



VARIATIONS ON PARTICIPATORY BUDGETING



VARIATIONS ON PARTICIPATORY BUDGETING **SIMON REY**



	Handwritten text: "Handwritten text"	Progress bar: 100% green	Line graph: A single peak	Bar chart: Three bars of increasing height
	Handwritten text: "Handwritten text"	Progress bar: 75% green, 25% yellow	Line graph: A jagged line	Bar chart: Three bars of varying heights
	Handwritten text: "Handwritten text"	Progress bar: 50% green, 50% red	Line graph: A smooth curve	Bar chart: Three bars of varying heights
	Handwritten text: "Handwritten text"	Progress bar: 25% green, 75% red	Line graph: A smooth curve	Bar chart: Three bars of varying heights
	Handwritten text: "Handwritten text"	Progress bar: 100% yellow	Line graph: A smooth curve	Bar chart: Three bars of varying heights
	Handwritten text: "Handwritten text"	Progress bar: 75% green, 25% red	Line graph: A smooth curve	Bar chart: Three bars of varying heights
	Handwritten text: "Handwritten text"	Progress bar: 100% yellow	Line graph: A smooth curve	Bar chart: Three bars of varying heights
	Handwritten text: "Handwritten text"	Progress bar: 50% green, 50% red	Line graph: A smooth curve	Bar chart: Three bars of varying heights
	Handwritten text: "Handwritten text"	Progress bar: 25% green, 75% red	Line graph: A smooth curve	Bar chart: Three bars of varying heights



INSTITUTE FOR LOGIC, LANGUAGE AND COMPUTATION

Variations on Participatory Budgeting

Simon Rey

Variations on Participatory Budgeting

ILLC Dissertation Series DS-2023-09



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 900
1090 GE Amsterdam
phone: +31-20-525 6051
e-mail: illc@uva.nl
homepage: 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).

Copyright © 2023 by Simon Rey

Cover design by Simon Rey based on an original idea of Marianne de Heer Kloots.
Printed and bound by Ipskamp Printing.

ISBN: 978-94-6473-211-5

Variations on Participatory Budgeting

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 vrijdag 13 oktober 2023, te 10.00 uur

door

Simon Jean Rey

geboren te Caen

Promotiecommissie

Promotor:	prof. dr. U. Endriss	<i>Universiteit van Amsterdam</i>
Copromotores:	dr. R. de Haan	<i>Universiteit van Amsterdam</i>
	dr. J.F. Maly	<i>Universiteit van Amsterdam</i>
Overige leden:	prof. dr. E. Elkind	<i>University of Oxford</i>
	prof. dr. S. Ghebreab	<i>Universiteit van Amsterdam</i>
	dr. D. Grossi	<i>Universiteit van Amsterdam</i>
	prof. dr. J. Lang	<i>CNRS & Université Paris Dauphine</i>
	dr. rer. nat. R.E.M. Reiffenhäuser	<i>Universiteit van Amsterdam</i>
	prof. dr. G. Schäfer	<i>Universiteit van Amsterdam</i>
	dr. P.M. Skowron	<i>Uniwersytet Warszawski</i>

Faculteit der Natuurwetenschappen, Wiskunde en Informatica

Samenvatting

Variaties op Burgerbegroting

Door de waarneembare afname van het vertrouwen in regeringen is het idee van een *crisis van de democratie* steeds meer aanwezig. In de literatuur over politicologie wordt nog voortdurend gedebatteerd over het optimale niveau van vertrouwen voor een goed functionerend democratisch proces, maar er bestaat in ieder geval een duidelijke consensus dat te weinig vertrouwen het proces in gevaar kan brengen. Het is dan ook niet verwonderlijk dat er in de laatste tijd veel innovatieve instrumenten zijn ontwikkeld om het democratisch proces te vernieuwen. Deze proefschrift gaat over een van die instrumenten, namelijk de burgerbegroting (BB), een reeks mechanismen bedoeld om collectieve en participatieve budgetteringsbeslissingen te nemen.

In dit proefschrift onderzoeken we BB-mechanismen als manieren om een collectieve budgetteringsbeslissing te krijgen. In het bijzonder onderzoeken we BB als een stemprocedure waarbij burgers worden gevraagd hun voorkeuren in te dienen om te beslissen welke projecten moeten worden gefinancierd, met inachtneming van een budgetbeperking. Ons onderzoek vindt zijn oorsprong in de literatuur over computationele socialekeuzetheorie, het onderzoeksgebied dat manieren bestudeert om vanuit individuele voorkeuren tot collectieve beslissingen te komen. Met de standaard gereedschapskist proberen we te begrijpen hoe individuen meningen kunnen worden samengevoegd tot een collectieve beslissing in verschillende BB contexten.

Gezien de talloze implementaties van BB is een holistische benadering een uitdaging. Ons onderzoek is gestructureerd langs twee assen, die elk een deel van het proefschrift bepalen en die nieuwe aspecten van BB in de analyse inbrengen.

Het eerste deel van het proefschrift is gewijd aan het zogenaamde *standaardmodel van BB*, de meest voorkomende wiskundige formalisering van BB processen in de literatuur. We hanteren twee nieuwe perspectieven om dit model te onderzoeken.

De studie van het standaardmodel in de literatuur houdt zich bijna uitsluitend bezig met de vraag hoe een rechtvaardige uitkomst kan worden verkregen. Meestal zijn beperkende hypothesen vereist, waarbij ofwel onredelijke veronderstellingen

worden gemaakt over het gedrag van de kiezers, en/of wordt geëist dat de kiezer onrealistische hoeveelheden informatie geeft. In het eerste hoofdstuk van dit deel van het proefschrift vechten we deze benaderingen aan en opperen we een nieuwe kijk op rechtvaardigheid voor BB die niet aan deze nadelen lijdt. We bestuderen een groot aantal nieuwe rechtvaardigheidsconcepten, bespreken hoe ze kunnen worden geïmplementeerd, en tonen de levensvatbaarheid van onze nieuwe benadering aan.

Hoewel het standaardmodel een groot aantal reële BB-processen omvat, beperkt het feit dat rechtvaardigheid het meest bestudeerde onderwerp is de toepasbaarheid ervan. Rechtvaardigheid is immers niet het hoofddoel van alle BB-processen; sommige processen zijn bijvoorbeeld georganiseerd om te ontdekken wat de beste alternatieven zijn. Op basis van deze vaststelling presenteert het tweede hoofdstuk van dit deel de eerste studie van het standaardmodel vanuit een epistemisch perspectief. Hier is het doel om enkele grondwaarheden over de intrinsieke kwaliteiten van de projecten bloot te leggen. We onderzoeken de epistemische vermogens van vele BB-regels, en tonen aan dat de meeste er eigenlijk geen hebben.

In het tweede deel gaan we over op de studie van BB-modellen die de standaardmodel uitbreiden om meer aspecten van BB-processen vast te leggen.

In het eerste hoofdstuk onderzoeken we multi-beperkte BB-modellen waarbij de budgetlimiet niet de enige beperking is die de haalbaarheid van een uitkomst bepaalt. We presenteren een algemene aanpak voor een dergelijke uitbreiding van het standaardmodel, waarbij we in detail aangeven hoe budgettoewijzingen moeten worden bepaald wanneer de structuur van de uitkomst complexer is.

Het tweede hoofdstuk behandelt het tijdsaspect van BB. BB-processen worden immers uitgevoerd over meerdere jaren, waarbij er elk jaar één verkiezing wordt gehouden. We nemen dit aspect op in de formele analyse, waarbij de nadruk ligt op rechtvaardigheid over de tijd. We introduceren noties van temporele rechtvaardigheid en onderzoeken onder welke voorwaarden deze kunnen worden toegepast.

Een ander belangrijk aspect van BB-processen dat nog niet in de analyse was opgenomen, is het feit dat de projecten in BB-processen door de burgers zelf worden voorgesteld. In het laatste hoofdstuk van dit deel breiden we het standaardmodel uit door een voorfase op te nemen waarin projectvoorstellen worden ingediend en vervolgens worden geselecteerd om de reeks projecten te vormen die in stemming worden gebracht. We richten ons eerst op deze eerste fase, waarbij we verschillende methoden voor het bepalen van de shortlist onderzoeken, en gaan vervolgens over tot het bestuderen van de interacties tussen de twee fasen.

In het algemeen bestudeert dit proefschrift procedures voor het selecteren van projecten die in BB-scenario's worden gefinancierd. We hebben ernaar gestreefd om een breder scala aan feitelijke implementaties van BB-processen in de formele analyse op te nemen dan tot nu toe was gedaan. Het eindproduct omvat een grote verscheidenheid aan toepassingen van BB-processen, bestudeerd vanuit verschillende invalshoeken. Ik hoop dat dit werk, op kleine schaal, kan bijdragen tot betere beslissingen voor BB, en daardoor de democratie op grotere schaal kan helpen verbeteren.

Summary

Variations on Participatory Budgeting

The third quarter of the 20th century has seen the rise of the idea of a *crisis of democracy*, both inside and outside of the academic literature. This crisis is often linked to the observation that confidence and trust in governments and governmental bodies is declining in many countries. There is a—still on-going—debate in the political science literature regarding the optimal level of trust for a well-functioning democratic process, but there is clear consensus that too little trust can endanger it. It is thus not surprising that the past 40 to 50 years have also seen the rise of a wide range of innovative tools developed to deepen and renew the democratic process. One such tool is *participatory budgeting* (PB), which encompasses a large range of mechanisms that aim to make budgeting decisions in a participatory and collective manner.

This thesis studies PB mechanisms. We view them as ways of obtaining a collective budgeting decision. More specifically, we investigate PB as a voting procedure in which citizens are asked to submit their preferences in order to decide which projects should be funded, subject to a budget constraint. Our investigation has its roots in the literature on *computational social choice*, the field of research that studies ways of reaching collective decisions from individual preferences. Equipped with the standard toolbox of the computational social choice scientist, we aim at understanding how to aggregate individual opinions into a collective decision in a wide variety of PB contexts.

There exists a myriad of different implementations of PB in real life. This makes taking a holistic approach to PB particularly challenging. Our investigation is structured along two axes, each defining a part of the thesis, each bringing new aspects of real-world PB processes into the analysis.

The first part of the thesis is dedicated to the so-called *standard model of PB*, *i.e.*, the most frequently encountered mathematical formalisation of PB processes in the literature. We consider two new perspectives on PB to investigate it.

The study of the standard model in the literature is primarily concerned with the

question of how to obtain a fair outcome. However, restrictive hypotheses are usually required, either making unreasonable assumptions about the voters' behaviours, and/or requiring the voters to provide unrealistic amounts of information. In the first chapter of this part of the thesis, we challenge these approaches and propose a new take on fairness for PB that does not suffer from these drawbacks. We study a large range of new fairness concepts, discuss how they can be enforced, and demonstrate the viability of our new approach.

The standard model captures a large and diverse set of real-life PB processes. However, focusing on fairness restricts the applicability of the formal analysis. Indeed, fairness is not the main objective of all PB processes; some processes are for instance organised to discover what the best alternatives are. Motivated by this observation, the second chapter of this part presents the first study of the standard model from an epistemic perspective. Here, the goal is to uncover some ground truths about the intrinsic qualities of the projects. We investigate the epistemic abilities of many PB rules, and show that most actually do not enjoy any.

In the second part of the thesis, we move to the study of PB models that extend the standard one, in order to capture additional aspects of real-life PB processes.

In the first chapter we investigate multi-constrained PB models. More specifically, we study models of PB in which the budget limit is not the only constraint that determines the feasibility of an outcome. Additional constraints can be used to model statements such as "at least € 10 000 have to be allocated to cycling infrastructure". In this chapter we present a general approach for such extension of the standard model, detailing how to determine budget allocations when the structure of the outcome is more complex.

The second chapter of this part tackles the temporal aspect of PB. Indeed, most PB processes are implemented over the course of several years, one election being ran each year. We incorporate this aspect into the formal analysis, and focus on providing fairness over time. We introduce several notions of temporal fairness and present conditions under which they can be enforced.

Another important aspect of PB processes that had not been incorporated in the analysis before is the fact that projects are not just voted on but also proposed by the citizens. In the last chapter of this part, we extend the standard model by including a preliminary stage in which project proposals are submitted and then shortlisted to form the set of projects that are brought to the vote. We first focus on the first stage, investigating different methods of determining the shortlist, and then move on to the study of the interactions between the two stages.

Overall, this thesis is concerned with procedures to select the projects to be funded in PB scenarios. Throughout our analysis, we aim at incorporating a wider range of actual implementations of PB processes into the formal analysis than had been done before. The end product covers a large diversity of implementations of PB processes, studied from various angles. I hope that this work, at its small scale, can help make better decisions for PB, and by that, improve the democratic process at a larger scale.

Contents

Acknowledgments	xiii
------------------------	-------------

Initial Words

1 Introduction	3
1.1 The Object: Participatory Budgeting	5
1.1.1 Definition	5
1.1.2 Implementation	6
1.2 The Question: Aggregating Preferences for Participatory Budgeting . . .	7
1.3 The Method: (Computational) Social Choice	8
1.4 The Contribution: Variations on Participatory Budgeting	9
1.4.1 Variations on the Method	10
1.4.2 Variations on the Model	11
1.4.3 Beyond the Technicalities	12

Part One: General Background

2 The Standard Model of Participatory Budgeting	15
2.1 Setting up the Voting Stage	15
2.2 Collecting the Votes	16
2.3 Determining Budget Allocations	17
2.3.1 Utilitarian Welfare Maximising Rules	18
2.3.2 The Sequential Phragmén Rule	20
2.3.3 The Methods of Equal Shares	21
2.4 Discussing Computational Complexity	22

3	The (Computational) Social Choice Take on Participatory Budgeting	23
3.1	Ballot Design	23
3.1.1	Terminology around the Voters	24
3.1.2	Cardinal Ballots	26
3.1.3	Ordinal Ballots	28
3.1.4	Comparison of Ballot Formats	29
3.1.5	Ballot-Based Satisfaction	31
3.2	Participatory Budgeting Rules	32
3.2.1	Welfare-Maximising Rules	32
3.2.2	The Sequential Phragmén Rule	34
3.2.3	The Maximin Support Rule	34
3.2.4	The Method of Equal Shares	35
3.2.5	Other Rules for Participatory Budgeting	36
3.3	Fairness in Indivisible Participatory Budgeting	36
3.3.1	Extended and Proportional Justified Representation	37
3.3.2	The Core	46
3.3.3	Priceability	49
3.3.4	Proportionality in Ordinal PB	51
3.3.5	Other Fairness Requirements	52
3.3.6	Fairness in Extended Settings	55
3.3.7	Taxonomies of Proportionality in PB	55
3.4	Axiomatic Analysis	59
3.4.1	Exhaustiveness	59
3.4.2	Monotonicity Requirements	60
3.4.3	Strategy-Proofness	62
3.4.4	Other Axioms	65
3.5	Algorithmic Considerations	65
3.5.1	Outcome Determination of Standard PB Rules	66
3.5.2	Maximising Social Welfare	66
3.5.3	Other Algorithmic Problems	68
3.6	Variations and Extensions of the Standard Model	68
3.6.1	Local versus Global Processes	70
3.6.2	Additional Distributional Constraints	70
3.6.3	Interaction Between Projects	71
3.6.4	Enriched Cost Functions	72
3.6.5	Uncertainty in PB	72
3.6.6	PB with Endogenous Funding	73
3.6.7	Weighted PB	73
3.7	Related Frameworks	73

Part Two: Variations on the Method

4	Defining Fairness via Equality of Resources	77
4.1	The Share	79

4.2	The <i>Fair</i> Share	80
4.2.1	Exact Fair Share	81
4.2.2	Fair Share up to One Project	83
4.2.3	Local Fair Share	87
4.3	The <i>Justified</i> Share	88
4.3.1	Achieving EJS	90
4.3.2	Relaxing EJS	95
4.4	Relationships between Criteria	100
4.4.1	Share-Based Fairness Criteria	100
4.4.2	Comparison with Priceability	105
4.4.3	Representation-Based Fairness Criteria	108
4.4.4	Share and Efficiency Requirements	110
4.5	Approaching Fair Share in Practice	111
4.5.1	Optimal Distance to Fair Share	111
4.5.2	Distance to Fair Share of Common PB Rules	116
4.6	Summary	118
5	Viewing Participatory Budgeting Rules through the Epistemic Lens	119
5.1	The Truth-Tracking Perspective	122
5.2	Proportional PB Rules	123
5.3	Monotonic Argmax Rules	125
5.3.1	Nash Social Welfare	126
5.3.2	Utilitarian Social Welfare	132
5.4	Summary	135

Part Three: Variations on the Model

6	A General Framework for Multi-Constraint Participatory Budgeting	139
6.1	Frameworks	141
6.1.1	Participatory Budgeting with Multiple Resources	142
6.1.2	Judgment Aggregation	142
6.1.3	Embedding PB into Judgement Aggregation	144
6.2	Efficient Embeddings of Participatory Budgeting	146
6.2.1	Tractable Language for Judgment Aggregation	146
6.2.2	DNNF Circuit Embeddings	148
6.2.3	Participatory Budgeting with Project Dependencies	152
6.2.4	Participatory Budgeting with Quotas on Types of Projects	159
6.3	Enforcing Exhaustiveness	165
6.3.1	Exhaustive Embeddings for Single-Resource Instances	166
6.3.2	Asymmetric Judgment Aggregation Rules	167
6.4	Axiomatic Analysis of Judgment Aggregation Rules	168
6.5	Summary	174

7	A Long-Term Approach to Participatory Budgeting	175
7.1	Perpetual Participatory Budgeting	177
7.2	A Fairness Theory for Perpetual Participatory Budgeting	180
7.2.1	Evaluation Functions	180
7.2.2	Fairness Criteria	183
7.3	Achieving Perfect Fairness: EQUAL- Φ	187
7.4	Optimising for Fairness: Φ -GINI	191
7.4.1	Among all Solutions	191
7.4.2	Among Exhaustive Solutions	195
7.5	Converging Towards Fairness: EQUAL- Φ -CONV	201
7.5.1	For the Cost Evaluation Function Φ^{cost}	201
7.5.2	For the Share Evaluation Function Φ^{share}	207
7.5.3	For the Relative Cost Evaluation Function $\Phi^{relcost}$	208
7.6	Summary	214
8	An End-to-End Model for Participatory Budgeting	217
8.1	The Formal End-to-End Model	221
8.1.1	Additional Notation and Terminology	221
8.1.2	The Shortlisting Stage	222
8.1.3	The Allocation Stage	222
8.1.4	Agent Preferences	222
8.2	Shortlisting Rules	224
8.2.1	The Equal Representation Shortlisting Rule	224
8.2.2	Median-Based Shortlisting Rules	225
8.3	End-to-End Example	226
8.4	Axioms for Shortlisting Rules	227
8.5	First-Stage Strategy-Proofness	230
8.5.1	Awareness-Restricted Manipulation	233
8.5.2	Unrestricted Manipulation	235
8.6	Summary	239

Final Words

9	Conclusion	245
9.1	Closing out the Thesis	245
9.2	Opening up New Perspectives	247
	Bibliography	249
	Index	267
	List of Symbols	271

Acknowledgments

I think every PhD student dreams of writing their acknowledgments—a manifestation that the work has been done—it was at least my case for the past four years. I am absolutely thrilled to have reached this moment.

