinner-voting adaptation of the Greedy$_\omega$ method from Chapter 4 as it satisfied WEJR-X-r which is one of this chapter's strongest axioms that is always satisfiable. Then, to account for varying interpretations of EJR in this setting, we studied an alternative relaxation of WEJR and we referred to this as low-WEJR. This axiom proved most suitable for seat-based rules as our weight-based WMWV rules did not have positive results when in came to this low-WEJR axiom. To end this chapter, we presented a weight-seat translation of priceability from multiwinner voting. This unsurprisingly led to investigating PJR-like axioms for this setting and we found a more positive proportionality outlook for both seq-Phragmén$_\omega$ and MES$_\omega$.

Detour. This single-chapter part of the thesis represented a deviation in the content as with **Chapter 6**, we studied the simulation of multiwinner voting rules using judgment aggregation. Multiwinner voting rules that were simulated includes among others, the k-Borda and k-Copeland rules. Then, we not only simulated multiwinner rules that satisfy the notion of weak Gehrlein stability, such as k-Kemeny and NED, but also, we observed a natural way to define novel multiwinner rules that satisfy weak Gehrlein stability. We followed this simulation study with a dive into more computational considerations related to the simulations. Through the use of DNNF circuits to encode constraints, we saw that some positive results as some of our simulations were as computationally efficient as their multiwinner-voting versions. However, there were also more negative results as when some feasibility constraints are used in a simulation, such as k-INDIFF-INDIFF and k-INCOMP-INCOMP, the approach using DNNF circuits does not lead to a retention of computational efficiency in general.

Part Two: Constraining the Feasible Committees. We then returned to proportionality in the following part. Here, **Chapter 7** dealt with the public decisions model and **Chapter 8** saw us work with the shortlisting model. We now describe these chapters' findings in more detail. In importing proportionality to the constrained public-decision model in **Chapter 7**, we adapted both the agreement-EJR and cohesiveness-EJR notions from the standard public-decision model. In the chapter's analysis, we also defined multiple restricted classes of constraints such as the k-restrictive constraints or those constraints that satisfy the NFD property. With the constrained adaptation of the latter notion, cohesiveness-EJR, we could only show mostly negative results as the weaker agreement-EJR proved to be more suitable for this setting. For our approximation axiom based on agreement-EJR, we found that our adaptations of MES, MeCorA and LS-PAV each provided strong proportional guarantees. We also adapted priceability to this constrained setting and showed that an adaptation of MeCorA was another viable candidate for providing proportional outcomes. The second chapter of this part, and the thesis' final chapter prior to concluding, was **Chapter 8** and here, we defined an EJR adaptation for the constrained shortlisting model. Now, we saw that this axiom, called EJR_{sh}, cannot always be satisfied (although, when we assume it is, shortlists providing EJR_{sh} can be found in polynomial time). This saw us follow by examining weakenings of EJR_{sh} that were based on candidate independence. Then, in looking for rules that satisfied these weaker axioms, we introduced a class of $(\boldsymbol{a}, \boldsymbol{d})$-NAV rules. This class was inspired by both the Thiele methods of standard (fixed-sized) multiwinner voting and the NAV rules from variable-sized multiwinner voting. These rules provided a mixed bag in terms off the axioms with the rule based on PAV performing most positively. These lacklustre findings continued in the subsequent computational complexity analysis as we found that both of the rules that we considered were computationally hard to compute (in general).

9.2 Taking It Further

When we concluded each of the chapters that presented research (so Chapters 4-8), we presented brief, chapter-specific discussions dedicated to future work. As these parts were specifically related to each chapter's content, this section will not be used to rehash these specific research directions. Rather, it outlines more broad ideas that (*i*) stem from the thesis' content, and (*ii*) we deem deserving of attention from within the research community.

Getting Closer to Reality with More Complex Domains: If the goal of (much of) the computational social choice community is to have decision-making methods used in practice, then we argue that it is imperative to not only expand on 'fundamental' research done on the standard models (such as apportionment and multiwinner voting), but to explore richer models. From the side of practice, the real world presents plenty of decision-making processes where using a more standard model *suffices* for decision-makers, but where using an enriched model would yield more *significant* insights through bringing said processes closer to reality. We argue that such instances should be searched for and systematically investigated. This would be beneficial as it allows for the design of notions that are better suited to talk about proportionality in the respective scenarios. As suggested by Chapter 3, there are likely many such research lines that are currently underway, and the sheer number of the works being produced should only further encourage researchers to explore more interesting domains where proportionality naturally fits and yet proportionality research is sorely lacked.

Voter Satisfaction: An offshoot of the above point on moving closer to reality, is getting a handle of the notion of voter satisfaction in the various complex domains. When the standard models are enriched with additional components, it becomes less and less clear what exactly voters place value towards when it come to the final outcome. And this satisfaction notion is vital in the proportionality notions of justified representation. In fact, this is already a topic that is drawing interest such as with participatory budgeting (Brill et al., 2023b). This certainly warrants attention in many more, similarly rich, domains. However, from a methodological standpoint, it is unclear how exactly to tackle this issue. One option is to take the approach of Brill et al. (2023b) and abstract away from particular satisfaction notions. Alternatively, the use of real-world experimental studies may be an interesting route to take in order to gain insights on what voters value within the complex domains.

Broaden the Proportionality Notions: A follow-up from the last point. The main notion we studied, that of justified representation, relies on the notion of satisfaction but a notion such as priceability does not. This leads to a more broad research avenue of looking at alternative proportionality notions that are more specifically designed for the complex domains in focus.

More Emphasis on Limiting Overrepresentation: We now wrap the thesis up by highlighting that the proportionality content of this thesis, along with the research conducted by those working in computational social choice, puts a significant focus on avoiding underrepresentation when the goal is to ensure proportional representation. However, we believe that it is also of great importance to prevent some voter groups from receiving too much representation. The prominence of the upper quota axiom in the apportionment literature and (to some extent) the envy-freeness notion of fair division show that this is an aspect that has been considered. We believe that greater emphasis should be placed on this when investigating proportionality for models such as standard multiwinner voting or more complex domains such as participatory budgeting.

Bibliography

Stéphane Airiau, Haris Aziz, Ioannis Caragiannis, Justin Kruger, Jérôme Lang and Dominik Peters. 2023. Portioning using ordinal preferences: Fairness and efficiency. *Artificial Intelligence* 314 (2023), 103809. (Cited on pages 28 and 30.)

Noga Alon, Robert Bredereck, Jiehua Chen, Stefan Kratsch, Rolf Niedermeier and Gerhard J. Woeginger. 2015. How to Put Through Your Agenda in Collective Binary Decisions. *ACM Transactions on Economics and Computation* 4, 5:1-28 (2015). (Cited on page 170.)

Georgios Amanatidis, Georgios Birmpas, Aris Filos-Ratsikas and Alexandros A. Voudouris. 2022. Fair Division of Indivisible Goods: A Survey. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*. 5385-5393. (Cited on pages 9, 30, 40, and 64.)

Sanjeev Arora and Boaz Barak. 2009. *Computational Complexity: A Modern Approach*. Cambridge University Press. (Cited on pages 14 and 174.)

Kenneth J. Arrow. 1951. *Social Choice and Individual Values*. Cowles Foundation. (Cited on page 7.)

Kenneth J. Arrow, Amartya Sen and Kotaro Suzumura (Eds). 2002. *Handbook of Social Choice and Welfare* Volume 1. North-Holland. (Cited on page 7.)

Kenneth J. Arrow, Amartya Sen and Kotaro Suzumura (Eds). 2011. *Handbook of Social Choice and Welfare* Volume 2. North-Holland. (Cited on page 7.)

Haris Aziz, Anna Bogomolnaia and Hervé Moulin. 2020. Fair Mixing: The Case of Dichotomous Preferences. *ACM Transactions on Economics and Computation* 8, 18:1-27 (2020). (Cited on pages 30 and 81.)

Haris Aziz, Markus Brill, Vincent Conitzer, Edith Elkind, Rupert Freeman and Toby Walsh. 2017a. Justified Representation in Approval-Based Committee Voting. *Social Choice and Welfare* 48, 461-485 (2017). (Cited on pages 18, 19, 22, and 158.)

Haris Aziz, Edith Elkind, Piotr Faliszewski, Martin Lackner and Piotr Skowron. 2017b. The Condorcet Principle for Multiwinner Elections: From Shortlisting to Proportionality. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*. 84-90. (Cited on pages 109, 110, 120, 122, and 130.)

Haris Aziz, Edith Elkind, Shenwei Huang, Martin Lackner, Luis Sánchez Fernández and Piotr Skowron. 2018. On the Complexity of Extended and Proportional Justified Representation. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*. 902-909. (Cited on pages 19 and 22.)

Haris Aziz, Serge Gaspers, Joachim Gudmundsson, Simon Mackenzie, Nicholas Mattei and Toby Walsh. 2015. Computational Aspects of Multi-Winner Approval Voting. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 107-115. (Cited on pages 126 and 173.)

Haris Aziz and Barton Lee. 2020. The expanding approvals rule: Improving proportional representation and monotonicity. *Social Choice and Welfare*. 54, 1-45 (2020). (Cited on page 29.)

Haris Aziz, Hervé Moulin and Fedor Sandomirskiy. 2020. A Polynomial-time Algorithm for Computing a Pareto Optimal and Almost Proportional Allocation. *Operations Research Letters*. 48, 573-578 (2020). (Cited on page 40.)

Moshe Babaioff, Tomer Ezra and Uriel Feige. 2021. Fair-Share Allocations for Agents with Arbitrary Entitlements. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*. 127. (Cited on page 40.)

Michel L. Balinski. 2005. What is just?. *The American Mathematical Monthly* 112, 502-511 (2005). (Cited on page 8.)

Michel L. Balinski and Gabrielle Demange. 1989. Algorithms for proportional matrices in reals and integers. *Mathematical Programming* 45, 193-210 (1989). (Cited on page 28.)

Michel L. Balinski and Gabrielle Demange. 1989. An axiomatic approach to proportionality between matrices. *Mathematics of Operations Research* 14, 700-719 (1989). (Cited on page 28.)

Michel L. Balinski and Svetlozar T. Rachev. 1993. Rounding Proportions: Rules of Rounding. *Numerical Functional Analysis and Optimization* 14, 475-501 (1993). (Cited on page 8.)

Michel L. Balinski and H. Peyton Young. 1975. The Quota Method of Apportionment. *The American Mathematical Monthly* 82, 701-730 (1975). (Cited on page 43.)

Michel L. Balinski and H. Peyton Young. 1982. *Fair Representation: Meeting the Ideal of One Man, One Vote.* Yale University Press. (Cited on pages 8, 13, 16, 17, 39, 42, 53, 56, 60, 65, 66, and 67.)

Siddhartha Banerjee, Vasilis Gkatzelis, Safwan Hossain, Billy Jin, Evi Micha and Nisarg Shah. 2023. Proportionally Fair Online Allocation of Public Goods with Predictions. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*. 20-28. (Cited on page 34.)

Salvador Barberà and Danilo Coelho. 2008. How to Choose a Non-Controversial List with k Names. *Social Choice and Welfare* 31, 79-96 (2008). (Cited on pages 120 and 122.)

Dorothea Baumeister, Gábor Erdélyi and Jörg Rothe. 2016. Judgment Aggregation. In *Economics and Computation: An Introduction to Algorithmic Game Theory, Computational Social Choice, and Fair Division.* Jörg Rothe (ed). Springer. 361-391. (Cited on page 110.)

Xiaohui Bei, Xinhang Lu and Warut Suksompong. 2022. Truthful Cake Sharing. In *Proceedings of the 36th AAAI Conference On Artificial Intelligence (AAAI)*. 4809-4817. (Cited on page 34.)

Gustavo M. Bodanza, Fernando A. Tohmé and Marcelo R. Auday. 2017. Collective Argumentation: A Survey of Aggregation Issues around Argumentation Frameworks. *Argument & Computation* 8, 1-34 (2017). (Cited on page 7.)

Niclas Boehmer, Markus Brill, Alfonso Cevallos, Jonas Gehrlein, Luis Sánchez Fernández and Ulrike Schmidt-Kraepelin. 2024. 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)*. 9519-9527. (Cited on page 103.)

Niclas Boehmer, Markus Brill and Ulrike Schmidt-Kraepelin. 2022. Proportional Representation in Matching Markets: Selecting Multiple Matchings under Dichotomous Preferences. In *Proceedings of the 21st International Conference on Autonomous Agents And Multiagent Systems (AAMAS)*. 136-144. (Cited on page 35.)

Anna Bogomolnaia, Hervé Moulin and Richard Stong. 2005. Collective choice under dichotomous preferences. *Journal of Economic Theory* 122, 165-184 (2005). (Cited on pages 30 and 31.)

Sirin Botan and Ulle Endriss. 2020. Majority-Strategyproofness in Judgment Aggregation. In *Proceedings of the 14th International Conference on Autonomous Agents And Multiagent Systems (AAMAS)*. 186-194. (Cited on page 121.)

Sylvain Bouveret, Yann Chevaleyre and Nicolas Maudet. 2016. Fair Allocation of Indivisible Goods. In *Handbook of Computational Social Choice*. Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang and Ariel D. Procaccia (Eds). Cambridge University Press. Chapter 12, 284-309. (Cited on pages 8 and 83.)

Steven Brams and D. Marc Kilgour. 2014. Satisfaction Approval Voting. *Voting Power and Procedures: Essays in Honour of Dan Felsenthal And Moshé Machover*. 323-346. (Cited on page 169.)

Florian Brandl, Felix Brandt, Matthias Greger, Dominik Peters, Christian Stricker and Warut Suksompong. 2022. Funding Public Projects: A Case for the Nash Product Rule. *Journal of Mathematical Economics* 99, 102585. (Cited on pages 30 and 81.)