Let me take a chronological approach here. Right after high-school I started studying economics and management science, excited by the prospect of understanding what is driving people's decisions. Three years later, I discovered that sophisticated techniques from theoretical computer science could be used to help the decision process, or even, to be the decision process. This revelation came through the course on operational research taught by Renaud Lacour. That was it for me as an economist, I would start again as a computer scientist.

From my economics days, I carried forward my interest for understanding the rational behind decision-making. So, when at the end of my first year of computer science, I saw an announcement for an internship in computational social choice (COMSOC)—and area of research at the intersection between economics and computer science—in the region I grew up, I could not let the opportunity pass, and I did not. Thank you, Umberto, for introducing me to the field of COMSOC, and for sharing your love of research with me. I still have fond memories of the six weeks we worked together. From that point on, I have always been looking for your friendly face in every event I attend. Getting back to our little story, all the elements were starting to fall in place: I discovered COMSOC through Umberto—Ulle's former student—by extending some work that Sirin did, while sharing an office with Arianna. The importance of all these names will be clarified soon.

During my Master's years I started working with Aurélie and Nicolas on different projects related to COMSOC. It all started as my first Master's thesis, and continued far beyond. I learnt so much from the two of you—not only about the technicalities of fair division, but also about academic research as a whole. Most importantly, you showed me how fun it is to work when you have fantastic collaborators. I am more than honoured to have had such academic mentors, thank you for everything.

The summer of my first year of Master's brought me to the Netherlands for the

first time. I visited Davide for three months in Groningen, exploring in depth whether my newly acquired taste for modal logic would survive further examination (spoiler alert, it did not). I will always remember what an amazing host you were, Davide, even though we had never interacted before. The work we did together eventually led to my first publication as the main contributor, thank you for making it happen. Thanks to you I also attended many relevant events throughout the Netherlands that summer, which, in turn, opened so many doors: Our first encounter took place at the ILLC where you were attending a workshop; On the occasion of a Dutch social choice colloquium for which Nicolas came to Amsterdam, I met Sirin, Ulle and Zoi for the first time; A later-aborted collaboration started with Marja, Sirin and Zoi during the workshop you (Davide) organised at the Lorentz center. I am deeply grateful for all you did for my integration into the COMSOC network.

When you work on fair division, and in COMSOC in general, there are few names that pop up all the time. One of them is Haris', and that pushed me to travel almost 17 000 km to do my (second) Master's thesis with him in Sydney. It was great fun working with you, Haris. Thank you for all the work you put in, but also, and mainly, for all the non-work related events you invited me to. From Sydney I still have vivid memories of almost getting lost in the bush (on my own); celebrating Serge's wedding; discovering bubble tea with Abdallah; (almost) petting a wombat with Ágnes; making crêpes with Christine, Haris and Sana; sharing a desk with Ayda; and so much more.

It is also during my Australian stay that I started working on a new version of PrefLib. Without prior notice, I eventually sent Nick a video of the new website I had designed. Fortunately Nick responded very positively to that email and we moved forward with the whole plan. I am very proud to be the maintainer of the website. I always enjoyed our calls, Nick, though I often felt we spend more time laughing than working... Thanks for taking me onboard, and thanks for not being upset by my arrogant "I did not like your design so I re-designed the website from scratch".

The most important thing that happened while I was upside down—at least for the present thesis—is that I got offered a position to do a PhD with Ulle. It had all started after that Dutch social choice colloquium I mentioned earlier when I very bluntly asked Ulle "What should I do to do a PhD with you?", the very first time we saw each other. One and a half years later I started working with Ulle, which eventually led to the writing up of this thesis. Working with you, Ulle, has been a real pleasure and I am grateful beyond words to have been your student for the past four years. You managed to find the very sweet spot between giving me the freedom I wanted while still steering me. Working closely with you taught me so much. You showed me countless new places where I could practice my predisposition for rigorous work (which one could call nitpicking). My bibliographies will never be the same, and I am delighted about it. Your urge to always produce high quality work contributed to many of my, and our, successes, and drastically impacted my way of working. Beyond work, I am also very grateful to the person you are. We definitely had a lot of fun together. I most definitely enjoyed a lot of our chats, laughs, gossips, and deeper conversations. Many thanks, Ulle, for these four years, and the one(s) to come.

I am also indebted to my co-supervisors Ronald and Jan. Ronald, you brought to me your enchantment for complexity theory and your constant positivity. You also trusted me for other aspects of the job. I was very happy to be your TA and to co-supervise Giovanni's thesis with you. Thank you for all of that. Jan, working with you has been a real privilege. I do believe that we make an excellent pair. Our collaboration is an absolute highlight of my PhD. You have all my admiration. And do not worry, you will receive further attention later regarding the other aspects of our friendship.

I also want to thank all the other collaborators I interacted with during my PhD. I was more than happy to collaborate with you, Martin, strengthening the link between Amsterdam and Vienna. I really benefited from your deep understanding of the topics we worked on, but also from your friendliness. Dorothea, Linus and Christian, many thanks go to you too. You were the perfect hosts and my visit to Düsseldorf really re-ignited my taste for research. Zoi and Sirin, I am so happy to count you as my collaborators. You already show up in so many parts of my life that it only made sense to also have you here as well.

In addition to the people I worked closely with, the community as a whole helped me enjoy my PhD. Seeing you at every event I attended was a great pleasure, Edith. I always looked forward to chatting with you. I also want to thank you for letting me discover the intriguing world of Oxford, it was delightful. Piotr (F.), you also attended more or less the same set of events as me, and I always discovered your participation with excitement. I cannot believe we never made a map together given how much training on the topic you provided me with. I also enjoyed discussing PB, PrefLib, or just the academic life in general with you, Dominik. Thank you for that and sorry for the misaddressed texts you received. I also have a friendly thought for you, Jérôme. As we discovered, you were here from the very beginning. For a wide variety of work, social, and/or personal interactions I also want to thank Chris, Guido, Jannik, Maaike, Niclas, Oliviero, Piotr (S.), and Rachael.

Finally, let me thank Edith, Sennay, Davide, Jérôme, Rebecca, Guido and Piotr (S.) for accepting to be part of my committee. Some of these names already appeared above, which makes me particularly happy. I am looking forward to discussing this thesis with all of you.

I think it is more than fair to say that the friends I made along the way may be the most important people. Let me turn to them now.

First, there are the ones that have been here forever and never cease to be around. Marion, we go way back together and I hope that our friendship will persist throughout the years. Every day I received one of your voice message was a great day. Wherever you are, know that you can always call me for a bowl of soup or just some comforting chat if you get sick. Camille, you are one of the most important persons in my life. I grew so much by sharing your life that I will be forever indebted to you. Thank you for being there then, and still now. Emma, living with you was a lot of fun and the beginning of a great friendship. Please keep on calling at random times

of the day, I enjoy our chats too much for them to stop.

Then, there are the friends from Amsterdam. Dean, you have been the first, and what a first! I only know Amsterdam with you in it (well, ok, Diemen), and even if I wish for you to find the position you dream of, I selfishly hope that you will stay in Amsterdam; I value our friendship way too much. Thank you for everything, all the fun, the bike rides, the concerts, the cooking, the dancing, the tea breaks, the robots... Call me if you ever need to move houses!

The jinx (not-all) believers deserve a whole paragraph for sure. You all have always been there, and especially when I needed you the most! Zoi, thank you for all the time we shared, in or out of Amsterdam. You have been a close friend, a confidant, and an incredible support. I would always pick you to either share a cheese plater, relax on some Greek islands, or enjoy incredible Mexican food; or all of them at the same time. I am so happy we will share the stage for my defence. Sirin, with you I have discovered the pleasure of a hectic life. I feel so fortunate to have shared some years with you. You know all the love I have for you and I hope that some day we settle in the vicinity of each other. I'm still missing your presence when I open the Dixit box. Arianna, you coming to my place for brunch solely based on Sirin's voucher was a great decision. Your unconditional support always was, and is, an incredible comfort that everyone should envy me for (or be jealous of me?). I think I gave you a heart attack by asking you to be my paranymph, sorry for that but believe me, I am thrilled to have you.

I am of course very grateful to all the members of the COMSoC group in Amsterdam. From the early ones let me thank Arthur for bringing some comforting Frenchness to my Amsterdam début. The group definitely improved with your arrival, Adrian. Your sense of humour is yet to be dethroned, and your impersonation skills will definitely never be. I am honoured to count you, Julian, as one of my close friends. I am still optimistically dreaming of finding this one song we will both root for. I hope you forgive me for all the office chit-chat, it was a real pleasure to share it with you. Federico, I hope that the rest of your stay in Amsterdam will be less rocky than this one specific group meeting... But it has been a pleasure discussing mathematical non-sense with you. I never thought that playing beach volleyball in dungarees would be a thing, thanks for creating the circumstances for this to happen, Markus. Last but not least, Jan, you were the game changer for my life in Amsterdam. I believe, as mentioned earlier, that you are a brilliant supervisor, but even more, you are an exceptional friend. I cannot thank you enough for all the evenings we spent together. I was honoured to be your official wedding photographer, please pick me again for any similar event!

Beyond the COMSoC group, the ILLC was and is full of amazing people. Lwenn was among my first close friends in Amsterdam, thank you for the fun we had. Daira, we became friends before we shared an office, and sharing one really brought us closer. I really got accustomed to our long discussions and I will always come to you to entangle my moral conundrums. I hope you appreciate the page count that this thesis has reached. Alina, you were the last one to enter my circle of close friends and

I do not want you out, ever. I am always happy when you are around. Chit-chatting with you, Roos and Caitlin, during my tea breaks was such fun! Thank you for all these moments and for all the support you provided throughout the years. Finally, Marianne, you deserve a special mention for creating the drawing that followed me for my whole PhD, and that eventually gave rise to the cover of this thesis.

Very close to my heart is my first Dutch friend, Abel. I value our friendship so much. Thanks for introducing me to the world of extinguished languages, I enjoyed discussing them more than you would think. But mostly, thanks for all our walks, I am always looking forward to them. And finally, thank you for taking the lead on the translation of my summary, your input was very valuable.

The cover of the thesis would never have seen the light without the draw club. I would never have imagined I could do it if it were not for you, Nicole and Lennart. Thank you for inviting me to join.

Sophie, you entering my life changed everything, and you know that I am choosing my words carefully when writing this. I already wrote in many places what you mean to me so I will not repeat myself here. Nakupenda. Do not forget it, LomL.

The final paragraph goes to the Philæ's fans. Après plusieurs essais, c'est simplement trop bizarre de m'adresser à vous en anglais. Zoé, j'ai adoré que tu viennes me voir à Amsterdam. Reviens, il faut qu'on fasse du paddle, ou un tour de l'Europe en vélo, je ne sais plus. J'ai toujours hâte que l'on s'appelle, et les quelques semaines où l'on ne trouve pas le temps sont étranges. Martin, mon frère favori, j'ai grandement profité de tes connaissances durant ma thèse. Merci d'avoir été mon conseiller quand j'en ai eu besoin, mais plus généralement, pour cette complicité qui rend n'importe quel sujet facile à aborder. Promis on ira punter un de ces quatre! Papa et maman, merci pour tout, les discussions sérieuses, les plus marrantes, les moments partagés à Amsterdam, nos coup de fils régulier, et bien plus. Je suis, encore maintenant, admiratif de tout ce que vous avez fait pour nous trois. Merci.

And as it happened for the thesis as a whole, these acknowledgments are getting too long so let me stop here. Pardon me if your name does not appear, it is most likely a me problem rather than a you one.

Silnon

Initial Words

Chapter 1

Introduction

Not very long after I started my PhD, my mother called me to discuss a problem she was facing at work. At the time, she was involved in the management of a research and development team. They wanted to try out new methods to determine which research projects should be carried out, that is, which ones should be *funded*. The process they envisioned involved two stages. In the first stage, the management team would collect research project proposals from the team members. In the second stage, the team members would be asked to submit their preferences regarding which projects they would want to work on. Based on the reported preferences, the management team would then decide which projects to fund. My mother wanted me to help her decide what procedure the management team should use for this last step. She was thus after a mechanism to find a *collective* allocation of public funds to a collection of projects. Coincidentally, studying such procedures ended up being the exact topic I would work on during my PhD, and is, four years later, at the very heart of this thesis.

This thesis studies scenarios in which a decision regarding how to allocate public funds is to be reached collectively. This very broad definition accounts for a wide variety of processes, including the one described above. As the title of the thesis suggests, we will specifically focus on the case of *participatory budgeting* (PB). Generally speaking, PB is a democratic tool used to allocate a given amount of money to a collection of projects based on a group of individuals' preferences over the projects. It is *participatory* in the sense that the decision is based on the opinion of the individuals who will most directly be impacted by it. This is completely opposite to more classical *expert-based* budgeting decisions where a committee of experts decides on the allocation. Another implication of the participatory objective of a PB process is that a large number of individuals are involved—typically all residents of a city—making PB a *collective* budgeting problem.

Reaching a collective decision is not an easy task. No one who ever tried to plan

a holiday with a group of ten friends would disagree on that. The participatory and the collective aspects of any PB process thus raise a lot of intricate questions. Indeed, when considering expert-based budgeting decisions, it is usually the case that, because of the small number of people involved, simple deliberation will be enough to reach a final consensus. Now, such a deliberative process cannot be run at the scale of a major city such as Paris where hundreds of thousands of citizens are asked to provide their opinion. Thankfully, we have known for centuries how to efficiently determine outcomes based on the opinions of thousands, or millions: by simply running an election. In this thesis, we will approach PB as a voting process.

To help us make the best budgeting decisions for PB, we can thus exploit the large academic literature that analyses voting procedures. This literature has mainly been developed within *social choice theory*, the research area investigating questions related to the aggregation of individual opinions into a collective decision (Arrow, 1951; Arrow, Sen and Suzumura, 2002). Although it has received a lot of attention in the past 70 years, social choice theory cannot be directly applied to PB, yet. Indeed, PB is a more complex voting setting than the ones typically studied in the literature (Arrow, Sen and Suzumura, 2002, 2011). This is due to the fact that in a budgeting scenario the outcomes we are after (*i*) are composed of several winners who can have different costs, and, (*ii*) should not cost more than a given budget limit. Even though these two requirements may sound straightforward, they do make the setting much more involved for a formal analysis. For instance, even the simpler setting of multi-winner voting—where outcomes should satisfy (*i*) and (*ii*) but with alternatives that all have the same cost—required many years of research until it reached a stable state, which only happened recently (Lackner and Skowron, 2023). The formal analysis of PB is still at a burgeoning stage and the literature, though growing at a fast pace, is still fairly sparse.

Overall, the question of how to aggregate preferences in a budgeting scenario is far from being answered, and I am not ashamed to admit that I had no concrete suggestion for my mother back in October 2019. The picture has changed since then, and even though I would still struggle to provide a definitive answer to my mother’s question, I would probably be more useful. The content of this thesis would actually be of great help for that, as the underlying motivation for all the pages to come is to provide an answer to a similar question, namely:

How should we aggregate preferences over costly alternatives into a collective decision regarding which of the alternatives to fund, given a budget limit?

There is no single answer to this question, and without additional context it is probably impossible to answer it. Throughout this dissertation, we will discuss and investigate several possible answers. We will analyse budgeting procedures from various angles, investigate numerous budgeting settings, draw connections between different formal frameworks, and much more.

Providing a definite answer to this question will not only help my mother but also positively impact society in general. The investigation presented in this thesis indeed

fits within a broader research agenda that aims at developing better ways of organising grassroots democracy initiatives. As we will explain later, PB processes are by their very nature grassroots as they were initially developed to reignite the Brazilian democracy. Years later, PB is still viewed as an important innovative tool to improve the democratic process. Getting a better understanding of how to make good, fair, efficient, and/or optimal decisions for such processes is thus of prime importance to sustain the development of participatory democracy. To the question “What are you working on?”, one of my dear colleagues once answered “We are fixing democracy here!”. This is, on a very small scale, what this thesis is about.

We have now outlined the general context of this thesis. Many important details have been omitted for the sake of simplifying the exposition, and the next point on our agenda is to clarify them all. This will be done in the rest of this chapter. We will also introduce the different topics covered in the thesis.

1.1 The Object: Participatory Budgeting

Participatory budgeting is an *innovative democratic institution* through which citizens are involved in the decision process of the allocation of public funds (Wampler, McNulty and Touchton, 2021). It has originally been developed by politicians in the city of Porto Alegre, Brazil, where it was first implemented in 1989. After years of dictatorship, the hope of the local politicians was to establish Brazil’s new representative democracy by increasing the number of administrative mechanisms involving citizens (Abers, 2000). After this initial success, PB spread to other Brazilian cities, and not long after was used worldwide (Porto de Oliveira, 2017; Dias, 2018; Dias, Enríquez and Júlio, 2019).

1.1.1 Definition

With the rise of PB processes throughout the world, more and more diverse mechanisms have been implemented under the name of PB, making it hard to provide a clear definition of what actually is a PB process. Instead of providing a direct definition, political scientists usually prefer to characterise PB processes through the properties they satisfy. Following this idea, Sintomer, Herzberg and Röcke (2008) present five criteria making any budgeting process involving non-elected citizens a PB process:

- ▶ It should be about the allocation of scarce resources;
- ▶ It should involve a public institution (city, district) with an elected body and power over administration and allocation of resources;
- ▶ It has to be repeated over the years;
- ▶ It has to allow for public deliberation phases;

- ▶ It should implement some mechanisms enforcing accountability on the result.

The above list informs us about the organisational aspects of a PB process. Complementing this approach, [Wampler \(2012\)](#) identifies five¹ core principles that need to be implemented (at least in part) by any PB processes in order to generate social change:

- ▶ **Voice:** citizens are offered a chance to voice their opinions and ideas;
- ▶ **Vote:** by voting in the PB process, citizens actively take part in state-sanctioned decision-making processes;
- ▶ **Social justice:** areas that are more in need are targeted to achieve a better redistribution of resources;
- ▶ **Social inclusion:** traditionally marginalised groups are offered additional opportunities to be represented;
- ▶ **Oversight:** citizens are involved at every step of the process to organise it, monitor the implementation of the projects, and so on.

These two sets of principles and criteria only offer us a general overview of the key components of any PB process, but do not touch on how PB processes are actually implemented. This is the topic of the next section.

1.1.2 Implementation

It is once again hard to provide a general description of how a PB process is organised given the multiplicity of actual implementations. Still, researchers have been able to single out several key steps that almost all PB processes follow ([Wampler, 2000](#); [Cabannes, 2004](#); [Shah, 2007](#)). We present them below.

- ▶ Regular meetings are held by the municipality to discuss potential projects that could be funded using the available budget. Typically, these projects are proposed by the citizens.
- ▶ A shortlist of potential projects is decided upon, usually, by collecting all proposals that are feasible and fit the requirements of the PB process. Additionally, the cost of each possible project is determined, either by experts from the municipality or by the citizens who submitted the project.
- ▶ Citizens vote on the shortlisted projects to determine which of them will be funded, given the budget constraint.
- ▶ The municipality reports back to the citizens on the advancement of the actual realisation of the selected projects.

¹Note that [Wampler \(2012\)](#) only discusses *four* core principles. Social inclusion was only added at a later stage (see, e.g., [Wampler, McNulty and Touchton, 2021](#)).

This description of a typical PB process is the one we adopt for the thesis. Specifically, when referring to a PB process, we will have in mind a mechanism implementing these steps. Note that we will mainly focus on the voting stage, *i.e.*, the third bullet point of the above.

As should be clear by now, this description only fits most PB processes, but not all. For instance, the above is phrased as if the organising entity was a municipality. This is the typical case; however, the scale of the process can vary significantly: from schools,² or housing communities,³ to neighbourhoods of a city,⁴ and even to subnational entities.⁵ It is also interesting to note that not all the processes include a voting stage. Indeed, some PB processes are organised as a simple deliberative mechanism throughout which the set of projects to implement is determined meeting after meeting. This was typically the case for some of the first PB processes implemented in Brazil (Cabannes, 2004). Any PB process considered in this thesis will be assumed to include a voting stage however.

1.2 The Question: Aggregating Preferences for Participatory Budgeting

The previous section specifically highlights the existence of a plethora of ways PB processes are implemented. Such variety in the object of study hinders the development of a unique formal approach for PB. Instead, it calls for a multi-faceted analysis, one that can investigate PB processes in their full variety. This is the approach adopted in this thesis, with a special focus on the voting stage of PB processes.

With this in mind, we can rephrase the central question we posed on page 4 in a more accurate way. Our goal is to understand ways of aggregating preferences in PB settings, with an emphasis on incorporating the variety of PB processes in the analysis. Essentially, this thesis tries to answer the following research question:

How should we aggregate the reported preferences in a PB setting, taking into account the variety of forms PB processes can take?

In other words, given a set of costly alternatives, we are interested in knowing how to aggregate the reported preferences over the alternatives into a collective decision regarding which of them to fund, subject to a budget constraint. In addition, we are interested in aggregation methods that make it possible to account for the multiplicity of actual implementations of PB processes.

²See participatorybudgeting.org/pb-at-ps139 for an example of PB within a primary school.

³See the example of social housing in Scotland for instance: sharedfuturecic.org.uk/participatory-budgeting-within-social-housing-ideas-for-better-engaging-with-tenants-and-residents-groups.

⁴*E.g.*, Amsterdam organises individual PB processes for each district (City of Amsterdam, 2022).

⁵PB processes were organised at the scale of a regional department in Peru (Shah, 2007).

1.3 The Method: (Computational) Social Choice

As we briefly mentioned earlier, we will provide an answer to our research question following the (*computational*) *social choice* method (Arrow, Sen and Suzumura, 2002, 2011; Brandt, Conitzer, Endriss, Lang and Procaccia, 2016a; Endriss, 2017).

Social choice is a research field that emerged when highly educated individuals started to methodologically study ways of determining the winner(s) of an election. The actual birth of the research field is usually attributed to two multi-faceted scholars—Nicolas de Condorcet and Jean-Charles de Borda—who argued in the 18th century about how to best determine the winner of an election (Borda, 1781; Condorcet, 1785; McLean and Urken, 1995). The concepts they discussed at the time—today known as the Borda rule and the Condorcet paradox—are now of fundamental importance to the formal analysis of voting. Social choice theory only blossomed many years later when Kenneth Arrow published his seminal book featuring his celebrated *impossibility theorem* (Arrow, 1951). Since then, social choice became an active field of research, mainly within the economics community. The last important turn for social choice theory happened in the '90s and early '00s when computer scientists started to investigate social choice problems using tools from theoretical computer science, and to use social choice tools for their own problems (Bartholdi, Tovey and Trick, 1989; Hudry, 1989; Ephrati and Rosenschein, 1993). This led, some years later, to the establishment of *computational* social choice as a research field (Endriss and Lang, 2006; Brandt, Conitzer, Endriss, Lang and Procaccia, 2016a), which is still very active now (Endriss, 2017; Laslier, Moulin, Sanver and Zwicker, 2018).

Though it started with the study of (single-winner) voting methods, the focus of (computational) social choice theory now includes a much broader set of applications including more complex voting settings, the fair allocation of items to agents, coalition formation between agents, and many more (Brandt, Conitzer, Endriss, Lang and Procaccia, 2016a; Endriss, 2017). What do all these topics have in common? They all are concerned with determining collective decisions based on the preferences reported by the agents involved in the process. Overall, social choice theory can be defined as the field of research that is concerned with the ways of *aggregating* individual opinions into collective decisions.

Social choice theory thus studies *aggregation methods*, trying to single out a set of most appealing ones for different settings. To do so, the standard method is to adopt the *axiomatic* approach. Following this approach, one would first devise a set of normative properties that are considered appealing for an aggregation method—the so-called *axioms*—and would then analyse different aggregation methods in light of these axioms (Thomson, 2001). This clearly is the standard method in social choice and it has led to many of its most famous results, such as impossibility theorems (Arrow, 1951; Gibbard, 1973; Satterthwaite, 1975), and characterisation results (Black, 1948; May, 1952; Young, 1974; Young and Levenglick, 1978; Moulin, 1980).

The computational take on social choice adds to the axiomatic approach the toolbox inherited from computational complexity theory (Arora and Barak, 2009). Broadly

speaking, complexity theory assesses how hard it would be to solve certain problems using a computer. It does so by classifying problems based on the number of elementary steps that would be required to solve them. Using this approach, researchers have been able to analyse the complexity of determining the outcome of aggregation methods (Hemaspaandra, Hemaspaandra and Rothe, 1997; Lipton, Markakis, Mossel and Saberi, 2004; Hemaspaandra, Spakowski and Vogel, 2005), of computing successful manipulations by the agents or the decision makers (Bartholdi, Tovey and Trick, 1989; Faliszewski, Hemaspaandra and Hemaspaandra, 2009; Conitzer and Walsh, 2016; Faliszewski and Rothe, 2016), and of verifying structural properties (Escoffier, Lang and Öztürk, 2008; Elkind, Lackner and Peters, 2022), among others. We will see in the coming chapters specific computational problems designed for the analysis of PB. Additional new tools came with the computer science perspective (numerical simulations, computer-aided methods), that we will also make use of later on.

Overall, our toolbox for investigating aggregation problems for PB will include both the standard axiomatic analysis from social choice theory, together with the computational complexity and numerical simulation analysis from computer science.

1.4 The Contribution: Variations on Participatory Budgeting

All the elements are now in place. We have presented the object of our study: participatory budgeting; the research question: aggregation problems in PB; and the method: computational social choice. The last step is now to describe the contribution of this thesis, that is, its actual content.

Let us first position this work with respect to the literature. This thesis is not the only work that uses tools from computational social choice to analyse aggregation problems for PB. Nonetheless, the literature on this topic is relatively recent, and only a handful of papers had been published when I started working on it in 2019. As we will see later (Chapter 3 is dedicated to surveying the literature), the literature has grown significantly since then, and many aspects of the aggregation problems for PB have already been studied. However, the focus has almost exclusively been on the voting stage of PB processes with the only constraint being the budget limit.

Overall, the research question posed on page 4 is already (partially) answered by the literature, but its refinement stated on page 7 is left mostly unanswered. The use of “mostly” here is due to the observation that several published papers contribute to the study of formal models that are extensions of the standard PB model (see Section 3.6). However, all these works—with the only exception of Hershkovitz, Kahng, Peters and Procaccia (2021)—do not capture alternative models of PB as implemented at a large scale in real life, but rather explore variations of the model that could *potentially* be implemented.⁶ By focusing on actually implemented variations of PB, the present

⁶This can be witnessed by the absence of real-life examples in all the papers cited in Section 3.6.

thesis contributes in closing this gap in the literature.

Let us now turn to the content of the thesis. All the chapters will be introduced below, providing a general overview of the contribution of the thesis.

The first item on our agenda is to provide general background information on PB. This will be done in Part **One**. We will first formalise our approach in Chapter 2. There, the standard model of PB will be defined, and all other relevant concepts will be introduced. More background information will be provided in Chapter 3, which presents a comprehensive survey of the computational social choice literature on PB.

Once the scene will have been set, we will turn to the technical contribution. The research question clearly highlights our will to account for variations in PB. We will do so following two general directions, each making up for one part of the thesis. In Part **Two** we will focus on variations on the *method*; that is, we will present alternative ways of investigating the aggregation problems in PB compared to the standard ones presented in the survey (Chapter 3). Subsequently, in Part **Three**, we will discuss variations on the *model*; that is, we will present alternative models for PB that capture aspects of PB processes that cannot be captured in the standard model introduced in Chapter 2. Finally, we will draw a conclusion of our journey in Chapter 9.

We will now provide more details regarding the two technical parts. References to the published material that serves as the basis of each chapter will be provided there as well. Note that the content of Chapter 3 is largely based on [Rey and Maly \(2023\)](#). Moreover, the material presented by both [Rey \(2022\)](#) and [Motamed, Soeteman, Rey and Endriss \(2022\)](#) has been useful in several places of the thesis, even though they do not contain results directly included here.

1.4.1 Variations on the Method

We first focus on the variations on the method that will be presented in Part **Two**.

Fairness in PB via equality of resources. Our research question is concerned with the ways of aggregating the reported preferences into a collective decision in PB scenarios. One of the fundamental aspects of this question is to investigate how to reach *fair* collective decisions. Fairness in PB has been the focus of most of the research on PB from (computational) social choice researchers (see Chapter 3). The most prominent concepts studied in this branch of the literature are based on measures of the satisfaction of a voter. However, in many cases, it is impossible for the decision maker to access the satisfaction of a voter. There are several explanations for that, one of the most compelling ones being the limited information reported by the voters in their ballots. In Chapter 4 (mainly based on the content presented by [Maly, Rey, Endriss and Lackner, 2023](#)) we present a novel approach for fairness in PB based, not on satisfaction, but on a measure of the amount of resources used to try to satisfy a voter. As we will see, this approach gives rise to a new fairness theory for PB based on equality of resources. In this chapter, we demonstrate that our new take on

fairness is fully viable and bypasses the limitations regarding fairness concepts based on the satisfaction of the voters. This first variation on the method allows us to talk about fairness, even in settings where the satisfaction of a voter cannot be accessed, thus increasing the applicability of the formal analysis.

Epistemic approach to PB. Not all PB scenarios directly correspond to the case of the municipality running a process to involve citizens in budgeting decisions. There is for instance a set of PB applications where the goal is to find the best alternatives, for some notion of “best”. This was the case in our initial example where some research projects are better than others (in the sense of having a higher probability of success) and the process is there to retrieve these best projects. In Chapter 5 (based on the material presented by [Rey and Endriss, 2023](#)) we present an *epistemic* approach to PB, the first of its kind. This epistemic view is based on two essential assumptions: (i) there exists an objectively best, but unknown, outcome called the ground truth, and (ii) the ballots of the voters are noisy estimates of the ground truth. Aggregation mechanisms should then aim at retrieving the ground truth from the ballots. In this chapter, we will look for aggregation mechanisms for PB that have good epistemic properties, a task that will prove difficult. This second variation on the method allows us to account for different kinds of PB processes, ones that have different goals than the canonical municipality-run PB process.

1.4.2 Variations on the Model

Let us now turn to the variations on the model, presented in Part [Three](#). Building on the standard model presented in Chapter 2, chapters in this part present alternative formal models of PB that capture new aspects of actually implemented PB processes.

PB with additional constraints. Throughout this introduction we have always been talking about a unique *budget limit* without ever mentioning any other constraints on the outcome of the PB process. This is a limitation of the model as many cities actually do include additional constraints. For instance, in the PB process organised by the city of Lyon, projects are grouped based on the neighbourhood they belong to. Even though voters are not constrained on the projects they can vote for, multiple budget constraints are imposed, one for each neighbourhood. This example is far from being unique and many more can be found in real-life implementations of PB, potentially with different types of extra constraints. To account for this variety of constraints on the outcome, we develop in Chapter 6 (based on the results of [Rey, Endriss and de Haan, 2020](#)) a general framework for aggregating reported preferences for PB when additional constraints are imposed on the outcome (on top of the global budget constraint). Developing a general approach will lead a number of technical difficulties and we will provide elaborate ways of circumventing them. Overall the study of this first variation on the model will allow us to account for a much larger set of PB processes, bringing us closer to answering our research question.

Long-term approach to PB. As already mentioned in Section 1.1.1 when we presented criteria characterising PB processes, it is in the essence of any PB process to be repeated over the years. Up to Chapter 7, this aspect of PB processes will not have been covered at all (both in the thesis and in the literature). In this chapter, we will develop an approach to analyse repeated PB processes from a fairness standpoint. This is, to this date, the first formal analysis of the repetitive aspect of PB processes. Of course, one could simply analyse long-term PB processes as a repetition of single-shot instances. However, the temporal aspect opens the door to a much deeper analysis. Focusing on fairness, we will take advantage of the temporal aspect of our model to introduce several novel ways of considering fairness for PB, aiming to enforce fairness over time (in addition to fairness for each single instance). This second variation on the model makes it possible to capture one of the key criteria defining PB processes. The technical material of this chapter is taken from [Lackner, Maly and Rey \(2021\)](#).

End-to-end analysis of PB. Getting back to the typical skeleton of a PB process we presented in Section 1.1.2, it is clear that citizens are involved in several steps of the process. So far, we only considered the voting stage, the one for which the question of the aggregation of the ballots is the most prominent. However, other stages of the process also call for a social choice analysis. In Chapter 8 (based on [Rey, Endriss and de Haan, 2021](#)) we present the first formal model incorporating the stage during which citizens submit project proposals in the analysis. We analyse both the aggregation problems inherent to this stage (how to determine a shortlist of proposals), and the impact of this two-stage model on the behaviour of the voters. With the contribution presented in this chapter, we are able to assess PB processes in their full length, once again bringing us closer to our goal.

1.4.3 Beyond the Technicalities

The tools we will use to analyse PB are deeply mathematical, making the exposition of the thesis very formal. However, as we discussed already, the end goal of the thesis is to support the development of participatory democracy. To reconcile these two aspects—the formal aspect of our analysis, in contrast with the societal end-goal of the thesis—one needs to see beyond the formal analysis.

It is true that the analysis in itself can be technical, but the object of the analysis needs not be. The beauty of the field comes from its very formal analysis of aggregation procedures, that can be easily used as-is for real-life instances. Focusing on our concrete problem, what I mean here is that even though the contribution of the thesis is mainly technical, most of the aggregation procedures and concepts we will study are relatively simple and can be explained to a general audience. There is thus a way to see beyond the technicalities and to focus on their actual implications.

With that in mind, let us now delve into the technicalities. While doing so, keep in mind that we are improving the democratic process here!

Part One

General Background

Chapter 2

The Standard Model of Participatory Budgeting

In this chapter we introduce what we call the *standard model* of PB. This corresponds to the model that is the most widely studied in the literature, and will also be the one that we will study throughout this thesis (with some variations in Part [Three](#)).

The notation we introduce in this chapter will be the one used in the thesis. Additional notation, specific to each chapter, will be introduced along the way.

To help the reader, an index of the concepts and a nomenclature of the symbols appearing in the thesis can be found at the end (pages [267](#) and [271](#) respectively).

2.1 Setting up the Voting Stage

The *voting stage*, sometimes referred to as the *allocation stage*, is part of a PB process that we are mainly interested in. What we call the standard model of PB actually only models the voting stage of real-life PB processes.