Felix Brandt. 2017. Rolling the Dice: Recent Results in Probabilistic Social Choice. *Trends in Computational Social Choice*. Ulle Endriss (Ed). AI Access. Chapter 1, 3-26. (Cited on pages 31 and 81.)

Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang and Ariel D. Procaccia (Eds). 2016. *Handbook of Computational Social Choice*. Cambridge University Press. (Cited on page 7.)

Markus Brill, Hayrullah Dindar, Jonas Israel, Jérôme Lang, Jannik Peters and Ulrike Schmidt-Kraepelin. 2023. Multiwinner Voting with Possibly Unavailable Candidates. In *Proceedings of the 37th AAAI Conference On Artificial Intelligence (AAAI)*. 5532-5539. (Cited on page 35.)

Markus Brill, Jonas Israel, Evi Micha and Jannik Peters. 2022. Individual Representation in Approval-Based Committee Voting. In *Proceedings of the 36th AAAI Conference On Artificial Intelligence (AAAI)*. 4892-4899. (Cited on page 19.)

Markus Brill, Stefan Forster, Martin Lackner, Jan Maly and Jannik Peters. 2023. Proportionality in Approval-Based Participatory Budgeting. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*. 5524-5531. (Cited on pages 30, 40, and 184.)

Markus Brill, Paul Gölz, Dominik Peters, Ulrike Schmidt-Kraepelin and Kai Wilker. Approval-Based Apportionment. 2020. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence*. 1854-1861. (Cited on page 28.)

Markus Brill, Paul Gölz, Dominik Peters, Ulrike Schmidt-Kraepelin and Kai Wilker. 2024. Approval-based apportionment. *Mathematical Programming* 203, 77-105 (2024). (Cited on page 28.)

Markus Brill, Jean-François Laslier and Piotr Skowron. 2017. Multiwinner Approval Rules as Apportionment Methods. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI).* 414-420. (Cited on pages 8 and 24.)

Markus Brill, Evangelos Markakis, Georgios Papasotiropoulos and Jannik Peters. 2023b. Proportionality Guarantees in Elections with Interdependent Issues. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI).* 2537-2545. (Cited on pages 32 and 135.)

Markus Brill and Jannik Peters. 2023. Robust and Verifiable Proportionality Axioms for Multiwinner Voting. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC).* 301. (Cited on page 19.)

Markus Brill and Jannik Peters. 2024. Completing Priceable Committees: Utilitarian and Representation Guarantees for Proportional Multiwinner Voting. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI).* 9528-9536. (Cited on page 25.)

Eric Budish. 2011. The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes. *Journal of Political Economy* 119, 1061-1103 (2011). (Cited on page 40.)

Laurent Bulteau, Noam Hazon, Rutvik Page, Ariel Rosenfeld and Nimrod Talmon. 2021. Justified Representation for Perpetual Voting. *IEEE Access* 9, 96598-96612 (2021). (Cited on pages 33, 135, and 142.)

Bundestag. 2023. Functions and Responsibilities of the Committees. (https://www.bundestag.de/en/committees/function-245820,2023). Last visited on 2023-07-31. (Cited on page 39.)

Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah and Junxing Wang. 2019. The Unreasonable Fairness of Maximum Nash Welfare. *ACM Transactions on Economics and Computation* 7, 1-32 (2019). (Cited on page 40.)

Alfonso Cevallos and Alistair Stewart. 2021. A verifiably secure and proportional committee election rule. In *Proceedings of the 3rd ACM Advances in Financial Technologies (AFT).* 29-42. (Cited on page 103.)

Mithun Chakraborty, Ulrike Schmidt-Kraepelin and Warut Suksompong. 2021. Picking sequences and monotonicity in weighted fair division. *Artificial Intelligence* 301, 103578 (2021). (Cited on pages 8, 40, 45, 55, 63, 64, 66, 67, 68, and 83.)

John R. Chamberlin and Paul N. Courant. 1983. Representative Deliberations and Representative Decisions: Proportional Representation and the Borda Rule. *The American Political Science Review* 77, 718-733 (1983). (Cited on page 23.)

Nikhil Chandak, Shashwat Goel and Dominik Peters. 2024. Proportional Aggregation of Preferences for Sequential Decision Making. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*. 9573-9581. (Cited on pages 33, 41, 135, 139, 142, 158, and 160.)

Zhi-Zhong Chen and Seinosuke Toda. 1995. The Complexity of Selecting Maximal Solutions. *Information and Computation* 119, 231-239 (1995). (Cited on page 175.)

Julian Chingoma. 2023. Fairness and Stability in Complex Domains: Doctoral Consortium. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*. 7079-7080. (Cited on page 10.)

Julian Chingoma, Ulle Endriss and Ronald de Haan. 2022. Simulating Multiwinner Voting Rules in Judgment Aggregation. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 263-271. (Cited on pages 11, 12, 135, 136, 167, and 169.)

Julian Chingoma, Ulle Endriss, Ronald de Haan, Adrian Haret and Jan Maly. 2024. Apportionment with Weighted Seats. In *arXiv*. 2405.18102. (Cited on page 11.)

Julian Chingoma, Umberto Grandi and Arianna Novaro. 2024. Proportionality for Constrained Public Decisions. In *arXiv*. 2409.02609. (Cited on page 12.)

Vincent Conitzer, Rupert Freeman and Nisarg Shah. 2017. Fair Public Decision Making. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*. 629-646. (Cited on pages 32, 41, and 135.)

Péter Csóka and Jean-Jacques P. Herings. 2018. Decentralized Clearing in Financial Networks. *Management Science* 64, 4681-4699 (2018). (Cited on page 8.)

Adnan Darwiche and Pierre Marquis. 2002. A Knowledge Compilation Map. *Journal of Artificial Intelligence Research* 17, 229–264 (2002). (Cited on page 126.)

Argyrios Deligkas, Aris Filos-Ratsikas and Alexandros A. Voudouris. 2023. Heterogeneous Facility Location with Limited Resources. *Games and Economic Behavior* 139, 200-215 (2023). (Cited on pages 34 and 81.)

Franz Dietrich. 2006. Judgment Aggregation: (Im)possibility Theorems. *Journal of Economic Theory* 126, 286-298 (2006). (Cited on page 110.)

Franz Dietrich. 2014. Scoring Rules for Judgment Aggregation. *Social Choice and Welfare* 42, 873-911 (2014). (Cited on pages 110, 112, and 126.)

Franz Dietrich and Christian List. 2007. Arrow's Theorem in Judgement Aggregation. *Social Choice and Welfare* 29, 19-33 (2007). (Cited on pages 110 and 113.)

Virginie Do, Matthieu Hervouin, Jérôme Lang and Piotr Skowron. 2022. Online Approval Committee Elections. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*. 251-257. (Cited on page 34.)

Elad Dokow and Ron Holzman. 2010. Aggregation of Binary Evaluations. *Journal of Economic Theory* 145, 495-511 (2010). (Cited on page 110.)

Jon Doyle and Michael P. Wellman. 1991. Impediments to Universal Preference-based Default Theories. *Artificial Intelligence* 49, 97-128 (1991). (Cited on page 7.)

Conal Duddy. 2015. Fair sharing under dichotomous preferences. *Mathematical Social Sciences* 73, 1-5 (2015). (Cited on page 31.)

Conal Duddy, Ashley Piggins and William S. Zwicker. 2016. Aggregation of Binary Evaluations: A Borda-like approach. *Social Choice and Welfare* 46, 301-333 (2016). (Cited on page 166.)

Edith Elkind, Piotr Faliszewski, Piotr Skowron and Arkadii Slinko. 2017. Properties of Multiwinner Voting Rules. *Social Choice and Welfare* 48, 599-632 (2017). (Cited on pages 29, 120, and 126.)

Edith Elkind, Minming Li and Houyu Zhou. 2022. Facility Location with Approval Preferences: Strategyproofness and Fairness. In *Proceedings of the 21st International Conference on Autonomous Agents And Multiagent Systems (AAMAS)*. 391-399. (Cited on pages 34 and 81.)

Edith Elkind, Svetlana Obraztsova and Nicholas Teh. 2024. Verifying Proportionality in Temporal Voting. In *Proceedings of the 23rd International Conference on Autonomous Agents And Multiagent Systems (AAMAS)*. 2246-2248. (Cited on page 33.)

Ulle Endriss. 2016. Judgment Aggregation. In *Handbook of Computational Social Choice*. Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang and Ariel D. Procaccia (Eds). Cambridge University Press. Chapter 17, 399-424. (Cited on pages 107, 110, 113, and 135.)

Ulle Endriss. 2018. Judgment Aggregation with Rationality and Feasibility Constraints. In *Proceedings of the 17th International Conference on Autonomous*

Agents and Multiagent Systems (AAMAS). 946-954. (Cited on pages 107, 108, 109, 110, 111, 115, and 136.)

Ulle Endriss. 2020. Collective Information. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*. 13520-13524, Blue Sky Ideas Track. (Cited on page 7.)

Ulle Endriss, Ronald de Haan, Jérôme Lang and Marija Slavkovik. 2020. The Complexity Landscape of Outcome Determination in Judgment Aggregation. *Journal of Artificial Intelligence Research* 69, 687-731 (2020). (Cited on page 126.)

Brandon Fain, Ashish Goel and Kamesh Munagala. 2016. The Core of the Participatory Budgeting Problem. In *Proceedings of the 12th International Conference on Web and Internet Economics (WINE)*. 384-399. (Cited on page 30.)

Piotr Faliszewski, Piotr Skowron, Arkadii Slinko and Nimrod Talmon. 2017. Multiwinner Voting: A New Challenge for Social Choice Theory. *Trends in Computational Social Choice*. Ulle Endriss (Ed). AI Access. Chapter 2, 27-47. (Cited on pages 8, 23, 107, 109, 120, and 122.)

Piotr Faliszewski, Arkadii Slinko and Nimrod Talmon. 2020. Multiwinner Rules with Variable Number of Winners. In *Proceedings of the 24th European Conference on Artificial Intelligence (ECAI)*. 67-74. (Cited on pages 166, 169, and 170.)

Alireza Farhadi, Mohammad Ghodsi Mohammad Taghi Hajiaghayi, Sébastien Lahaie, David M. Pennock, Masoud Seddighin, Saeed Seddighin and Hadi Yami. 2019. Fair Allocation of Indivisible Goods to Asymmetric Agents. *Journal of Artificial Intelligence Research* 64, 1-20 (2019). (Cited on page 40.)

Michael F. Feldkamp. 2023. Datenhandbuch zur Geschichte des Deutschen Bundestages. (Deutscher Bundestag, 2023). Available at https://www.bundestag.de/dokumente/parlamentsarchiv/datenhandbuch_archiv/datenhandbuch_archiv-196998. Last visited on 2024-08-23. (Cited on pages 69 and 70.)

Peter C. Fishburn and Aleksandar Pekeč. 2004. Approval Voting for Committees: Threshold Approaches. *Technical Report* (2004). (Cited on page 170.)

Rupert Freeman, Anson Kahng, and David M. Pennock. 2020. Proportionality in Approval-Based Elections With a Variable Number of Winners. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*. 132-138. (Cited on pages 32, 133, 135, 139, 141, and 142.)

William V. Gehrlein. 1985. The Condorcet Criterion and Committee Selection. *Mathematical Social Sciences* 10, 199-209 (1985). (Cited on page 120.)

Andrey Goder and Vladimir Filkov. 2008. Consensus Clustering Algorithms: Comparison and Refinement. In *Proceedings of the 10th Workshop on Algorithm Engineering and Experiments (ALENEX)*. 109-117. (Cited on page 7.)

Umberto Grandi and Ulle Endriss. 2013. Lifting integrity constraints in binary aggregation. *Artificial Intelligence* 199-200, 45-66 (2013). (Cited on page 111.)

Davide Grossi and Gabriella Pigozzi. 2014. *Judgment Aggregation: A Primer*. Morgan & Claypool Publishers. (Cited on pages 107 and 110.)

Ronald de Haan. 2018. Hunting for Tractable Languages for Judgment Aggregation. In *Proceedings of the 16th International Conference on Principles of Knowledge Representation and Reasoning (KR)*. 194-203. (Cited on pages 108, 126, and 127.)

Daniel Halpern, Gregory Kehne, Ariel D. Procaccia, Jamie Tucker-Foltz, and Manuel Wüthrich. 2023. Representation with Incomplete Votes. In *Proceedings Of The 37th AAAI Conference On Artificial Intelligence (AAAI)*. 5657-5664. (Cited on page 35.)

Adrian Haret, Martin Lackner, Andreas Pfandler and Johannes Peter Wallner. 2020. Proportional Belief Merging. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*. 2822-2829. (Cited on pages 32 and 135.)

Paul Harrenstein, Marie-Louise Lackner and Martin Lackner. 2022. A Mathematical Analysis of an Election System Proposed by Gottlob Frege. *Erkenntnis* 87, 2609-2644 (2022). (Cited on page 28.)

Philemon Huising. 2023. Multiwinner Voting with Priority Candidates. *MSc Thesis*. University of Amsterdam. (Cited on page 36.)

Jonas Israel and Markus Brill. 2021. Dynamic Proportional Rankings. In *Proceedings Of The 30th International Joint Conference On Artificial Intelligence (IJCAI)*. 261-267. (Cited on page 34.)

Svante Janson. 2016. Phragmén's and Thiele's Election Methods. *arXiv*, 1611.08826. (Cited on page 23.)

Dimitris J. Kavvadias and Martha Sideril. 1998. The Inverse Satisfiability Problem. *SIAM Journal of Computing* 28, 152-163 (1998). (Cited on page 152.)

Marc D. Kilgour. 2016. Approval Elections with a Variable Number of Winners. *Theory and Decision* 81, 199-211 (2016). (Cited on page 166.)