An *instance* of the voting stage, or simply an instance, is a tuple $I = \langle \mathcal{P}, c, b \rangle$ where $\mathcal{P} = \{p_1, \dots, p_m\}$ is the *set of projects*; $c : \mathcal{P} \rightarrow \mathbb{R}_{>0}$ is the *cost function*, associating every project $p \in \mathcal{P}$ with its cost $c(p) \in \mathbb{R}_{>0}$; and $b \in \mathbb{R}_{>0}$ is the *budget limit*. For any subset of projects $P \subseteq \mathcal{P}$, we denote by $c(P)$ its total cost $\sum_{p \in P} c(p)$. Note here that we make the common assumption for the costs and the budget limit to be non-negative.

An instance $I = \langle \mathcal{P}, c, b \rangle$ is said to have *unit costs* if for every project $p \in \mathcal{P}$, we have $c(p) = 1$ and $b \in \mathbb{N}_{>0}$. In the *unit-cost setting*, we restrict our attention to unit-cost instances. This setting is particularly interesting as it corresponds to the multi-winner voting setting ([Lackner and Skowron, 2023](#)), at least if we also require budget allocations to use up all the budget.

The outcome of the voting stage $I = \langle \mathcal{P}, c, b \rangle$ is a *budget allocation* $\pi \subseteq \mathcal{P}$ such that $c(\pi) \leq b$. We will denote by $\text{FEAS}(I)$ the set of all the *feasible* budget allocations for instance I , defined as:

$$\text{FEAS}(I) = \{\pi \subseteq \mathcal{P} \mid c(\pi) \leq b\}.$$

A feasible budget allocation $\pi \in \text{FEAS}(I)$ is said to be *exhaustive* if there is no project $p \in \mathcal{P} \setminus \pi$ such that $c(\pi \cup \{p\}) \leq b$, that is, if π is cost-maximal. We will denote by $\text{FEAS}_{\text{EX}}(I)$ the set of all feasible and exhaustive budget allocations for an instance I :

$$\text{FEAS}_{\text{EX}}(I) = \{\pi \in \text{FEAS}(I) \mid \pi \text{ is exhaustive}\}.$$

2.2 Collecting the Votes

Let $\mathcal{N} = \{1, \dots, n\}$ be the *set of voters* involved in the PB process. When facing an instance $I = \langle \mathcal{P}, c, b \rangle$, they are asked to submit their preferences over the projects in \mathcal{P} . Several ballot formats have been considered for PB as we shall see in Chapter 3. Except for that one chapter, we will only focus on approval ballots throughout this thesis. When submitting an *approval ballot*, a voter indicates the projects that they approve of. For agent⁷ $i \in \mathcal{N}$, we denote by $A_i \subseteq \mathcal{P}$ the approval ballot that agent $i \in \mathcal{N}$ is submitting. Whenever we have $p \in A_i$, we will say that agent i is a *supporter* of project $p \in \mathcal{P}$.

Importantly, in Chapter 3 (and only there), we will use a different mathematical object to refer to approval ballots. We will still use A_i to denote the approval ballot of agent $i \in \mathcal{N}$, however A_i will then be a function $A_i : \mathcal{P} \rightarrow \{0, 1\}$ mapping projects into $\{0, 1\}$ where $A_i(p) = 1$ indicates that agent i approves of project p , *i.e.*, that $p \in A_i$ in the notation introduced above. This is done to highlight the relationship between approval and cardinal ballots, as explained in detail there.

Since an approval ballot is a subset of projects, any property of budget allocations (or subset of projects) is also well-defined for approval ballots. In that respect, we will sometimes discuss feasible, or exhaustive ballots, for instance. Note that feasible approval ballots are sometimes called *knapsack ballots* (Goel, Krishnaswamy, Sakshuwong and Aitamurto, 2019).

The vector $\mathbf{A} = (A_1, \dots, A_n)$ of the ballots of the agents is called a *profile*. Given two profiles \mathbf{A} and \mathbf{A}' , we use $\mathbf{A} + \mathbf{A}'$ to denote the profile obtained by concatenating them. Strictly speaking, this is not well-defined as it requires us to work with electorates of different sizes (the set \mathcal{N} is fixed in our notation). Because the concatenation operation is only needed in Chapter 5, we hope that the reader can forgive this slight informality.

The number of supporters of a given project $p \in \mathcal{P}$ in a profile \mathbf{A} , is the *approval score* of p , formally defined as $\text{app}(p, \mathbf{A}) = |\{i \in \mathcal{N} \mid p \in A_i\}|$.

⁷Note that we use the terms “voters” and “agents” interchangeably, purely for stylistic reasons. The reader may indeed get bored to always read the same terminology all the time.

We make the important assumption that all projects have at least one supporter. Said differently, we assume that the approval score of every project is at least 1. This assumption simplifies a number of definitions as it reduces the need for additional conditions. We will make it explicit when this assumption is crucially needed.

We will often need to discuss the satisfaction of a voter. However, when using approval ballots, there is no obvious way to define a measure of the satisfaction of a voter. Brill, Forster, Lackner, Maly and Peters (2023) introduced the concept of *approval-based satisfaction functions*, which are functions translating a budget allocation into a satisfaction level for the agents, given their approval ballots. This concept will be used extensively in the different chapters. Let us provide the definition.

Definition 2.2.1 (Approval-Based Satisfaction Functions). *For a given instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile \mathbf{A} , an (approval-based) satisfaction function is a mapping $sat : 2^{\mathcal{P}} \rightarrow \mathbb{R}_{\geq 0}$ satisfying the following two conditions:*

- ▶ $sat(P) \geq sat(P')$ for all subsets of projects $P, P' \subseteq \mathcal{P}$ such that $P \supseteq P'$: the satisfaction is inclusion-monotonic;
- ▶ $sat(P) = 0$ if and only if $P = \emptyset$: only the empty set yields a satisfaction of 0.

The satisfaction of agent $i \in \mathcal{N}$ for a budget allocation $\pi \in \text{FEAS}(I)$ is defined as:

$$sat_i(\pi) = sat(\pi \cap A_i).$$

Several satisfaction functions have been introduced in the literature and two of them are now standard. The first one is the *cardinality satisfaction function* (Talmon and Faliszewski, 2019), denoted by sat^{card} , which measures the satisfaction of the voters as the number of selected and approved projects:

$$sat^{card}(P) = |P|.$$

The second one is the *cost satisfaction function* (Talmon and Faliszewski, 2019), denoted by sat^{cost} , which measures the satisfaction of the voters as the cost of the selected and approved projects:

$$sat^{cost}(P) = c(P).$$

2.3 Determining Budget Allocations

Budget allocations are determined through the use of PB rules. A *PB rule* F is a function that takes as input an instance I and a profile \mathbf{A} and that returns a set of feasible budget allocations $F(I, \mathbf{A}) \subseteq \text{FEAS}(I)$. PB rules that always return a single budget allocation are called *resolute*. For simplicity, we will denote the output $\{\pi\}$ of a resolute PB rule by just π . PB rules that are not resolute are called *irresolute*. They potentially

return several *tied* budget allocations. Unless explicitly stated to the contrary, we will assume rules to be *resolute*.

Note that most of the properties we will introduce in the coming chapters concern budget allocations. To avoid unnecessary definitions, we use the same definitions for PB rules. For a given property \mathcal{X} of a budget allocation, we say that a resolute PB rule F satisfies \mathcal{X} if for every instance I and profile \mathbf{A} , the outcome $F(I, \mathbf{A})$ satisfies \mathcal{X} . When needed, we will explicitly specify how properties for budget allocations are lifted to *irresolute* rules.

We now introduce the rules that will be most commonly encountered throughout this thesis. Additional rules will be introduced in Section 3.2.

2.3.1 Utilitarian Welfare Maximising Rules

We first discuss rules that aim at maximising the utilitarian social welfare. The latter is usually defined as the total satisfaction of the agents. Since we are using approval ballots, we define it via satisfaction functions. Given a satisfaction function sat , the *utilitarian social welfare* of a budget allocation $\pi \in \text{FEAS}(I)$ given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile \mathbf{A} is defined as:

$$\text{UTIL-SW}[sat](I, \mathbf{A}, \pi) = \sum_{i \in \mathcal{N}} sat_i(\pi).$$

Remember that sat_i is defined as $sat_i(\pi) = sat(\pi \cap A_i)$.

Returning the budget allocation that maximises the utilitarian social welfare thus constitutes the first concrete example of a PB rule we have seen so far. Let us introduce two specific such rules, based on sat^{card} and sat^{cost} .

The *cardinality welfare maximising rule* MAXCARD is defined for any instance I and approval profile \mathbf{A} as:

$$\begin{aligned} \text{MAXCARD}(I, \mathbf{A}) &= \arg \max_{\pi \in \text{FEAS}(I)} \text{UTIL-SW} [sat^{card}] (I, \mathbf{A}, \pi) \\ &= \arg \max_{\pi \in \text{FEAS}(I)} \sum_{i \in \mathcal{N}} |\pi \cap A_i|. \end{aligned}$$

Similarly, the *cost welfare maximising rule* MAXCOST is defined for any instance I and approval profile \mathbf{A} as:

$$\begin{aligned} \text{MAXCOST}(I, \mathbf{A}) &= \arg \max_{\pi \in \text{FEAS}(I)} \text{UTIL-SW} [sat^{cost}] (I, \mathbf{A}, \pi) \\ &= \arg \max_{\pi \in \text{FEAS}(I)} \sum_{i \in \mathcal{N}} c(\pi \cap A_i). \end{aligned}$$

These definitions give rise to *irresolute rules*. Remember that our default is to work with resolute rules in this thesis. They can be obtained by using some fixed tie-breaking mechanism between all budget allocations maximising UTIL-SW.

It should be clear to any reader who is familiar with the concept that the output of these two cannot be computed in polynomial time (unless $P = NP$).

Interestingly, we can also reinterpret MAXCARD and MAXCOST in terms of approval score. Indeed, for any instance I and profile \mathbf{A} the following holds:

$$\begin{aligned}\text{MAXCARD}(I, \mathbf{A}) &= \arg \max_{\pi \in \text{FEAS}(I)} \sum_{p \in \pi} \text{app}(p, \mathbf{A}), \\ \text{MAXCOST}(I, \mathbf{A}) &= \arg \max_{\pi \in \text{FEAS}(I)} \sum_{p \in \pi} \text{app}(p, \mathbf{A}) \cdot c(p).\end{aligned}$$

These two formulations will prove useful to draw parallels with the knapsack problem (Kellerer, Pferschy and Pisinger, 2004).

Exploiting the connection between the maximisation of UTIL-SW and various knapsack problems, we can use the prolific literature on the topic to derive PB rules approximating the maximum utilitarian social welfare. We will introduce next two examples of such rules. Let us first define the general scheme of a greedy rule.

Definition 2.3.1 (Greedy Scheme). *Consider an instance $I = \langle \mathcal{P}, c, b \rangle$ and a strict ordering \triangleright over \mathcal{P} . The greedy scheme $\text{GREED}(I, \triangleright)$ is a procedure selecting a budget allocation π iteratively as follows. The budget allocation π is initially empty. Projects are considered in the order defined by \triangleright . When considering project p for current budget allocation π , p is selected (added to π) if and only if $c(\pi \cup \{p\}) \leq b$. If there is a next project according to \triangleright , it is considered; otherwise π is the output of $\text{GREED}(I, \triangleright)$.*

With that scheme in mind, we are now ready to define the two greedy variants of MAXCARD and MAXCOST , initially introduced by Talmon and Faliszewski (2019). These two rules will appear regularly in our analysis.

Let us first consider the *greedy cardinality welfare rule*, GREEDCARD . Given an instance I and a profile \mathbf{A} , we say that an ordering of the projects \triangleright is compatible with app/c if we have $p \triangleright p'$ if and only if $\text{app}(p, \mathbf{A})/c(p) \geq \text{app}(p', \mathbf{A})/c(p')$ holds. In other words, the projects are ordered in \triangleright according to their approval score divided by their cost. For any I and \mathbf{A} , GREEDCARD is then defined as:

$$\text{GREEDCARD}(I, \mathbf{A}) = \{\text{GREED}(I, \triangleright) \mid \triangleright \text{ is compatible with } \text{app}/c\}.$$

Similarly, given I and \mathbf{A} , an ordering of the projects \triangleright is compatible with app if we have $p \triangleright p'$ if and only if $\text{app}(p, \mathbf{A}) \geq \text{app}(p', \mathbf{A})$ holds, that is, if the projects are ordered in \triangleright according to their approval score. The *greedy cost welfare rule* GREEDCOST is then defined for any I and \mathbf{A} as:

$$\text{GREEDCOST}(I, \mathbf{A}) = \{\text{GREED}(I, \triangleright) \mid \triangleright \text{ is compatible with } \text{app}\}.$$

As before, the above definitions introduce irresolute rules. To make them resolute one would need to simply select a single suitable ordering of the projects. Note that this can also be interpreted in terms of breaking ties between projects.

As the names suggest, these rules are indeed approximation of MAXCARD and MAXCOST respectively. Indeed, as the definitions of the rules in terms of approval score highlight, they both consider the projects ordered according to their “score per cost” (for their respective notion of score). Standard results from the knapsack literature then ensure that such greedy rule would approximate their respective social welfare objective within a factor 2 (Kellerer, Pferschy, and Pisinger, 2004, Chapter 2).⁸

It should also be stated that the outcomes of these two rules can always be determined in polynomial time.

2.3.2 The Sequential Phragmén Rule

We now leave the world of rules based on measures of social welfare and turn to other kinds of rules. The first one we present is the *sequential Phragmén rule*, an adaptation of a rule introduced at the end of the 19th century by the Swedish mathematician Lars Edvard Phragmén (Janson, 2016). This rule was formally studied in the multi-winner voting literature by Brill, Freeman, Janson and Lackner (2017), and has then been adapted to the PB setting by Los, Christoff and Grossi (2022).

Definition 2.3.2 (Sequential Phragmén). *Given an instance I and a profile A , the Sequential Phragmén rule, SEQPHRAG, constructs budget allocations using the following continuous process.*

Voters receive money in a virtual currency. They all start with a budget of 0 and that budget continuously increases as time passes. At time t , a voter will have received an amount t of money. For any time t , let P_t^ be the set of projects $p \in \mathcal{P}$ for which the supporters of p altogether have more than $c(p)$ money available. As soon as, for a given t , P_t^* is non-empty, if there exists a $p \in P_t^*$ such that $c(\pi \cup \{p\}) > b$, the process stops; otherwise one project from P_t^* is selected, the budget of its supporters is set to 0, and the process resumes.*

Breaking the ties among the projects in any P_t^* in the above definition will lead to a resolute rule. In the irresolute variant, one would consider all possible ways of breaking such ties.

The termination condition we stated above can be surprising at first sight. It is needed for the rule to satisfy priceability, which however comes at the cost of exhaustiveness (see Section 3.4.1).⁹

An equivalent formulation—the so-called *discrete formulation*—will be presented in Definition 3.2.1. This alternative formulation makes it easier to see that the output of SEQPHRAG can always be determined in polynomial time.

⁸Note that for the factor 2 approximation to be formally correct, one needs to either take the outcome of the rules as we defined them, or the most valuable item, whichever has the highest score.

⁹Phrasing the termination condition as it is here also implies that none of the results rely on the way ties are being broken. If one were to use the stopping condition “the rule stops as soon as it would select a project leading to a violation of the budget constraint”, priceability would only be satisfied when ties are broken in favour of the most expensive project.

2.3.3 The Methods of Equal Shares

This section is devoted to the *method of equal shares*. It is similar to SEQPHRAG except that agents receive all their money from the outset. For a description aimed at non-experts, see equalshares.net, a website maintained by Dominik Peters.

This rule has been introduced for PB by [Peters, Pierczyński and Skowron \(2021\)](#)¹⁰ based on the version for multi-winner voting introduced by [Peters and Skowron \(2020\)](#). We provide below the definition for approval ballots of [Brill, Forster, Lackner, Maly and Peters \(2023\)](#).

The inner mechanism of the rule requires a notion of the satisfaction of the agents. Since we are using approval ballots, we will parametrise the rule by a satisfaction function. That way, we actually introduce a family of rules, one for each satisfaction function.¹¹

Definition 2.3.3 (Method of Equal Shares). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile \mathbf{A} the method of equal shares for the satisfaction function sat , $MES[sat]$, constructs a budget allocation π , initially empty, iteratively as follows.*

A load $\ell_i : 2^{\mathcal{P}} \rightarrow \mathbb{R}_{\geq 0}$, is associated with every agent $i \in \mathcal{N}$, initialised as $\ell_i(\emptyset) = 0$ for all $i \in \mathcal{N}$. The load represents how much virtual money the agents have spent.

Given π and a scalar $\alpha \geq 0$, the contribution of agent $i \in \mathcal{N}$ for project $p \in \mathcal{P} \setminus \pi$ is defined by:

$$\gamma_i(\pi, \alpha, p) = \mathbb{1}_{p \in A_i} \cdot \min(b/n - \ell_i(\pi), \alpha \cdot sat(\{p\})).$$

This is the amount i would pay to buy project p for a given α . Importantly, i only contributes to p if $p \in A_i$, i.e., if i approves of p . Note that the above means that agents are initially provided b/n units of the virtual currency.

Given a budget allocation π , a project $p \in \mathcal{P} \setminus \pi$ is said to be α -affordable, for $\alpha \geq 0$, if the total contribution of the agents given α is at least $c(p)$, i.e., if we have:

$$\sum_{i \in \mathcal{N}} \gamma_i(\pi, \alpha, p) \geq c(p).$$

At a given round with current budget allocation π , if no project is α -affordable for any α , $MES[sat]$ terminates. Otherwise, it selects a project $p \in \mathcal{P} \setminus \pi$ that is α^ -affordable where α^* is the smallest α such that one project is α -affordable (π is updated to $\pi \cup \{p\}$). The agents' loads are then updated: $\ell_i(\pi \cup \{p\}) = \ell_i(\pi) + \gamma_i(\pi, \alpha, p)$. A new round then starts.*

Notice that in the above, sat is only ever used on singletons. Notably, this implies that even if the satisfaction function sat is not additive with respect to the projects, $MES[sat]$ is still well-defined.

¹⁰Note that for all references to [Peters, Pierczyński and Skowron \(2021\)](#) we advise the reader to consider the extended version, updated in October 2022 and available at arxiv.org/abs/2008.13276.

¹¹This is not strictly speaking true, as many satisfaction functions would lead to the same rule.

It is important to note that for any *sat*, the outcome of $\text{MES}[\text{sat}]$ can be computed in polynomial time (as long as *sat* can also be). Specific ways of doing so are explained by [Peters, Pierczyński and Skowron \(2021\)](#).

2.4 Discussing Computational Complexity

A large part of the coming technical analysis is dedicated to the study of PB from a computational complexity perspective. We do not wish to provide a full introduction to the topic here, but will still provide some necessary definitions. We assume the reader to be familiar with the basic notions of computational complexity such as the idea of NP-completeness and all the notions necessary to understand its definition (decision problems, polynomial-time solvability, the \mathcal{O} notation,¹² the complexity class NP, polynomial-time reductions,...). The proper definitions can be found in the excellent textbook by [Arora and Barak \(2009\)](#).

Our complexity analysis will not involve many sophisticated notions. The three complexity classes we will use are P, NP and coNP. Remember that a computational problem is in coNP if its complement is in NP. We will usually distinguish between *weakly* and *strongly* NP-complete problems. In case a problem is solvable in polynomial time when the input is represented in unary, we will say that the problem is *solvable in pseudo-polynomial time*, or weakly-polynomial time. A problem is *weakly NP-complete* if it is NP-complete but solvable in pseudo-polynomial time. A problem is *strongly NP-complete* if it is NP-complete even when the input is represented in unary. Moreover, a problem is *solvable in FPT time* on an instance x and for a parameter $k \in \mathbb{N}$ if it is solvable in time in $\mathcal{O}(f(k) \cdot \text{poly}(|x|))$, where $f(\cdot)$ is a computable function, $\text{poly}(\cdot)$ is a polynomial, and $|x|$ is the size of the instance x (see [Downey and Fellows, 2013](#), for more details).

¹²Throughout the thesis, we will interpret the \mathcal{O} notation in a set manner. We will use statement such as “ $f(n)$ is in $\mathcal{O}(g(n))$ ” interpreting $\mathcal{O}(g(n))$ as the class of functions $h(n)$ such that $h(n) \leq C \cdot g(n)$ for some constant C and all $n \geq n_0$ for a given n_0 . See the discussion in [Graham, Knuth and Patashnik \(1994\)](#) for more details as why it matters.

Chapter 3

The (Computational) Social Choice Take on Participatory Budgeting

This chapter presents a review of the computational social choice literature dedicated to (indivisible) PB. The goal of this chapter is to provide a comprehensive set of definitions and to unify concepts and notation that appeared in different publications. It is not necessary for the reader to go through all of it to understand the rest of the thesis. However, we will regularly refer to the concepts introduced here. More specifically:

- ▶ Many fairness concepts presented in Section 3.3 will be used in Chapter 4;
- ▶ Social welfare notions from Section 3.5.2 will be considered in Chapter 5;
- ▶ The monotonicity axioms defined in Section 3.4.2 will be part of the investigation presented in Chapter 6.

The model and notation in this chapter follow that introduced in Chapter 2, except for the ballots, as will be made explicit. Ballot formats will actually be our first point of focus (Section 3.1). Once the design of the ballots will be clarified, we will discuss rules for aggregating said ballots (Section 3.2). We will then present how to assess the quality of these rules in terms of fairness (Section 3.3) and other axiomatic properties (Section 3.4). After that, we will look at the algorithmic aspects of PB (Section 3.5). After discussing the standard model for PB, we will present its variations and extensions that have been introduced in the literature (Section 3.6). We will finally provide interesting pointers to related frameworks (Section 3.7).

3.1 Ballot Design

Ballot design is an important part of the research on PB. Indeed, the outcome space being combinatorial in nature, the design of the ballots is critical to achieve a good

balance between the amount of information elicited and the practical usability of said ballot. To get the maximum amount of information, we would want to offer the agents the possibility to submit their preferences over all possible budget allocations. These could take the forms of orderings over $\text{FEAS}(I)$, or utility functions associating a score to every feasible budget allocation $\pi \in \text{FEAS}(I)$. This approach clearly cannot be implemented in real life as the size of $\text{FEAS}(I)$ is exponential in the number of projects, which in itself might already be quite large (in 2023 there were 138 projects in the Warsaw PB process¹³).

Several ballot formats have then been designed in the pursuit of the best trade-off between the amount of information that is elicited and the usability of the ballot. All of these formats are *project-based* ballots, *i.e.*, the information collected concerns the projects and not the feasible budget allocations. This is mainly because the set of all the feasible budget allocations can be huge. The approval ballots format (introduced in Section 2.2) is one example. In what follows, we introduce other ballot formats, distinguishing between cardinal ballots (Section 3.1.2)—which associate a score to each project—and ordinal ballots (Section 3.1.3)—which require agents to rank the projects. Note that only approval ballots will be considered in this thesis (outside of the present chapter); we include the others here for the sake of completeness.

To get an overview of the different ballot formats that have been introduced and the papers studying them, we present in Table 3.1.1 a classification of the papers we have reviewed, based on the ballot format they are considering.

3.1.1 Terminology around the Voters

Let us start by clarifying some terminology. By going through the literature on PB, and more generally on computational social choice, it appears that the terms preferences, utility, satisfaction, and ballots are used in a somewhat interchangeable fashion. In the following we suggest exact definitions for each of those, hoping that it will help to clarify and unify the use of these terms.

One distinction that seems important to us is that of the private and public information of the voters. The information submitted by the voters, their *ballots*, is the only information that is publicly available, especially to the decision maker. In no case can the ballots be assumed to represent the internal preference model of the voters. Hopefully, the ballots reflect some aspects of the preferences of the voters, but one cannot claim that they capture them entirely. This observation is based on the following two main arguments. First, we know that almost none of the rules we are studying prevent voters from rationally behaving strategically, so there is no reason to assume their ballots to be truthful (Gibbard, 1973; Satterthwaite, 1975; Dietrich and List, 2007b; Meir, 2018; Peters, 2018). Second, even if voters try to vote truthfully, it is debatable whether they would be able to produce a ballot that faithfully represents

¹³See the data hosted on pabulib.org (Stolicki, Szufa and Talmon, 2020) and the specific Warsaw 2023 file: [poland_warszawa_2023_pb](#).

Cardinal Ballots	
Generic	Benadè, Nath, Procaccia and Shah (2021) – Chen, Lackner and Maly (2022) – Los, Christoff and Grossi (2022) – Fairstein, Benadè and Gal (2023) – Fluschnik, Skowron, Triphaus and Wilker (2019) – Hershkowitz, Kahng, Peters and Procaccia (2021) – Jiang, Munagala and Wang (2020) – Laruelle (2021) – Los, Christoff and Grossi (2022) – Munagala, Shen and Wang (2022) – Munagala, Shen, Wang and Wang (2022) – Patel, Khan and Louis (2021) – Peters, Pierczyński and Skowron (2021)
Approval	Aziz and Ganguly (2021) – Aziz, Gujar, Padala, Suzuki and Vollen (2022) – Aziz, Lee and Talmon (2018) – Baumeister, Boes and Hillebrand (2021) – Baumeister, Boes and Laußmann (2022) – Baumeister, Boes and Seeger (2020) – Brill, Forster, Lackner, Maly and Peters (2023) – Jain, Sornat and Talmon (2020) – Jain, Sornat, Talmon and Zehavi (2021) – Los, Christoff and Grossi (2022) – Motamed, Soeteman, Rey and Endriss (2022) – Sreedurga, Bhardwaj and Narahari (2022) – Talmon and Faliszewski (2019)
<i>t</i>-Approval	Fairstein, Benadè and Gal (2023)
Knapsack	Benadè, Nath, Procaccia and Shah (2021) – Fairstein, Benadè and Gal (2023) – Goel, Krishnaswamy, Sakshuwong and Aitamurto (2019)
<i>t</i>-Threshold	Benadè, Nath, Procaccia and Shah (2021) – Fairstein, Benadè and Gal (2023)
Cumulative	Skowron, Slinko, Szufa and Talmon (2020)
Ordinal Ballots	
Strict Orders	Lu and Boutilier (2011) – Peters, Pierczyński and Skowron (2021)
Weak Orders	Aziz and Lee (2021) – Laruelle (2021)
Value-for Money	Benadè, Nath, Procaccia and Shah (2021) – Goel, Krishnaswamy, Sakshuwong and Aitamurto (2019) – Fairstein, Benadè and Gal (2023)
Value	Benadè, Nath, Procaccia and Shah (2021) – Fairstein, Benadè and Gal (2023)

Table 3.1.1: Papers studying PB organised by the type of ballots considered.

their true internal preferences due to their *bounded rationality* (Dhillon and Peralta, 2002; Bendor, Diermeier, Siegel and Ting, 2011). It is therefore questionable to assume that a voter’s ballot represents their true preferences, even if voters behave truthfully.

We thus urge researchers to always clarify the assumption they are making about the voters, their internal state, and how they cast their ballots. To help with that, we present below what we believe to be the best way to use this terminology.

- ▶ **Preferences:** The preferences are private information accessible only to the voters themselves, reflecting their views on the possible outcomes of the decision making scenario. Remember from the above that this information may not be accessible in full to the voters (notably because of bounded rationality). In economic theory, it is usually assumed that preferences take the form of *weak* or *incomplete rankings* over the different outcomes (Lewin, 1996), though other representations of the preferences can be argued for (see, e.g., Hansson, 2001). Note that the term “preferences” sometimes indicates that the preferences are *ordinal*, i.e., that they are based on rankings of the outcomes.
- ▶ **Utility:** The utility of a voter is a specific type of preference for which every outcome can be mapped to a specific numerical value. These preferences are sometimes referred to as *cardinal preferences*.
- ▶ **Satisfaction:** The satisfaction of a voter and their utility are often used synonymously. In computational social choice, it is also often used when ballots do not allow agents to report their full utility functions (because of the limited expressiveness of the ballots). In this case, it represents an approximation of the utility of a voter that would be compatible with the ballot submitted. We claim that it is important to always be clear that such satisfaction functions can at most be *proxies* to the utilities of the agents, and in no case their actual level of satisfaction or utility (even if the ballots would allow voters to submit their full preferences). In the following, we use satisfaction as meaning “the satisfaction that the decision maker is assuming the voter enjoys”.
- ▶ **Ballots:** The ballot of an agent is the information they submitted, formatted according to the specified type of ballot. Let us emphasise once again that a ballot is the sole information submitted by the (potentially strategically-behaving) voter and not necessarily a representation of their private information.

This terminology and those definitions are the ones used throughout the thesis.

We are now ready to actually discuss ballot formats. Note that regardless of the format, we denote by A_i the ballot submitted by agent $i \in \mathcal{N}$.

3.1.2 Cardinal Ballots

Let us start with cardinal ballots. Loosely speaking, when these ballots are used, agents are asked to submit a score for all projects. Additional constraints are some-

times imposed on the scores. Note that we refer to this ballot format as *cardinal ballots* and not utility functions or cardinal preferences as they are usually called, in line with our discussion in Section 3.1.1.

Formally, a *cardinal ballot* $A_i : \mathcal{P} \rightarrow \mathbb{R}_{\geq 0}$ for agent $i \in \mathcal{N}$ is a mapping from projects to non-negative scores. Note that in our definition cardinal ballots associate scores with projects and not budget allocations. Of course the definition can easily be adapted to allow voters to submit scores over budget allocations, but since almost no paper (the only potential exception being Jain, Sornat, and Talmon, 2020) is working with cardinal ballots over budget allocations, we decided to keep the simpler definition.

A common assumption (see, e.g., Peters, Pierczyński, and Skowron, 2021) is that the score of a budget allocation for an agent is simply the sum of the scores of the projects it contains. We call this the *additivity* assumption.

Even though cardinal ballots can be used as is for PB, several important variations have been introduced that we discuss below.

Approval Ballots

In this chapter, we represent approval ballots as cardinal ballots by requiring the score of each project to either be 0 or 1, i.e., for agent $i \in \mathcal{N}$, their approval ballot $A_i : \mathcal{P} \rightarrow \{0, 1\}$ is a mapping from \mathcal{P} to $\{0, 1\}$, where for any $p \in \mathcal{P}$, $A_i(p) = 1$ indicates that agent i approves of project p , and $A_i(p) = 0$ that i does not approve of p . Note that this differs from the definition presented in Chapter 2.

It is important to state that approval ballots are the most widely used ballot format in real life PB processes. At the same time, and potentially for that exact reason, it is also the most studied format in the literature (see Table 3.1.1).

One of the main drawbacks of approval ballots is that they are semantically weak: not much information is communicated. In particular, it is unclear what an agent intends to communicate when not approving a project (setting $A_i(p) = 0$ for project p). It is notably ambiguous whether this case should be treated as stating a rejection of the project, or simply stating an indifference status regarding the project. One way of circumventing this issue is to enforce additional constraints on the ballots that allow us to interpret them more accurately.

Semantically Enriched Approval Ballots

As explained above, the semantics of approval ballots is not well defined. This leads to various problems and has prompted researchers to introduce some additional constraints on the approval ballots to correct this.

In practice, it is often the case that voters can only approve of a limited number of projects. When asked for *t-approval ballots*, agents can only approve up to $t \in \mathbb{N}_{>0}$ different projects. This is formalised by imposing the constraint $\sum_{p \in \mathcal{P}} A_i(p) \leq t$ on the ballot A_i of each agent $i \in \mathcal{N}$. This allows us to get some understanding of the

projects that are not approved: they are not part of the top- t projects of the voter (assuming that voters can actually order the projects based on their preferences).

We already mentioned *knapsack ballots*, those are approval ballots that are feasible. Phrasing it differently, when submitting knapsack ballots, agents are asked to provide their most preferred feasible budget allocation. In this sense, knapsack ballots have a clear meaning that can be used to make potentially better decisions.

Another variation of approval ballot are *t -threshold approval ballots* (Benadè, Nath, Procaccia and Shah, 2021). Here, agents are assumed to have private additive utility functions that they are aware of, and they are asked to submit an approval ballot, approving of a project if and only if it provides them with utility at least $t \in \mathbb{R}$.

Cumulative Ballots

When cumulative ballots are used (Skowron, Slinko, Szufa and Talmon, 2020), agents are asked to distribute a certain amount of money (usually b/n , i.e., their share of the budget) over all the projects. Formally, a cumulative ballot A_i is a cardinal ballot such that $\sum_{p \in \mathcal{P}} A_i(p) \leq 1$. The idea behind cumulative ballots is that agents control some share of the budget and indicate how they would want to use that share.

Note that one could also assume that $A_i(p)$ represents the fraction of the budget limit b that voters i believes should be allocated to project p (in total). This interpretation however does not fit with the assumption that projects are indivisible.

3.1.3 Ordinal Ballots

The second main category of ballots that have been studied for PB are ordinal ballots. In this context, the ballot of an agent is an ordering over the projects. Formally, agent i 's ballot A_i is a strict linear order over \mathcal{P} . We will typically denote it by \succ_i where for two projects $p, p' \in \mathcal{P}$, $p \succ_i p'$ indicates that agent i prefers p over p' .

Ordinal ballots can be used as is for aggregation purposes. However, because projects have different costs, the exact semantics of the ordering is not always clear. Several specific ways of ranking the projects have thus been proposed.

When submitting *ranking-by-value ballots* (Benadè, Nath, Procaccia and Shah, 2021), agents are assumed to provide a strict total order over the projects such that a project p is ranked above another project p' if and only p is preferred to p' .

Similarly, *ranking by value-for-money ballots* (Goel, Krishnaswamy, Sakshuwong and Aitamurto, 2019) requires agents to provide rankings of the projects based on their value for money. Note that this is only well-defined when agents are assumed to have private utility functions that they are aware of.

We have only mentioned *strict* rankings above, but weak rankings have also been considered (Aziz and Lee, 2021). A weak ranking will typically be denoted by \succsim with \succ being the strict part of the ranking and \sim the indifference relation, defined as $p \succ p'$ if $p \succsim p'$ but not $p' \succsim p$; and $p \sim p'$ if $p \succsim p'$ and $p' \succsim p$, for any two projects p and p' . Rankings by value or value for money can also be considered with weak rankings.

Bound	Deterministic Distortion		Randomised Distortion	
	Lower	Upper	Lower	Upper
Knapsack	$\Omega(2^m/\sqrt{m})$	$\mathcal{O}(m \cdot 2^m)$	$\Omega(m)$	m
Rankings by Value	$\Omega(m^2)$	$\mathcal{O}(m^2)$	$\Omega(\sqrt{m})$	$\mathcal{O}(\sqrt{m} \cdot \log(m))$
Rankings by Value-for-Money	Unbounded		$\Omega(\sqrt{m})$	$\mathcal{O}(\sqrt{m} \cdot \log(m))$
Det. t -Threshold Approval*	Unbounded		$\Omega(\sqrt{m})$	m
Rand. t -Threshold Approval*	—		$\Omega(\log(m)/\log \log(m))$	$\mathcal{O}(\log^2(m))$

*For t -threshold approval ballots, [Benadè, Nath, Procaccia and Shah \(2021\)](#) distinguish between two cases. In the deterministic case (Det.) the threshold t is chosen arbitrarily by the decision maker once for all the agents. In the randomised (Rand.) case, for each agent, the threshold t is sampled at random from a given distribution. Note that this distinction makes little sense in the deterministic case.