Bettina Klaus, David F. Manlove and Francesca Rossi. 2016. Matching under Preferences. In *Handbook of Computational Social Choice*. Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang and Ariel D. Procaccia (Eds). Cambridge University Press. Chapter 14, 333-354. (Cited on page 7.)

Mark W. Krentel. 1988. The Complexity of Optimization Problems. *Journal of Computer and System Sciences* 36, 490-509 (1988). (Cited on page 175.)

Lewis A. Kornhauser and Lawrence G. Sager. 1993. The One and the Many: Adjudication in Collegial Courts. *California Law Review* 81, 1-59 (1993). (Cited on pages 7, 110, and 111.)

Martin Lackner. 2020. Perpetual Voting. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*. 2103-2110. (Cited on pages 32, 33, and 135.)

Martin Lackner and Jan Maly. 2021. Approval-Based Shortlisting. In *Proceedings of the 20th International Conference on Autonomous Agents And Multiagent Systems (AAMAS)*. 737-745. (Cited on pages 165 and 166.)

Martin Lackner and Jan Maly. 2021. Perpetual Voting: The Axiomatic Lens. In *Proceedings of the 8th International Workshop on Computational Social Choice (COMSOC)*. (Cited on page 33.)

Martin Lackner and Jan Maly. 2023. Proportional Decisions in Perpetual Voting. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*. 5722-5729. (Cited on page 33.)

Martin Lackner, Jan Maly and Oliviero Nardi. 2023. Free-Riding in Multi-Issue Decisions. In *Proceedings of the 22nd International Conference on Autonomous Agents And Multiagent Systems (AAMAS)*. 2040-2048. (Cited on page 33.)

Martin Lackner, Jan Maly and Simon Rey. 2021. Fairness in Long-Term Participatory Budgeting. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*. 299-305. (Cited on page 33.)

Martin Lackner and Piotr Skowron. 2023. *Multi-Winner Voting with Approval Preferences*. Springer. (Cited on pages 8, 13, 17, 19, 20, 21, 22, 23, 101, 103, and 164.)

Jérôme Lang and Piotr Skowron. 2018. Multi-Attribute Proportional Representation. *Artificial Intelligence* 263, 74-106 (2018). (Cited on pages 15 and 27.)

Jérôme Lang and Marija Slavkovik. 2013. Judgment Aggregation Rules and Voting Rules. In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT)*. 230-243. (Cited on page 110.)

Jérôme Lang and Lirong Xia. 2016. Voting in Combinatorial Domains. In *Handbook Of Computational Social Choice*. Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang and Ariel D. Procaccia (Eds). Cambridge University Press. Chapter 9, 197-221. (Cited on pages 133 and 135.)

Christian List and Phillip Pettit. 2002. Aggregating Sets of Judgments: An Impossibility Result. *Economics and Philosophy* 18, 89-110 (2002). (Cited on pages 107 and 110.)

Maaike Los, Zoé Christoff and Davide Grossi. 2022. Proportional Budget Allocations: Towards a Systematization. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*. 398-404. (Cited on page 30.)

Tyler Lu and Craig Boutilier. 2011. Budgeted Social Choice: From Consensus to Personalized Decision Making. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*. 280-286. (Cited on pages 8 and 174.)

Tyler Lu and Craig Boutilier. 2015. Value-Directed Compression of Large-Scale Assignment Problems. In *Proceedings of the 29th AAAI Conference On Artificial Intelligence (AAAI)*. 1182-1190. (Cited on page 8.)

Xinhang Lu, Jannik Peters, Haris Aziz, Xiaohui Bei and Warut Suksompong. 2023. Approval-Based Voting with Mixed Goods. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*. 5781-5788. (Cited on page 35.)

Tomás Masarík and Grzegorz Pierczynski and Piotr Skowron. 2023. Group Fairness in Social Choice. *arXiv*, 2307.06077. (Cited on pages 31, 32, 34, 41, 133, 135, 160, 166, and 167.)

Ivan-Aleksandar Mavrov, Kamesh Munagala and Yiheng Shen. 2023. Fair Multiwinner Elections with Allocation Constraints. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC)*. 964-990. (Cited on pages 31 and 32.)

Michael K. Miller and Daniel Osherson. 2009. Methods for Distance-Based Judgment Aggregation. *Social Choice And Welfare* 32, 575-601 (2009). (Cited on pages 110 and 112.)

Marcin Michorzewski, Dominik Peters and Piotr Skowron. 2020. Price of Fairness in Budget Division and Probabilistic Social Choice. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*. 2184-2191. (Cited on page 31.)

Klaus Nehring and Marcus Pivato. 2019. Majority Rule in the Absence of a Majority. *Journal of Economic Theory* 183, 213-257 (2019). (Cited on page 121.)

Klaus Nehring, Marcus Pivato and Clement Puppe. 2014. The Condorcet Set: Majority Voting over Interconnected Propositions. *Journal of Economic Theory* 151, 268-303 (2014). (Cited on page 112.)

Klaus Nehring and Clement Puppe. 2008. Consistent Judgement Aggregation: The Truth-functional Case. *Social Choice and Welfare* 31, 41-57 (2008). (Cited on page 110.)

Christos H. Papadimitriou. 1994. *Computational Complexity*. Addison-Wesley. (Cited on page 174.)

Marc Pauly and Martin van Hees. 2006. Logical Constraints on Judgement Aggregation. *Journal of Philosophical Logic* 35, 569-585 (2006). (Cited on page 110.)

Dominik Peters, Grzegorz Pierczynski, Nisarg Shah and Piotr Skowron. 2021a. Market-Based Explanations of Collective Decisions. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*. 5656-5663. (Cited on pages 20 and 21.)

Dominik Peters, Grzegorz Pierczynski and Piotr Skowron. 2021b. Proportional Participatory Budgeting with Additive Utilities. In *Proceedings of the 34th Annual Conference on Neural Information Processing Systems (NeurIPS)*. 12726-12737. (Cited on pages 19, 21, 24, 29, 30, 40, 46, and 81.)

Dominik Peters and Piotr Skowron. 2020. Proportionality and the Limits of Welfarism. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*. 793-794. (Cited on pages 19, 20, 21, 23, 24, 25, 86, 99, and 100.)

Ariel D. Procaccia. 2016. Cake Cutting Algorithms. In *Handbook of Computational Social Choice*. Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang and Ariel D. Procaccia (Eds). Cambridge University Press. Chapter 13, 311-328. (Cited on pages 8 and 34.)

Ariel D. Procaccia, Jeffrey S. Rosenschein and Aviv Zohar. 2008. On the Complexity of Achieving Proportional Representation. *Social Choice and Welfare* 30, 353-362 (2008). (Cited on page 174.)

Friedrich Pukelsheim. 2014. *Proportional Representation: Apportionment Methods and their Application*. Springer. (Cited on pages 8, 13, 16, 17, 39, 66, and 67.)

Thomas C. Ratliff. 2003. Some Startling Inconsistencies when Electing Committees. *Social Choice and Welfare* 21, 433-454 (2003). (Cited on pages 120 and 122.)

Simon Rey, Ulle Endriss, and Ronald de Haan. 2020. Designing Participatory Budgeting Mechanisms Grounded in Judgment Aggregation. In *Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning (KR)*. 692-702. (Cited on pages 126, 127, and 151.)

Simon Rey, Ulle Endriss, and Ronald de Haan. 2023. A General Framework for Participatory Budgeting with Additional Constraints. *Social Choice and Welfare* 1-37 (2023). (Cited on pages 126, 127, and 151.)

Simon Rey and Jan Maly. 2023. The (Computational) Social Choice Take on Indivisible Participatory Budgeting. *arXiv*, 2303.00621. (Cited on pages 7, 9, 29, and 81.)

Federica Ricca, Andrea Scozzari, and Paolo Serafini. 2017. A Guided Tour of the Mathematics of Seat Allocation and Political Districting. *Trends in Computational Social Choice*. Ulle Endriss (Ed). AI Access. Chapter 3, 49-68. (Cited on page 28.)

Luis Sánchez-Fernández, Edith Elkind, Martin Lackner, Norberto Fernández García, Jesús Arias-Fisteus, Pablo Basanta-Val and Piotr Skowron. 2017. Proportional Justified Representation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*. 670-676. (Cited on page 19.)

Luis Sánchez Fernández, Norberto Fernández García, Jesús A. Fisteus and Markus Brill. 2024. The Maximin Support Method: An Extension of the D'Hondt Method to Approval-based Multiwinner Elections. *Mathematical Programming* 203, 107-134 (2024). (Cited on page 103.)

Piotr Skowron. 2021. Proportionality Degree of Multiwinner Rules. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*. 820-840. (Cited on page 19.)

Piotr Skowron and Adrian Górecki. 2022. Proportional Public Decisions. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*. 5191-5198. (Cited on pages 32, 133, 135, and 154.)

Piotr Skowron, Martin Lackner, Markus Brill, Dominik Peters, and Edith Elkind. 2017. Proportional Rankings. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*. 409-415. (Cited on pages 8, 19, and 34.)

Piotr Skowron, Piotr Faliszewski and Jérôme Lang. 2016. Finding a Collective Set of Items: From Proportional Multirepresentation to Group Recommendation. *Artificial Intelligence* 241, 191-216 (2016). (Cited on page 8.)

Warut Suksompong. 2025. Weighted Fair Division of Indivisible Items: A Review. *Information Processing Letters* 187, 106519 (2025). (Cited on page 40.)

George Szpiro. 2010. *Numbers Rule: The Vexing Mathematics of Democracy, from Plato to the Present*. Princeton University Press. (Cited on page 66.)

Nimrod Talmon and Rutvik Page. 2021. Proportionality in Committee Selection with Negative Feelings. *arXiv*, 2101.01435. (Cited on page 34.)

Thorvald N. Thiele. 1895. Om flerfoldsvalg. *Oversigt Over Det Kongelige Danske Videnskabernes Selskabs Forhandlinger*. 414-441. (Cited on page 21.)

William Thomson. 2016. Introduction to the Theory of Fair Allocation. In *Handbook of Computational Social Choice*. Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang and Ariel D. Procaccia (Eds). Cambridge University Press. Chapter 11, 261-282. (Cited on pages 7 and 8.)

Klaus W. Wagner. 1990. Bounded Query Classes. *SIAM Journal on Computing* 19, 833-846 (1990). (Cited on page 175.)

Wikipedia. 2024. Entry "Bundestagsauschuss". (https://de.wikipedia.org/wiki/Bundestagsausschuss, 2024). Last visited on 2024-08-23. (Cited on page 69.)

William S. Zwicker. 2016. Introduction to the Theory of Voting. In *Handbook of Computational Social Choice*. Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang and Ariel D. Procaccia (Eds). Cambridge University Press. Chapter 2, 23-56. (Cited on pages 7 and 109.)

Index

A

Aggregation rule . . . 111
- Additive Majority Rule . . . 121
- Max-num . . . 112
- Max-sum . . . 112
- The u-max-sum rule . . . 124

Apportionment . . . 8
Apportionment method . . . 15
Approval ballot . . . 8

B

Binary election instance . . . 137

C

Candidate independence . . . 168
Cohesive group
- Constrained public decisions . . 141
- Constrained shortlisting . . . 168
- Multiwinner voting . . . 18
- Variant for weighted multiwinner voting . . . 93
- Weighted multiwinner voting . . . 85

Committee . . . 17
Committee monotonicity . . . 120
Constrained election instance . . . 167
Constrained shortlisting rule . . . 167
Constraint . . . 31
Core . . . 18

D

Decomposable negation normal form circuit . . . 126
Divisor method . . . 16
- Adams method . . . 16
- D'Hondt method . . . 16

E

Envy-freeness . . . 61

F

Feasibility constraint . . . 111
Full supply . . . 15

H

House monotonicity . . . 66

I

Incomparability . . . 109
Indifference . . . 109

J

Judgment aggregation . . . 107, 110
Justified representation . . . 18, 19
- Extended . . . 18
- Proportional . . . 19

L

Largest Remainder Method (LRM) . 17
Lower Quota (LQ) . . . 16

M

Multiwinner voting . . . 8
Multiwinner voting rule for approval ballots . . . 21
- Greedy Cohesive Rule (GCR) . . 21

Local-search Proportional Approval Voting 22
Method of Equal Shares (MES) 24
Phragmén's Sequential rule 23
Thiele method 21
Approval Voting 22
Chamberlin Courant for approval ballots 23
Proportional Approval Voting 22
Multiwinner voting rule for preorder preferences 109
k-Borda 109
k-Copeland 110
k-Copeland^0 110
k-Copeland^1 110
k-Kemeny 122

N
Net Approval Voting (NAV) rule .. 169
Number of External Defeats rule .. 122

O
Obtainable Weighted Lower Quota . 46
Obtainable Weighted Upper Quota . 58
Outcome determination problem .. 126

P
Participatory budgeting 9, 29
Party-list profile 17
Preference preorder 109
Priceability
Constrained public decisions .. 160
Multiwinner voting 20
Weighted multiwinner voting ... 95
Profile
Approval ballots 17
Constrained public decisions .. 136
Judgment aggregation 111
Preferences as preorders 109
Public decisions 32

R
Rationality constraint 111

S
Saint-Laguë method 16
Satisfaction 17
Seat assignment 15
Social choice theory 7

U
Upper Quota (LQ) 16

W
Weak Gehrlein stability 120
Weighted-seat apportionment method 42
Weighted divisor method 43
Adams$_\omega$ method 43
D'Hondt$_\omega$ method 43
Weighted-seat approval-based multiwinner voting rule 83
GreedyX$_\omega$ 92
MES$_\omega$ 84
seq-Phragmén$_\omega$ 83
Seat-based rule 83
Weight-based rule 83
Weighted GCR (GCR$_\omega$) 86
Welfarist rule 23