Table 3.1.2: Summary of the results on the distortion of some of the ballot formats obtained by [Benadè, Nath, Procaccia and Shah \(2021\)](#). The deterministic distortion corresponds to the situation where only deterministic PB rules are considered. In the randomised distortion setting, randomised PB rules are also considered.

Finally, it is worth mentioning that in practice voters are only asked to submit incomplete ordinal ballots, typically ranking a small number of projects. We are not aware of any work studying this ballot format, which we could call *t-ordinal ballots*.

3.1.4 Comparison of Ballot Formats

Comparing the merits of different ballot formats is not an easy task. Two approaches have been explored in the literature focusing either on theoretical or empirical results.

Comparison via Distortion

One way to compare different ballot formats is via the *distortion* ([Procaccia and Rosenschein, 2006](#)) they induce. It is a measure of the amount of information communicated by a ballot format for the purpose of identifying a budget allocation that maximises utilitarian social welfare. Specifically, under the assumption that agents have cardinal preferences, the distortion of a ballot format measures the ratio between the maximum social welfare achievable in the knowledge of the full preferences of the agents, to the maximal social welfare achievable when agents submit their ballots according to the specific format.

[Benadè, Nath, Procaccia and Shah \(2021\)](#) provide a complete analysis of the distortion induced by four of the ballot formats we introduced: knapsack and t -threshold approval ballots, rankings by value and rankings by value for money. Table 3.1.2 presents their findings for both deterministic and randomised¹⁴ rules. Note that they

¹⁴A randomised PB rule returns for any instance I and profile \mathbf{A} , not a budget allocation, but a probability distribution over $\text{FEAS}(I)$.

Rank Value-for-Money	780	3.36	3.72
Cardinal Ballot	454	4.07	4.08
t -Threshold Approval	443	3.88	4.05
Rank Value	431	3.94	4.12
Knapsack Ballot	386	4.00	4.15
t -Approval	345	4.2	4.21
	Voting Time in Seconds	Reported Ease of Use	Reported Expressiveness

Figure 3.1.1: Some of the experimental findings of Fairstein, Benadè and Gal (2023) comparing different ballot formats. The voting time column indicates the time in seconds it took participants to submit their opinion for each ballot format. The reported ease of use and expressiveness columns represents the average value reported by the participants about the ease of use and the expressiveness of each ballot format, on a scale from 1 to 5 (the higher the better). The figures have been reproduced with the authorisation of the authors, using the data available in the GitHub repository github.com/rfire01/Participatory-Budgeting-Experiment.

also complemented their theoretical approach with an empirical one on real-life data. Their findings suggest that approval ballots, and more specifically knapsack ballots, may not be the best ballot format when it comes to PB.¹⁵

Comparison via Real-Life Experiments

Another approach to compare ballot formats for PB is to run experiments with human participants who will be asked to use different formats. This is the approach that Fairstein, Benadè and Gal (2023) followed. They recruited 1800 participants on Amazon Mechanical Turk who were then asked to cast their ballot in a format which was selected from a set of 6 for a specific PB instance (selected from a set of 4 instances). For each participant, the time they needed to vote was measured. Additionally, they asked the participants to self-report on the ease of use of the different formats.

Some of the findings from Fairstein, Benadè and Gal (2023) are presented in Figure 3.1.1. They studied the following ballot formats: generic cardinal ballots, 5-approval ballots, knapsack ballots, 10-threshold approval ballots, rankings by value and rankings by value-for-money. Summarising, all the ballot formats they study require a similar amount of time for the participants to cast, except for ranking by value for money for which participants take significantly longer. The results are the same for the self-reported measures. Notably, for all measures t -approval ballots outperform all the other ballot formats, though not by a large margin.

¹⁵The intuition as to why knapsack ballots do not behave well with respect to distortion is that in the worst case, when all projects cost exactly the budget limit b , knapsack ballots only elicit the favourite project of each agent, and it is well understood that this is not enough to make a high-quality decision.

3.1.5 Ballot-Based Satisfaction

Before we consider PB rules, let us discuss how to model satisfaction based on the different ballot formats we have introduced.

Generic Cardinal Ballots

When asked for cardinal ballots, voters are asked to report their satisfaction level for each project. There is thus no need to consider anything else than the ballot, at least as long as we are operating under the *additivity* assumption. This means the satisfaction of a voter is the sum of the scores they submitted for the projects that have been selected.

Approval Ballots

When it comes to approval ballots, we use the concept of satisfaction function as we introduced in Definition 2.2.1. On top of sat^{card} and sat^{cost} that we already defined, we introduce other satisfaction functions than exist in the literature. First, note that with indivisible projects, sat^{cost} is equivalent to the *overlap satisfaction function* of Goel, Krishnaswamy, Sakshuwong and Aitamurto (2019).

- ▶ **The Chamberlin-Courant satisfaction function** (Talmon and Faliszewski, 2019) measures the satisfaction of the voters as being 1 if at least one approved project was selected, and 0 otherwise:

$$sat^{CC}(P) = \mathbb{1}_{P \neq \emptyset}.$$

- ▶ **The square root and log satisfaction functions** (Brill, Forster, Lackner, Maly and Peters, 2023) measure the satisfaction of the voters as (marginally) diminishing when the cost of a project increases:

$$sat^{\log}(P) = \log(1 + c(P)) \qquad sat^{\sqrt{\cdot}}(P) = \sqrt{c(P)}.$$

In general, all the satisfaction functions we introduced apply seamlessly to all approval-like ballot formats (*t*-approval, knapsack, *t*-threshold...). Some are however more meaningful with some ballot formats than others.

One might wonder what the difference between an approval profile together with a satisfaction function sat , and a cardinal profile is. Assuming sat is additive, an approval profile with a satisfaction function is a special case of a cardinal profile in which every agent approving a project p has the same satisfaction for p . This is a natural assumption, given the limited information about the voters' preferences. However, some authors have proposed to model the satisfaction of voters in a way that also takes additional information into account, for example the non-approved projects in

the winning bundle. This cannot be modelled with a satisfaction function as defined by Brill, Forster, Lackner, Maly and Peters (2023).

For instance, Goel, Krishnaswamy, Sakshuwong and Aitamurto (2019) introduce a measure of the dissatisfaction of the voters in terms of the L_1 distance between a given budget allocation and their ballot. This cannot be modelled by approval-based satisfaction (even though they introduce it in a framework with knapsack ballots) as the satisfaction of a voters depends on projects that are outside of the selected and approved ones. It is important to keep in mind however that the authors deem this measure of satisfaction to be of limited relevance when the projects are indivisible.

Ordinal Ballots

To measure satisfaction with ordinal ballots, one can associate each project in the ordering with a given satisfaction level. This is usually done through *positional scoring functions* that associate each project with a score that only depends on the position of the project in the ranking. That is the approach followed by Laruelle (2021) for instance.

Satisfaction with ordinal ballots can also be defined in more general terms (not simply mapping projects to scores). For instance, Aziz and Lee (2021) compare sets of projects according to the cost of the projects ranked above a certain threshold, where the threshold is context-dependent. Note that this assumption is never explicitly stated and that this reflects our understanding of their definitions.

3.2 Participatory Budgeting Rules

We now turn to PB rules. We introduce additional rules (on top of the ones defined in Section 2.3) that are mainly used together with cardinal ballots.

Our exposition will start with welfare-maximising rules (Section 3.2.1). Thereafter, we will discuss three rules based on the idea of finding budget allocations that spread the cost of the selected projects nicely among the voters: the sequential Phragmén rule (Section 3.2.2), the maximin support rule (Section 3.2.3), and the method of equal shares (Section 3.2.4). A brief overview of the other rules that have been introduced in PB will conclude this part of the chapter (Section 3.2.5).

3.2.1 Welfare-Maximising Rules

In a purely utilitarian view, agents are assumed to have cardinal preferences over budget allocations and the aim is to select a budget allocation that maximises the overall utility of the agents. That is, utilitarian rules aim to achieve high utilitarian social welfare, where the utilitarian social welfare—which we denote by UTIL-SW—is defined for a given instance $I = \langle \mathcal{P}, c, b \rangle$, budget allocation $\pi \in \text{FEAS}(I)$ and a utility

function $\mu_i : 2^P \rightarrow \mathbb{R}_{\geq 0}$ for every agent $i \in \mathcal{N}$ as:

$$\text{UTIL-SW}(I, (\mu_i)_{i \in \mathcal{N}}, \pi) = \sum_{i \in \mathcal{N}} \mu_i(\pi).$$

Here, $\mu_i(\pi)$ denotes the utility of agent i for allocation π . As already mentioned, the decision maker does not have access to the utility of the agents, so welfare-maximising rules have to be defined in terms of the assumed satisfaction for a given ballot.

When using cardinal ballots we usually assume that the satisfaction of an agent is equivalent to their cardinal ballot. Therefore the above definition directly induces a PB rule if, in a slight abuse of notation, we equate the ballot of a voter with their utility. For a given I and \mathbf{A} , the rule selects the budget allocation that maximises:¹⁶

$$\text{UTIL-SW}(I, \mathbf{A}, \pi) = \sum_{i \in \mathcal{N}} \sum_{p \in P} A_i(p).$$

This measures the total satisfaction of the voters (assuming additivity of the ballots).

Still with cardinal ballots, [Fluschnik, Skowron, Triphaus and Wilker \(2019\)](#) study the utilitarian Chamberlin-Courant social welfare (that aims at finding *diverse* knapsacks in their terminology) with cardinal ballots. They also study the maximisation of the Nash social welfare, defined as the product of the satisfaction of the agents (once again defined formally in Section 3.5.2). Their motivation is more algorithmic, however, and they don't necessarily aim to devise PB rules.

When it comes to approval ballots, the main rules considered are MAXCARD, MAXCOST, GREEDCARD and GREEDCOST as introduced in Section 2.3.1. On top of these four rules, [Talmon and Faliszewski \(2019\)](#) introduce five extra rules. They additionally consider welfare defined in terms of sat^{CC} (see Section 3.1.5), and another greedy scheme to approximate the maximum social welfare (proportional greedy rules). [Baumeister, Boes and Seeger \(2020\)](#) complemented the work of [Talmon and Faliszewski \(2019\)](#), showing that two of their rules are actually equivalent, and introducing another greedy scheme (hybrid greedy rules).

Another measure of social welfare was studied by [Sreedurga, Bhardwaj and Narahari \(2022\)](#) in the context of PB with approval ballots: maximin social welfare—which we call egalitarian social welfare in Section 3.5.2—that measures the welfare of a society as the satisfaction of its least satisfied member. [Sreedurga, Bhardwaj and Narahari \(2022\)](#) consider the maximisation of the egalitarian social welfare as a PB rule, studying its computation and its axiomatic properties.

Finally, coming to ordinal ballots, [Laruelle \(2021\)](#) studies welfare-maximising rules with weak ordinal ballots where positional scoring functions are used to measure

¹⁶Note that even though the signature of the functions may look the same, there is a clear conceptual difference between the social welfare defined with utility functions, and UTIL-SW for cardinal ballots: the former uses private information of the voters, while the latter is only defined with respect to public information provided in the ballots.

the satisfaction of a voter (thus obtaining something equivalent to cardinal ballots). Within this framework, [Laruelle \(2021\)](#) defines greedy approximations of the utilitarian social welfare, and one greedy approximation for Chamberlin-Courant social welfare (there called Rawlsian social welfare) that aims at providing every agent with at least one satisfactory project (see Section 3.5.2 for a formal definition).

3.2.2 The Sequential Phragmén Rule

We briefly come back to the SEQPHRAG rule that we introduced in Definition 2.3.2. We provided there its *continuous* formulation. We will now introduce the *discrete* formulation where the *loads* of the voters are to be balanced (see, e.g., [Brill, Forster, Lackner, Maly, and Peters, 2023](#)). The two formulations are equivalent.

Definition 3.2.1 (Sequential Phragmén, Discrete Formulation). *Given an instance I and a profile A of approval ballots, the sequential Phragmén rule, SEQPHRAG, constructs a budget allocation π , initially empty, iteratively as follows. A load $\ell_i : 2^{\mathcal{P}} \rightarrow \mathbb{R}_{\geq 0}$, is associated with every agent $i \in \mathcal{N}$, initialised as $\ell_i(\emptyset) = 0$ for all $i \in \mathcal{N}$. Given π , the new maximum load for selecting project $p \in \mathcal{P} \setminus \pi$ is defined as:*

$$\ell^*(\pi, p) = \frac{c(p) + \sum_{i \in \mathcal{N}} A_i(p) \cdot \ell_i(\pi)}{|\{i \in \mathcal{N} \mid A_i(p) = 1\}|}.$$

At a given round with current budget allocation π , let $P^* \subseteq \mathcal{P}$ be such that:

$$P^* = \arg \min_{p \in \mathcal{P} \setminus \pi} \ell^*(\pi, p).$$

If there exists $p \in P^*$ such that $c(\pi \cup \{p\}) > b$, SEQPHRAG terminates and outputs π . Otherwise, a project $p \in P^*$ is selected (π is updated to $\pi \cup \{p\}$) and the agents' loads are updated: If $A_i(p) = 0$, then $\ell_i(\pi \cup \{p\}) = \ell_i(\pi)$, and otherwise $\ell_i(\pi \cup \{p\}) = \ell^*(\pi, p)$.

To obtain a resolute rule one needs to break the ties among the projects in any P^* . The irresolute rule is obtained by considering all possible ways of breaking such ties.

3.2.3 The Maximin Support Rule

The sequential Phragmén rule can be adapted to allow for redistributing the loads in each round. This leads to the *maximin support rule*. This rule was first introduced by [Aziz, Lee and Talmon \(2018\)](#) in the PB setting. Note that, as observed by [Brill, Forster, Lackner, Maly and Peters \(2023\)](#), they named it sequential Phragmén though they actually generalised the maximin support rule from multi-winner voting ([Sánchez-Fernández, Fernández-García, Fisteus and Brill, 2022](#)). It is defined as follows.

Definition 3.2.2 (Maximin Support Rule). *Given an instance I and a profile A of approval ballots, the maximin support rule, MAXIMINSUPP, constructs a budget allocation π , initially empty, iteratively as follows.*

Given I , \mathbf{A} and a subset of projects $P \subseteq \mathcal{P}$, a load distribution $\ell = (\ell_i)_{i \in \mathcal{N}}$ for P is a collection of functions $\ell_i : 2^P \rightarrow \mathbb{R}_{\geq 0}$ for every agent $i \in \mathcal{N}$ such that $\sum_{i \in \mathcal{N}} \ell_i = c(p)$ for all projects $p \in P$ and $\ell_i(p) = 0$ for all agent $i \in \mathcal{N}$ and project $p \in \mathcal{P}$ for which $A_i(p) = 0$. Omitting I and \mathbf{A} , let $\mathcal{L}(P)$ be the set of all the load distributions for $P \subseteq \mathcal{P}$.

At a given round with current budget allocation π , let $P^* \subseteq \mathcal{P}$ be such that:

$$P^* = \arg \min_{p \in \mathcal{P} \setminus \pi} \max_{\substack{\ell \in \mathcal{L}(\pi \cup \{p\}) \\ i \in \mathcal{N}}} \sum_{p' \in \pi \cup \{p\}} \ell_i(p').$$

If there exists $p \in P^*$ such that $c(\pi \cup \{p\}) > b$, the maximin support rule terminates and outputs π . Otherwise, a project $p \in P^*$ is selected (π is updated to $\pi \cup \{p\}$) and a new round begins.

Once again, to obtain a resolute rule one needs to break the ties among the projects in any P^* . The irresolute variant is obtained by considering all possible ways of breaking such ties.

Note that in their definition, [Aziz, Lee and Talmon \(2018\)](#) provide a linear program to compute efficiently the optimum load distribution in each round.

Interestingly, we know from the multi-winner voting literature that MAXIMIN-SUPP provides approximation guarantees (to the optimum load distribution) that SEQPHRAG does not ([Cevallos and Stewart, 2021](#)). This makes it a rule that deserves further investigation.

3.2.4 The Method of Equal Shares

Let us now discuss MES, a rule we introduced in Section 2.3.3. This rule can also be used with profiles of cardinal ballots, we provide the definition below.

Definition 3.2.3 (Method of Equal Shares for Cardinal Ballots). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile \mathbf{A} of cardinal ballots, the method of equal shares, MES, constructs a budget allocation π , initially empty, iteratively as follows.*

A load $\ell_i : 2^P \rightarrow \mathbb{R}_{\geq 0}$, is associated with every agent $i \in \mathcal{N}$, initialised as $\ell_i(\emptyset) = 0$ for all $i \in \mathcal{N}$. The load represents how much virtual money the agents have spent.

Given π and a scalar $\alpha \geq 0$, the contribution of agent $i \in \mathcal{N}$ for project $p \in \mathcal{P} \setminus \pi$ is defined by:

$$\gamma_i(\pi, \alpha, p) = \min(b/n - \ell_i(\pi), \alpha \cdot A_i(p)).$$

This is the amount i would pay to buy project p for a given α . Note that the above means that agents are initially provided b/n units of the virtual currency.

Given a budget allocation π , a project $p \in \mathcal{P} \setminus \pi$ is α -affordable, for $\alpha \geq 0$, if

$$\sum_{i \in \mathcal{N}} \gamma_i(\pi, \alpha, p) \geq c(p).$$

A project p is thus α -affordable if all the agents can contribute enough to afford p for the given α .

At a given round with current budget allocation π , if no project is α -affordable for any α , MES terminates.

Otherwise, it selects a project $p \in \mathcal{P} \setminus \pi$ that is α^* -affordable where α^* is the smallest α such that one project is α -affordable (π is updated to $\pi \cup \{p\}$). The agents' loads are then updated: $l_i(\pi \cup \{p\}) = l_i(\pi) + \gamma_i(\pi, \alpha, p)$. A new round then starts.

The above definition gives rise to a resolute rule (when ties among α^* -affordable projects are broken arbitrarily). For the irresolute variant of the rule, one would need to consider all ways to break the ties between α^* -affordable projects at each round.