List of Symbols

General

A_i Approval ballot of voter i . . 17

c Candidate 17

C Set of candidates 17

i Voter 17

k Number of parliamentary or committee seats 15

N Set of voters 15

$N(c)$ Set of voters approving c . . 17

$\boldsymbol{v}$ Vector of the parties' vote counts 15

Apportionment

$g_p^\omega(t)$ Number of seats assigned to party p before round t of divisor method 16

M Apportionment method 15

p Party 15

$q(p)$ Quota of party p 15

$r(p, \boldsymbol{s})$ Party p's representation from seat assignment $\boldsymbol{s}$ 15

s Seat assignment 15

$(\boldsymbol{v}, k)$ Election instance 15

Multiwinner Voting

$\boldsymbol{A}$ Profile of voters' approval ballots 17

$(\boldsymbol{A}, k)$ Election instance 17

b Budget per voter (for priceability) 20

ps Price system 20

W Committee 17

Weighted Apportionment

$\delta(\boldsymbol{s})$ Average distance to the weighted quota 71

$\delta^-(\boldsymbol{s})$ Average distance below the weighted lower quota 71

$\delta^+(\boldsymbol{s})$ Average distance above the weighted upper quota 71

$g_p^\omega(t)$ Sum of weights assigned to party p before round t of weighted divisor method . . . 43

$l^{\#}(p)$ Lower quota of seats for party p . . . 46

$l^{o}(p)$ Obtainable lower quota of weights for party p . . . 46

M_ω Weighted-seat apportionment method . . . 42

$q_\omega(p)$ Weighted quota of party p . . . 42

$r_\omega(p, \boldsymbol{s})$ Weighted representation of party p from $\boldsymbol{s}$ . . . 42

$R(\boldsymbol{w})_{[h]}$ Set of all possible representation values . . . 42

$\boldsymbol{s}(p)$ Vector of seats (in increasing order) assigned to party p in seat assignment $\boldsymbol{s}$ . . . 42

$(\boldsymbol{v}, \boldsymbol{w})$ Weighted apportionment instance . . . 42

$\boldsymbol{w}$ Weight vector . . . 41

ω Total weight of weight vector $\boldsymbol{w}$ . . . 41

$u^{\#}(p)$ Upper quota of seats for party p . . . 56

$u^{o}(p)$ Obtainable upper quota of weights for party p . . . 57

Weighted Multiwinner Voting

$(\boldsymbol{A}, \boldsymbol{w})$ Weighted multiwinner election instance . . . 82

C_{ℓ_ω} Set of candidates approved by some overrepresented ℓ_ω-cohesive group . . . 89

F_ω Weighted-seat approval-based multiwinner voting rule . . . 83

low-$\boldsymbol{\mathcal{S}}(\boldsymbol{w})$ Vector of restricted, possible weighted satisfaction values . . . 93

$N_{>\ell_\omega}$ Set of overrepresented ℓ_ω-cohesive groups . . . 89

$sat_\omega(A, \boldsymbol{s})$ Weighted satisfaction from seat assignment $\boldsymbol{s}$ for candidate set A . . . 82

$\boldsymbol{s}(A)$ Vector of the positions within $\boldsymbol{s}$ of candidates . . . 86

$\mathcal{S}(\boldsymbol{w})$ Set of all possible weighted satisfaction values . . . 82

Simulating MWV Rules with JA

Φ Agenda of propositions (items) . . . 111

$\Phi^{C}_{\succcurlyeq}$ Preference agenda . . . 113

$\succcurlyeq_i$ Preference for voter i as a preorder . . . 109

$\text{Agr}(J, J')$ Set of agenda items that the judgments J and J' agree on . . . 111

$\text{Dis}(J, J')$ Set of agenda items that the judgments J and J' disagree on . . . 111

F Multiwinner voting rule for preorder preferences . . . 109

F_{JA} Aggregation rule . . . 111

J Judgment . . . 111

$\text{Maj}(\boldsymbol{J})$ Majority judgment . . . 111

$\text{Mod}(\Gamma)$ Set of all judgments that satisfy constraint Γ . . . 111

$n_{(\boldsymbol{J},\varphi)}$ Support of an issue φ in profile $\boldsymbol{J}$ . . . 111

$\textsc{Outcome}(F_{\mathrm{JA}})$ Outcome determination problem for F_{JA} . . . 126

$\boldsymbol{P}$ Preference profile 109

$(\boldsymbol{P}, k)$ Election instance 109

$\mathcal{R}(C)$ Set of all possible preorders over candidates 109

Constrained Public Decision

$\mathrm{Agr}(\boldsymbol{v}, \boldsymbol{v}')$ Agreement between vectors $\boldsymbol{v}$ and $\boldsymbol{v}'$ 137

$\boldsymbol{B}$ Ballot profile 136

$\boldsymbol{b}_i$ Ballot of voter i 136

$(\boldsymbol{B}, \mathcal{C}))$ Election instance 136

$\mathcal{C}$ Constraint 136

D_t Domain of issue t 136

$G_{\boldsymbol{w}}$ Outcome implication graph for outcome $\boldsymbol{w}$ 151

$\mathcal{I}$ Set of issues 136

$\mathcal{I}_{\mathrm{fix}}(C)$ Set of fixed decisions for set of outcomes C 143

$N(a_t, d)$ Set of voters that agree with decision d for issue t 137

$\boldsymbol{w}$ Outcome 136

Constrained Shortlisting

$(\boldsymbol{A}, \mathcal{C})$ Election instance 167

$\mathcal{C}$ Constraint 167

F_{sh} Shortlisting rule 167

$\boldsymbol{u}^m$ Scoring vector 169

Titles in the ILLC Dissertation Series:

ILLC DS-2018-08: **Srinivasan Arunachalam**
Quantum Algorithms and Learning Theory

ILLC DS-2018-09: **Hugo de Holanda Cunha Nobrega**
Games for functions: Baire classes, Weihrauch degrees, transfinite computations, and ranks

ILLC DS-2018-10: **Chenwei Shi**
Reason to Believe

ILLC DS-2018-11: **Malvin Gattinger**
New Directions in Model Checking Dynamic Epistemic Logic

ILLC DS-2018-12: **Julia Ilin**
Filtration Revisited: Lattices of Stable Non-Classical Logics

ILLC DS-2018-13: **Jeroen Zuiddam**
Algebraic complexity, asymptotic spectra and entanglement polytopes

ILLC DS-2019-01: **Carlos Vaquero**
What Makes A Performer Unique? Idiosyncrasies and commonalities in expressive music performance

ILLC DS-2019-02: **Jort Bergfeld**
Quantum logics for expressing and proving the correctness of quantum programs

ILLC DS-2019-03: **András Gilyén**
Quantum Singular Value Transformation & Its Algorithmic Applications

ILLC DS-2019-04: **Lorenzo Galeotti**
The theory of the generalised real numbers and other topics in logic

ILLC DS-2019-05: **Nadine Theiler**
Taking a unified perspective: Resolutions and highlighting in the semantics of attitudes and particles

ILLC DS-2019-06: **Peter T.S. van der Gulik**
Considerations in Evolutionary Biochemistry