We observe that MES does not necessarily spend the whole budget, *i.e.*, it is not *exhaustive*. Indeed, it is possible that no project is α -affordable for any α , in which case MES returns the empty set. For this reason, in practice MES nearly always needs to be combined with a completion method. We discuss this point in more detail in Section 3.4.1.

3.2.5 Other Rules for Participatory Budgeting

We have introduced what we believe to be the most prominent rules in the literature for PB. These are obviously not the only ones that have been defined. We briefly comment on other rules.

In the multi-winner literature, Thiele methods play an important role (Lackner and Skowron, 2023). It can thus be surprising that this is not the case in the PB setting. It turns out that these rules do not behave as nicely in PB as they did in multi-winner voting. In particular, *Proportional Approval Voting* (PAV) that provides interesting proportionality guarantees in multi-winner voting (Aziz, Brill, Conitzer, Elkind, Freeman and Walsh, 2017), no longer enjoys them in the PB setting as observed by Peters, Pierczyński and Skowron (2021) and Los, Christoff and Grossi (2022).

Among the other rules that have been defined, Skowron, Slinko, Szufa and Talmon (2020) propose an adaptation of the multi-winner variant of the *Single Transferable Vote* rule (STV) in the PB setting with cumulative ballots.

When considering ordinal ballots, Aziz and Lee (2021) introduced the *expanding-approvals rule* for PB. Peters, Pierczyński and Skowron (2021) proposed an ordinal version of MES, showing that it is an expanding-approvals rule.¹⁷

3.3 Fairness in Indivisible Participatory Budgeting

Throughout this section we will study different PB rules in terms of their fairness guarantees. This represents the largest share of the literature devoted to the analysis of PB rules and has proved to be a rich and fruitful research direction.

This section is mainly organised around the different types of fairness requirements that have been introduced. We will start with the concepts revolving around

¹⁷This part is only available in the extended version, available at arxiv.org/abs/2008.13276.

justified representation (Section 3.3.1) which will naturally lead us to the idea of the core (Section 3.3.2). We will then discuss the idea of priceability (Section 3.3.3). Broadening our perspective, our next focus will be fairness in ordinal PB (Section 3.3.4), and more generally all the other notions of fairness that have been introduced (Section 3.3.5). Fairness in extended models of PB will then be discussed (Section 3.3.6). We will conclude by attempting to unify everything by drawing taxonomies linking the different requirements to each other (Section 3.3.7).

3.3.1 Extended and Proportional Justified Representation

The main part of the research on fairness in PB focuses on adapting to PB the well studied concept of *justified representation* from the multi-winner voting literature (Aziz, Brill, Conitzer, Elkind, Freeman and Walsh, 2017; Aziz, Elkind, Huang, Lackner, Sanchez-Fernandez and Skowron, 2018; Bredereck, Faliszewski, Kaczmarczyk and Niedermeier, 2019; Peters and Skowron, 2020; Lackner and Skowron, 2023). The idea behind justified representation is that groups of agents that are large enough and similar enough—the so-called *cohesive groups*—deserve to be satisfied with a fraction of the outcome that is proportional to their size.

In the following we will define the most important concepts related to justified representation and present the main results from the literature. Figures 3.3.1 and 3.3.2 summarise (among others) the results presented in this section.

Justified Representation with Cardinal Ballots

We first consider the general case of cardinal ballots. A special focus on approval ballots will come later.

Cohesive groups for cardinal ballots. Let us start with the definition of cohesive groups. We follow the definition of Peters, Pierczyński and Skowron (2021).

Definition 3.3.1 ((α, P) -Cohesive Groups). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile \mathbf{A} of cardinal ballots, a non-empty group of agents $N \subseteq \mathcal{N}$ is said to be (α, P) -cohesive, for a function $\alpha : \mathcal{P} \rightarrow [0, 1]$ and a set of projects $P \subseteq \mathcal{P}$, if the following two conditions are satisfied:*

- ▶ $\alpha(p) \leq A_i(p)$ for all $i \in N$ and $p \in P$, that is, α is lower-bounding the score of the agents in N for the projects in P ;
- ▶ $\frac{|N|}{n} \cdot b \geq c(P)$, that is, N 's share of the budget is enough to afford P .

Overall, for any (α, P) -cohesive group of agents $N \subseteq \mathcal{N}$, it should be that N is (i) large enough to afford the projects in P , and, (ii) cohesive enough to “deserve” the satisfaction they receive from the projects in P , measured by α .

We will make use of one specific function α denoted by α^{\min} and defined for any profile \mathbf{A} and subset of agents $N \subseteq \mathcal{N}$ as:

$$\alpha_{N,\mathbf{A}}^{\min}(p) = \min_{i \in N} A_i(p), \text{ for all } p \in \mathcal{P}.$$

This function simply takes the minimum score submitted by any agent in N for project p . Note that for every group of agents $N \subseteq \mathcal{N}$ and subset of projects $P \subseteq \mathcal{P}$, if $|N|/n \cdot b \geq c(P)$ then N is $(\alpha_{N,\mathbf{A}}^{\min}, P)$ -cohesive.

Extended justified representation for cardinal ballots. We want to ensure that cohesive groups receive what they deserve. But what exactly do cohesive groups deserve? Consider an (α, P) -cohesive group N . Since agents in N control enough share of the budget to afford P , the most natural idea would be to guarantee all agents in N at least as much satisfaction as they all agree P would offer them (captured by α). This idea is captured by the following axiom: *strong extended justified representation*.¹⁸

Definition 3.3.2 (Strong Extended Justified Representation). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ a profile \mathbf{A} of cardinal ballots, a budget allocation $\pi \in \text{FEAS}(I)$ is said to satisfy strong extended justified representation (Strong-EJR) if for all $P \subseteq \mathcal{P}$, all $N \subseteq \mathcal{N}$ that is $(\alpha_{N,\mathbf{A}}^{\min}, P)$ -cohesive groups, and all $i \in N$, we have:*

$$\sum_{p \in \pi} A_i(p) \geq \sum_{p \in P} \min_{i \in N} A_i(p).$$

Remember that when using cardinal ballots, we made the assumption that the satisfaction of an agent behaves additively, so the left-hand side of the inequality above represents the agent's satisfaction.

Even though Strong-EJR is quite appealing (or at least somewhat natural), it is unsatisfiable in general. This was already observed in multi-winner voting (Aziz, Brill, Conitzer, Elkind, Freeman and Walsh, 2017).

Given this impossibility, the focus is usually put on (simple) *extended justified representation* (Aziz, Brill, Conitzer, Elkind, Freeman and Walsh, 2017; Peters, Pierczyński and Skowron, 2021). This is a weakening of Strong-EJR requiring only one member of each cohesive group to reach the satisfaction threshold. We thus switch one quantifier from a universal one to an existential one in the definition.

Definition 3.3.3 (Extended Justified Representation). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ a profile \mathbf{A} of cardinal ballots, a budget allocation $\pi \in \text{FEAS}(I)$ is said to satisfy extended justified representation (EJR) if for all $P \subseteq \mathcal{P}$ and all $N \subseteq \mathcal{N}$ that is $(\alpha_{N,\mathbf{A}}^{\min}, P)$ -cohesive groups, there exists $i \in N$ such that:*

$$\sum_{p \in \pi} A_i(p) \geq \sum_{p \in P} \min_{i \in N} A_i(p).$$

¹⁸Note here that we slightly differ from the definition of Peters, Pierczyński and Skowron (2021). Indeed, in the definition of Strong-EJR (and EJR) they consider any (α, P) -cohesive group while we only use a specific α , namely α^{\min} . The two definitions are however equivalent and we believe this one to be clearer since it requires one less universal quantifier.

The first thing to note is that EJR does not suffer the same drawback as Strong-EJR: it can always be satisfied.

Theorem 3.3.4 (Peters, Pierczyński and Skowron 2021). *For every instance I , there exists a budget allocation $\pi \in \text{FEAS}(I)$ that satisfies EJR.*

What Peters, Pierczyński and Skowron (2021) actually prove is that a *greedy cohesive rule*¹⁹ always returns a feasible budget allocation that satisfies EJR (it even satisfies *full justified representation*, see Section 3.3.5). This rule is interesting at a theoretical level but is quite artificial and thus not really appealing at a practical level. One of its main drawbacks is that it runs in exponential time. This however, seems to be unavoidable to satisfy EJR.

Theorem 3.3.5 (Peters, Pierczyński and Skowron 2021). *There is no strongly-polynomial time algorithm that computes, given an instance I and a profile \mathbf{A} of cardinal ballots, a budget allocation $\pi \in \text{FEAS}(I)$ satisfying EJR unless $\text{P} = \text{NP}$, even if there is only one voter.*

Interestingly, the hardness proof shows that the running time of an algorithm finding an EJR budget allocation has to be exponential in $\log(b)$, while the greedy cohesive rule mentioned above runs in time exponential in n , the number of voters. Closing this gap is an interesting open problem.

Let us quickly mention another computational result: checking whether a given budget allocation satisfies EJR is a coNP-complete problem. This is because it was already the case in the unit-cost setting with approval ballots (Aziz, Brill, Conitzer, Elkind, Freeman and Walsh, 2017).

In the hope of achieving polynomial-time computability, a relaxation of EJR has been introduced: EJR *up to one project* (EJR-1). It relaxes EJR in the following way: for each cohesive group N , it can be the case that no agent in N enjoys enough satisfaction, but, at least one agent would reach the desired level of satisfaction if we were to select an additional project. This concept can be interpreted as requiring that the satisfaction can only be at most one project away from the objective.

Definition 3.3.6 (Extended Justified Representation up to One Project). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ a profile \mathbf{A} of cardinal ballots, a budget allocation $\pi \in \text{FEAS}(I)$ is said to satisfy extended justified representation up to one project (EJR-1) if for all $P \subseteq \mathcal{P}$ and all $N \subseteq \mathcal{N}$ that is $(\alpha_{N,\mathbf{A}}^{\min}, P)$ -cohesive groups, there exists $i \in N$ such that either $\sum_{p \in \pi} A_i(p) \geq \sum_{p \in P} \min_{i \in N} A_i(p)$, or for some $p^* \in P \setminus \pi$ we have:*

$$A_i(p^*) + \sum_{p \in \pi} A_i(p) > \sum_{p \in P} \min_{i \in N} A_i(p).$$

¹⁹The idea behind a greedy cohesive rule is to consider all cohesive groups and to greedily select sets of projects P for which there is a suitable (α, P) -cohesive group with “maximum” α . This defines a general scheme to devise procedures as the notion of “suitable cohesive group” differs depending on the goal. Such procedures have notably been defined and used by Aziz, Lee and Talmon (2018) and Peters, Pierczyński and Skowron (2021).

One might be surprised by the strict inequality in the definition above. It is there for technical reasons: It ensures that EJR and EJR-1 coincide in the unit-cost setting when used with approval ballots.²⁰ It also has interesting consequences in terms of the fairness properties that EJR-1 implies.²¹

One of the main results from [Peters, Pierczyński and Skowron \(2021\)](#) is that MES does satisfy EJR-1. Given that MES runs in strongly-polynomial-time, this shows that a budget allocation satisfying EJR-1 can always be computed in polynomial time.

Theorem 3.3.7 ([Peters, Pierczyński and Skowron 2021](#)). *For every instance I and profile A of cardinal ballots, $\text{MES}(I, A)$ satisfies EJR-1.*

Proportional justified representation for cardinal ballots. Going down the list of weakenings of Strong-EJR, we have now reached proportional justified representation (PJR) ([Sánchez-Fernández, Elkind, Lackner, Fernández, Fisteus, Val and Skowron, 2017](#)). While EJR required at least one member of each cohesive group to enjoy the required satisfaction, PJR requires the group altogether—acting as a single agent—to achieve the required satisfaction. We provide below the definition from [Los, Christoff and Grossi \(2022\)](#) who defined it for PB with cardinal ballots.

Definition 3.3.8 (Proportional Justified Representation). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ a profile A of cardinal ballots, a budget allocation $\pi \in \text{FEAS}(I)$ is said to satisfy proportional justified representation (PJR) if for all $P \subseteq \mathcal{P}$ and all $N \subseteq \mathcal{N}$ that is $(\alpha_{N,A}^{\min}, P)$ -cohesive groups N we have:*

$$\sum_{p \in \pi} \max_{i \in N} A_i(p) \geq \sum_{p \in P} \min_{i \in N} A_i(p).$$

It should be quite obvious that any budget allocation satisfying EJR also satisfies PJR. From this, we can derive that [theorems 3.3.4 and 3.3.5](#) also apply to PJR. Specifically, we know that (i) for every instance, there exists a feasible budget allocation that satisfies PJR, and (ii) there exists no polynomial-time algorithm computing PJR budget allocations unless $P = NP$. The second point holds because PJR and EJR coincide when there is only a single agent and [Theorem 3.3.5](#) holds for one-agent profiles.

Interestingly, the problem of checking whether a budget allocation satisfies PJR or not is still coNP-complete (remember that this was already the case for EJR), and that, already on unit-cost instances with approval ballots ([Aziz, Elkind, Huang, Lackner, Sanchez-Fernandez and Skowron, 2018](#)).

²⁰EJR and EJR-1 do not coincide in the unit-cost setting with generic cardinal ballots as presented by [Peters, Pierczyński and Skowron \(2021\)](#) in Footnote 8 of the ArXiv version.

²¹Notably, having a strict inequality ensures that EJR-1 implies a property that could be called basic proportionality, which requires that if for a group of agents N there exists a $P \subseteq \mathcal{P}$ such that $|N|/n \cdot b \geq c(P)$ and $A_i(p) = A_j(p) > 0$ if and only if $p \in P$ for all $i, j \in N$, then P must be selected. This is not the case if EJR-1 is defined with a weak inequality.

Los, Christoff and Grossi (2022) show that a PB adaption of the PAV rule fails to satisfy PJR. This might come as a surprise since PAV satisfies EJR in the case of multi-winner voting elections.

This last property concludes our section on cardinal ballots.

Justified Representation with Approval Ballots

All we presented above for cardinal ballots also applies in the case of approval ballots. However, since approval ballots are a special case of cardinal ballots, the definitions can be simplified and stronger results can be obtained.

Cohesive groups for approval ballots. With cardinal ballots we had to introduce the α parameter to the definition of a cohesive group, since agents could assign different scores to the projects. This is not necessary with approval ballots.

Definition 3.3.9 (*P*-cohesive groups). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile \mathbf{A} of approval ballots, a non-empty group of agents $N \subseteq \mathcal{N}$ is said to be *P*-cohesive, for a set of projects $P \subseteq \mathcal{P}$, if the following two conditions are satisfied:*

- ▶ for all $p \in P$ and $i \in N$, $A_i(p) = 1$, that is, every agent in N approves all projects in P ;
- ▶ $\frac{|N|}{n} \cdot b \geq c(P)$, that is, N 's share of the budget is enough to afford P .

Remember the interpretation we had of cohesive groups: these are groups of agents that deserve some satisfaction in the final outcome. When using approval ballots, we will use satisfaction functions as introduced in Definition 2.2.1 as measures of satisfaction.

Extended justified representation for approval ballots. We are now ready to introduce concepts based on justified representation for approval ballots. Note that they are all parameterised by a satisfaction function. We start with Strong-EJR.

Definition 3.3.10 (Strong-EJR for Approval Ballots). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ a profile \mathbf{A} of approval ballots, and a satisfaction function sat , a budget allocation $\pi \in \text{FEAS}(I)$ is said to satisfy strong extended justified representation for sat (Strong-EJR[sat]) if for all $P \subseteq \mathcal{P}$ and all *P*-cohesive groups N , we have $sat_i(\pi) \geq sat_i(P)$ for all agents $i \in N$.*

As for cardinal ballots, Strong-EJR[sat] is appealing, but not satisfiable in general.²²

²²We are not aware of this result existing in the literature. The proof is rather simple. It relies on a counterexample using three projects p_1, p_2 , and p_3 , all of cost 1. The budget limit is 2. There are four agents with the following ballots: Agent 1 approves only of p_1 , agent 2 approves only of p_2 , agent 3 approves only of p_3 , and agent 4 approves of p_1, p_2 and p_3 . Recall that we assume for any satisfaction function sat that $sat(P) = 0$ if and only if $P = \emptyset$. Therefore, the only way to satisfy Strong-EJR[sat] is to select p_1, p_2 , and p_3 , which is not possible with $b = 2$.

Proposition 3.3.11. *For any satisfaction function sat , there exists an instance I such that no budget allocation $\pi \in \text{FEAS}(I)$ satisfies Strong-EJR[sat].*

EJR can also be adapted quite naturally to this setting and can be shown to be always satisfiable in exponential time (using some greedy cohesive rule).

Definition 3.3.12 (Extended Justified Representation for Approval Ballots). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ a profile \mathbf{A} of approval ballots, and a satisfaction function sat , a budget allocation $\pi \in \text{FEAS}(I)$ is said to satisfy extended justified representation for sat (EJR[sat]) if for all $P \subseteq \mathcal{P}$ and all P -cohesive groups N , there exists $i \in N$ such that $sat_i(\pi) \geq sat_i(P)$.*

Theorem 3.3.13 (Brill, Forster, Lackner, Maly and Peters 2023). *For every instance I and every satisfaction function sat , there exists a budget allocation $\pi \in \text{FEAS}(I)$ that satisfies EJR[sat].*

Unfortunately, for large classes of satisfaction functions, it is not possible to satisfy EJR in polynomial time.

Theorem 3.3.14 (Brill, Forster, Lackner, Maly and Peters 2023). *Let I be an instance and sat be a satisfaction function such that for all $P, P' \subseteq \mathcal{P}$ such that $c(P) < c(P')$ we have $sat(P) < sat(P')$. Then, there is no algorithm running in strongly-polynomial-time that computes, given an instance I and a profile \mathbf{A} of approval ballots, a budget allocation $\pi \in \text{FEAS}(I)$ satisfying EJR-1[sat] unless $\text{P} = \text{NP}$, even if there is only one voter.*

It is important to note that sat^{card} is not captured by the above statement, and indeed, budget allocations satisfying EJR[sat^{card}] can always be computed in polynomial time using $\text{MES}[sat^{card}]$ (Peters, Pierczyński and Skowron, 2021; Los, Christoff and Grossi, 2022). This is because for sat^{card} , EJR[sat^{card}] and EJR-1[sat^{card}] coincide (the latter is defined below).