ILLC DS-2019-07: **Frederik Möllerström Lauridsen**
Cuts and Completions: Algebraic aspects of structural proof theory

ILLC DS-2020-01: **Mostafa Dehghani**
Learning with Imperfect Supervision for Language Understanding

ILLC DS-2020-02: **Koen Groenland**
Quantum protocols for few-qubit devices

ILLC DS-2020-03: **Jouke Witteveen**
Parameterized Analysis of Complexity

ILLC DS-2020-04: **Joran van Apeldoorn**
A Quantum View on Convex Optimization

ILLC DS-2020-05: **Tom Bannink**
Quantum and stochastic processes

ILLC DS-2020-06: **Dieuwke Hupkes**
Hierarchy and interpretability in neural models of language processing

ILLC DS-2020-07: **Ana Lucia Vargas Sandoval**
On the Path to the Truth: Logical & Computational Aspects of Learning

ILLC DS-2020-08: **Philip Schulz**
Latent Variable Models for Machine Translation and How to Learn Them

ILLC DS-2020-09: **Jasmijn Bastings**
A Tale of Two Sequences: Interpretable and Linguistically-Informed Deep Learning for Natural Language Processing

ILLC DS-2020-10: **Arnold Kochari**
Perceiving and communicating magnitudes: Behavioral and electrophysiological studies

ILLC DS-2020-11: **Marco Del Tredici**
Linguistic Variation in Online Communities: A Computational Perspective

ILLC DS-2020-12: **Bastiaan van der Weij**
Experienced listeners: Modeling the influence of long-term musical exposure on rhythm perception

ILLC DS-2020-13: **Thom van Gessel**
Questions in Context

ILLC DS-2020-14: **Gianluca Grilletti**
Questions & Quantification: A study of first order inquisitive logic

ILLC DS-2020-15: **Tom Schoonen**
Tales of Similarity and Imagination. A modest epistemology of possibility

ILLC DS-2020-16: **Ilaria Canavotto**
Where Responsibility Takes You: Logics of Agency, Counterfactuals and Norms

ILLC DS-2020-17: **Francesca Zaffora Blando**
Patterns and Probabilities: A Study in Algorithmic Randomness and Computable Learning

ILLC DS-2021-01: **Yfke Dulek**
Delegated and Distributed Quantum Computation

ILLC DS-2021-02: **Elbert J. Booij**
The Things Before Us: On What it Is to Be an Object

ILLC DS-2021-03: **Seyyed Hadi Hashemi**
Modeling Users Interacting with Smart Devices

ILLC DS-2021-04: **Sophie Arnoult**
Adjunction in Hierarchical Phrase-Based Translation

ILLC DS-2021-05: **Cian Guilfoyle Chartier**
A Pragmatic Defense of Logical Pluralism

ILLC DS-2021-06: **Zoi Terzopoulou**
Collective Decisions with Incomplete Individual Opinions

ILLC DS-2021-07: **Anthia Solaki**
Logical Models for Bounded Reasoners

ILLC DS-2021-08: **Michael Sejr Schlichtkrull**
Incorporating Structure into Neural Models for Language Processing

ILLC DS-2021-09: **Taichi Uemura**
Abstract and Concrete Type Theories

ILLC DS-2021-10: **Levin Hornischer**
Dynamical Systems via Domains: Toward a Unified Foundation of Symbolic and Non-symbolic Computation

ILLC DS-2021-11: **Sirin Botan**
Strategyproof Social Choice for Restricted Domains

ILLC DS-2021-12: **Michael Cohen**
Dynamic Introspection

ILLC DS-2021-13: **Dazhu Li**
Formal Threads in the Social Fabric: Studies in the Logical Dynamics of Multi-Agent Interaction

ILLC DS-2022-01: **Anna Bellomo**
Sums, Numbers and Infinity: Collections in Bolzano's Mathematics and Philosophy

ILLC DS-2022-02: **Jan Czajkowski**
Post-Quantum Security of Hash Functions

ILLC DS-2022-03: **Sonia Ramotowska**
Quantifying quantifier representations: Experimental studies, computational modeling, and individual differences

ILLC DS-2022-04: **Ruben Brokkelkamp**
How Close Does It Get?: From Near-Optimal Network Algorithms to Suboptimal Equilibrium Outcomes

ILLC DS-2022-05: **Lwenn Bussière-Carae**
No means No! Speech Acts in Conflict

ILLC DS-2023-01: **Subhasree Patro**
Quantum Fine-Grained Complexity

ILLC DS-2023-02: **Arjan Cornelissen**
Quantum multivariate estimation and span program algorithms

ILLC DS-2023-03: **Robert Paßmann**
Logical Structure of Constructive Set Theories

ILLC DS-2023-04: **Samira Abnar**
Inductive Biases for Learning Natural Language

ILLC DS-2023-05: **Dean McHugh**
Causation and Modality: Models and Meanings

ILLC DS-2023-06: **Jialiang Yan**
Monotonicity in Intensional Contexts: Weakening and: Pragmatic Effects under Modals and Attitudes

ILLC DS-2023-07: **Yiyan Wang**
Collective Agency: From Philosophical and Logical Perspectives

ILLC DS-2023-08: **Lei Li**
Games, Boards and Play: A Logical Perspective

ILLC DS-2023-09: **Simon Rey**
Variations on Participatory Budgeting

ILLC DS-2023-10: **Mario Giulianelli**
Neural Models of Language Use: Studies of Language Comprehension and Production in Context

ILLC DS-2023-11: **Guillermo Menéndez Turata**
Cyclic Proof Systems for Modal Fixpoint Logics

ILLC DS-2023-12: **Ned J.H. Wontner**
Views From a Peak: Generalisations and Descriptive Set Theory

ILLC DS-2024-01: **Jan Rooduijn**
Fragments and Frame Classes: Towards a Uniform Proof Theory for Modal Fixed Point Logics

ILLC DS-2024-02: **Bas Cornelissen**
Measuring musics: Notes on modes, motifs, and melodies

ILLC DS-2024-03: **Nicola De Cao**
Entity Centric Neural Models for Natural Language Processing

ILLC DS-2024-04: **Ece Takmaz**
Visual and Linguistic Processes in Deep Neural Networks: A Cognitive Perspective

ILLC DS-2024-05: **Fatemeh Seifan**
Coalgebraic fixpoint logic Expressivity and completeness result

ILLC DS-2024-06: **Jana Sotáková**
Isogenies and Cryptography

ILLC DS-2024-07: **Marco Degano**
Indefinites and their values

ILLC DS-2024-08: **Philip Verduyn Lunel**
Quantum Position Verification: Loss-tolerant Protocols and Fundamental Limits

ILLC DS-2024-09: **Rene Allerstorfer**
Position-based Quantum Cryptography: From Theory towards Practice

ILLC DS-2024-10: **Willem Feijen**
Fast, Right, or Best? Algorithms for Practical Optimization Problems

ILLC DS-2024-11: **Daira Pinto Prieto**
Combining Uncertain Evidence: Logic and Complexity