For the same reasons as when we were considering cardinal ballots, checking whether a budget allocation satisfies EJR or not is a coNP problem.

EJR-1 can also be adapted quite naturally to the approval setting. Remember that Peters, Pierczyński and Skowron (2021) proved that MES always returns a budget allocation satisfying EJR-1. Since additive satisfaction functions can be interpreted as cardinal ballots, one can always compute an EJR-1[sat] budget allocations, for an additive satisfaction function sat , by running MES with the cardinal ballots corresponding to sat . In the approval setting, we can go further and get the same result for EJR up to any project.

Definition 3.3.15 (EJR up to Any Project for Approval Ballots). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ a profile \mathbf{A} of approval ballots, and a satisfaction function sat , a budget allocation $\pi \in \text{FEAS}(I)$ is said to satisfy extended justified representation up to any project for sat (EJR-X[sat]) if for all $P \subseteq \mathcal{P}$ and all P -cohesive groups N , there exists $i \in N$ such that for all $p^* \in P \setminus \pi$ we have $sat_i(\pi \cup \{p^*\}) > sat_i(P)$.*

EJR-X is a strengthening of EJR-1 that uses a universal quantifier on the project that bounds the distance between the justified and the actual satisfaction of an agent, instead of an existential one.

One of the main results from [Brill, Forster, Lackner, Maly and Peters \(2023\)](#) is that for a natural class of satisfaction functions, namely decreasing normalised satisfaction functions, the outcome of $\text{MES}[sat]$ always satisfies EJR-X[sat].

Definition 3.3.16 (DNS Function). *We say a satisfaction function sat has weakly decreasing normalised satisfaction (DNS) if for all projects $p, p' \in \mathcal{P}$ with $c(p) \leq c(p')$, we have:*

$$sat(p) \leq sat(p') \quad \text{and} \quad \frac{sat(p)}{c(p)} \geq \frac{sat(p')}{c(p')}.$$

In this case, we call sat a DNS function.

DNS functions ensure that more expensive projects are (weakly) better than cheaper ones; and that more expensive projects do not provide more satisfaction per cost than cheaper ones. Crucially, sat^{cost} and sat^{card} are DNS functions.

Theorem 3.3.17 ([Brill, Forster, Lackner, Maly and Peters 2023](#)). *Let sat be a DNS function. Then, for any instance I and profile \mathbf{A} of approval ballots, $\text{MES}[sat](I, \mathbf{A})$ satisfies EJR-X[sat].*

Before moving on to PJR for approval ballots, let us touch on an additional topic. [Fairstein, Vilenchik, Meir and Gal \(2022\)](#) study the consequences of satisfying EJR on measures of social welfare and representation.²³ Specifically, they provide bounds on the social welfare and representation guarantees of rules satisfying EJR[sat^{card}]. In other words, they compare the maximally achievable social welfare with respect to sat^{card} and sat^{CC} to the social welfare achieved by rules satisfying EJR.²⁴

Theorem 3.3.18 ([Fairstein, Vilenchik, Meir and Gal 2022](#)). *Let F be a PB rule that satisfies EJR[sat^{card}]. Then for any instance $I = \langle \mathcal{P}, c, b \rangle$ and profile \mathbf{A} of approval ballots, we have:*

$$\frac{c_{\min}}{n \cdot b} \left\lfloor \frac{b}{c_{\max}} \right\rfloor \leq \frac{\sum_{i \in \mathcal{N}} sat_i^{card}(F(I, \mathbf{A}))}{\max_{\pi \in \text{FEAS}(I)} \sum_{i \in \mathcal{N}} sat_i^{card}(\pi)} \leq \frac{4}{\sqrt{n}} - \frac{1}{n},$$

where $c_{\min} = \min_{p \in \mathcal{P}} c(p)$ and $c_{\max} = \max_{p \in \mathcal{P}} c(p)$.

²³[Fairstein, Vilenchik, Meir and Gal \(2022\)](#) also perform a similar analysis for specific rules. However, these rules are not part of the standard set of rules we study in this paper.

²⁴Let us also mention that [Lackner and Skowron \(2020\)](#) studied the same questions in the multi-winner voting setting.

Moreover, for any instance $I = \langle \mathcal{P}, c, b \rangle$ and profile \mathbf{A} of approval ballots, we have:

$$\frac{1}{n} \leq \frac{\sum_{i \in \mathcal{N}} \text{sat}^{CC}(\mathbf{F}(I, \mathbf{A}))}{\max_{\pi \in \text{FEAS}(I)} \sum_{i \in \mathcal{N}} \text{sat}^{CC}(\pi)} \leq \frac{1}{n-1},$$

where the upper bound holds if $n \geq b/c_{\min}$ with c_{\min} defined as above.

Proportional justified representation for approval ballots. Let us now move on to PJR. Three main sets of authors have adapted PJR in the context of PB with approval ballots: [Aziz, Lee and Talmon \(2018\)](#), [Los, Christoff and Grossi \(2022\)](#) and [Brill, Forster, Lackner, Maly and Peters \(2023\)](#).

We first provide the definition of PJR as stated by [Brill, Forster, Lackner, Maly and Peters \(2023\)](#).

Definition 3.3.19 (Proportional Justified Representation for Approval Ballots). *For an instance $I = \langle \mathcal{P}, c, b \rangle$ a profile \mathbf{A} of approval ballots, and a satisfaction function sat , a budget allocation $\pi \in \text{FEAS}(I)$ is said to satisfy proportional justified representation for sat ($\text{PJR}[\text{sat}]$) if for all $P \subseteq \mathcal{P}$ and all P -cohesive groups N , we have:*

$$\text{sat} \left(\bigcup_{i \in N} \{p \in \pi \mid A_i(p) = 1\} \right) \geq \text{sat}(P).$$

Similar adaptations of PJR to the PB setting have also been studied. $\text{PJR}[\text{sat}^{\text{cost}}]$ is equivalent to the BPJR-L property introduced by [Aziz, Lee and Talmon \(2018\)](#).²⁵ [Aziz, Lee and Talmon \(2018\)](#) also defined variants of (B)PJR based on the relative budget, which will be discussed in Section 3.3.5. Finally, $\text{PJR}[\text{sat}^{\text{card}}]$ has been introduced by [Los, Christoff and Grossi \(2022\)](#).

For now, let us focus on $\text{PJR}[\text{sat}]$. It should be clear that for any satisfaction function sat , $\text{EJR}[\text{sat}]$ implies $\text{PJR}[\text{sat}]$. Thus, for any instance I and profile \mathbf{A} of approval ballots, there exists a budget allocation satisfying $\text{PJR}[\text{sat}]$, however for a large class of satisfaction functions, it cannot be computed in polynomial time (see Theorem 3.3.14 for the exact condition on the satisfaction function). Finally, checking $\text{PJR}[\text{sat}]$ is coNP-complete for any sat that is neutral with respects to projects with the same cost, and that already holds in the unit-cost setting.

As we did for EJR, we can then introduce PJR-X.

Definition 3.3.20 (PJR up to Any Project for Approval Ballots). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ a profile \mathbf{A} of approval ballots, and a satisfaction function sat , a budget allocation $\pi \in \text{FEAS}(I)$ is said to satisfy proportional justified representation up to*

²⁵Note that the definition of BPJR-L proposed by [Aziz, Lee and Talmon \(2018\)](#) looks more involved than $\text{PJR}[\text{sat}^{\text{cost}}]$ as they do not use the notion of cohesive groups. Close inspection should convince the reader that these two definitions are equivalent.

any project for sat ($PJR\text{-}X[sat]$) if for all $P \subseteq \mathcal{P}$ and all P -cohesive groups N and any $p^* \in P \setminus \pi$, we have:

$$sat \left(\{p^*\} \cup \bigcup_{i \in N} \{p \in \pi \mid A_i(p) = 1\} \right) > sat(P).$$

Remember that we know for any DNS function sat (Definition 3.3.16), that $EJR\text{-}X[sat]$ can be satisfied (Theorem 3.3.17). Since $PJR\text{-}X[sat]$ is implied by $EJR\text{-}X[sat]$, this result also applies to $PJR\text{-}X[sat]$. Brill, Forster, Lackner, Maly and Peters (2023) actually prove something stronger: $PJR\text{-}X[sat]$ can be satisfied *simultaneously* for every DNS function sat .

Theorem 3.3.21 (Brill, Forster, Lackner, Maly and Peters 2023). *Let I be an instance and \mathbf{A} a profile. Then, $SEQPHRAG(I, \mathbf{A})$, $MAXIMINSUPP(I, \mathbf{A})$ and $MES[sat^{card}](I, \mathbf{A})$ satisfy $PJR\text{-}X[sat]$ for all DNS functions sat simultaneously.*

Interestingly, Brill, Forster, Lackner, Maly and Peters (2023) actually proved that this result holds for all rules satisfying a certain strengthening of priceability, as we will see later on (in Section 3.3.3).

This result is rather far-reaching given its generality. Note that it generalises the result of Los, Christoff and Grossi (2022) who prove that $SEQPHRAG$ satisfies $PJR\text{-}1[sat^{card}]$. It also generalises the result of Aziz, Lee and Talmon (2018) that $MAXIMINSUPP$ satisfies a property called $Local\text{-}BPJR\text{-}L[sat^{cost}]$ as explained below.

Finally, note that this result cannot be generalised to $EJR\text{-}X$, as Brill, Forster, Lackner, Maly and Peters (2023) show that there are instances where $EJR\text{-}1[sat^{cost}]$ and $EJR\text{-}1[sat^{card}]$ are incompatible.

Before Brill, Forster, Lackner, Maly and Peters (2023) introduced their definition of PJR parameterised by a satisfaction function, Aziz, Lee and Talmon (2018) defined PJR for PB with approval ballots. As we have mentioned before, they introduced an axiom called $BPJR\text{-}L$ —that is equivalent to $PJR[sat^{cost}]$ —and proved that budget allocations satisfying it could not be found in polynomial time (unless $P = NP$). Due to this observation, they introduced $Local\text{-}BPJR\text{-}L$, a weakening of $PJR[sat^{cost}]$. Let us provide the definition of this axiom. Note that we use here the definition of Brill, Forster, Lackner, Maly and Peters (2023) who extended it to work with arbitrary satisfaction functions. The original definition of Aziz, Lee and Talmon (2018) would correspond to $Local\text{-}BPJR\text{-}L[sat^{cost}]$.

Definition 3.3.22 (Local Budget PJR for the Budget Limit). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ a profile \mathbf{A} of approval ballots, and a satisfaction function sat , a budget allocation $\pi \in FEAS(I)$ is said to satisfy $Local\text{-}BPJR\text{-}L[sat]$ if for all $P \subseteq \mathcal{P}$ and all P -cohesive groups N , it is the case that for every $P^* \subseteq P$ such that $\{p \in \pi \mid \exists i \in N, A_i(p) = 1\} \not\subseteq P^*$ we have:*

$$P^* \notin \arg \max_{\substack{P' \subseteq \{p \in \mathcal{P} \mid \forall i \in N, A_i(p) = 1\} \\ c(P') \leq c(P)}} sat(P').$$

One of the main results of [Aziz, Lee and Talmon \(2018\)](#) is that `MAXIMINSUPP` satisfies `Local-BPJR-L[sat^{cost}]`. Later on, [Brill, Forster, Lackner, Maly and Peters \(2023\)](#) explored further the relationship between different properties and proved that any budget allocation satisfying `PJR-X[sat]` also satisfies `Local-BPJR-L[sat]` (so `SEQPHRAG`, `MAXIMINSUPP` and `MES[sat^{cost}]` all satisfy `Local-BPJR-L[sat^{cost}]`). In addition they showed that in the unit-cost case, `Local-BPJR-L` does not coincide with `PJR`, while `PJR-X` does.

It is also worth mentioning that [Aziz, Lee and Talmon \(2018\)](#) also introduced another axiom called *Strong-BPJR-L*. It is satisfied by a budget allocation π if for every $\ell \in [1, b]$, and for every group of voters N that controls ℓ units of budget, i.e., $|N|/n \cdot b \geq \ell$, and that unanimously approve projects of total cost more than ℓ , i.e., $c(\{p \in \mathcal{P} \mid \forall i \in N, A_i(p) = 1\}) \geq \ell$, we have $c(\bigcup_{i \in N} \{p \in \pi \mid A_i(p) = 1\}) \geq \ell$. Because of the indivisibility of the projects, this axiom cannot always be satisfied. Note that this definition implicitly uses the satisfaction function sat^{cost} as the groups of voter claiming ℓ units of budget need to enjoy collectively a cost-satisfaction of at least ℓ . Because of this limited applicability, we chose not to focus on this notion. Note that *Strong-BPJR-L* is a strengthening of `PJR[sat^{cost}]` (which is equivalent to `BPJR-L`) as the condition on the group of agents N is weaker.

3.3.2 The Core

Intuitively, `EJR` guarantees that in every cohesive group there is at least one voter that receives as much satisfaction as the group could guarantee each member if the group could spend their part of the budget as they wish. We now introduce a property that is similar in spirit, called the *core*, though it does not rely on cohesive groups.

The Core with Cardinal Ballots

We start by providing the definition of the core. Note that it was first introduced by [Fain, Goel and Munagala \(2016\)](#) for PB with divisible projects. The definition below, though adapted to the indivisible PB setting, is very similar.

Definition 3.3.23 (The Core with Cardinal Ballots). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile \mathbf{A} of cardinal ballots, a budget allocation $\pi \in \text{FEAS}(I)$ is in the core of I if for every group of voters $N \subseteq \mathcal{N}$ and subset of projects $P \subseteq \mathcal{P}$ such that $|N|/n \geq c(P)/b$, there exists a voter $i^* \in N$ with:*

$$\sum_{p \in \pi} A_{i^*}(p) \geq \sum_{p \in P} A_{i^*}(p).$$

The core can be seen as a kind of stability condition which guarantees that no groups of agents can “deviate” by taking their part of the budget to fund a set of projects P that gives each agent in the group a higher satisfaction than π . The core of PB

is inspired by the concept of the core in cooperative game theory (Scarf, 1967; Fain, Goel and Munagala, 2016), but there is no direct technical link.

Interestingly, EJR can be viewed as a restriction of the core where only cohesive groups are allowed to deviate. Therefore, the core can be seen as a generalisation of EJR to arbitrary groups of agents.

It is known that there are instances where no budget allocation is in the core. In this case, we say that the core of the instance is *empty*. Peters, Pierczyński and Skowron (2021) present an instance with cardinal ballots in the unit-cost setting for which the core is empty. They strengthen a first counter-example provided by Fain, Munagala and Shah (2018) without the unit-cost assumption.²⁶

Proposition 3.3.24 (Peters, Pierczyński and Skowron 2021). *There exists an instance I with unit costs and profile \mathbf{A} of cardinal ballots such that no budget allocation $\pi \in \text{FEAS}(I)$ is in the core, even if for every agent $i \in \mathcal{N}$ and project $p \in \mathcal{P}$ we have $A_i(p) \in \{0, 1, 2\}$.*

Approximating the Core with Cardinal Ballots

We now know that the core can be empty. This raises the question whether it is always possible to find budget allocations that are close to the core. We will present some recent answers to this question below.

We start with a multiplicative approximation to the core as defined by Peters, Pierczyński and Skowron (2021). This approximates the core by bounding the satisfaction the agents would enjoy when deviating.

Definition 3.3.25 (The α -sat Approximate Core with Cardinal Ballots). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$, a profile \mathbf{A} of cardinal ballots, and a scalar $\alpha \geq 1$, a budget allocation $\pi \in \text{FEAS}(I)$ is in the α -sat approximate core of I if for every group of voters $N \subseteq \mathcal{N}$ and subset of projects $P \subseteq \mathcal{P}$ such that $|N|/n \geq c(P)/b$, there exists a voter $i^* \in N$ and a project $p^* \in P$ with:*

$$\sum_{p \in \pi \cup \{p^*\}} A_{i^*}(p) \geq \frac{\sum_{p \in P} A_{i^*}(p)}{\alpha}.$$

Note that the above is actually an additive and multiplicative approximation of the core as an extra project is also added. This follows from the known impossibility of a (simply) multiplicative approximation of the core (Fain, Munagala and Shah, 2018; Cheng, Jiang, Munagala and Wang, 2020; Munagala, Shen, Wang and Wang, 2022).

Using the above definition of an approximation of the core, Peters, Pierczyński and Skowron (2021) showed that MES is never too far from the core.

²⁶This counterexample is described in the appendix on endowment-based core of Fain, Munagala and Shah (2018), available at arxiv.org/abs/1805.03164.

Theorem 3.3.26 (Peters, Pierczyński and Skowron 2021). Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile \mathbf{A} of cardinal ballots, let u_{\max} and u_{\min} be the highest and lowest possible satisfaction of a voter, defined as:

$$u_{\min} = \min_{i \in \mathcal{N}} \min_{\substack{\pi \in \text{FEAS} \\ \exists p \in \pi, A_i(p) > 0}} \sum_{p \in \pi} A_i(p) \quad \text{and} \quad u_{\max} = \max_{i \in \mathcal{N}} \max_{\pi \in \text{FEAS}(I)} \sum_{p \in \pi} A_i(p).$$

Then, $\text{MES}(I, \mathbf{A})$ is in the α -sat approximate core of I for $\alpha = 4 \log(2 \cdot u_{\max}/u_{\min})$.

The previous result shows that the $\mathcal{O}(\log(|\text{FEAS}(I)|))$ -sat approximate core is always non-empty for any instance I , and that a suitable budget allocation can be found in polynomial-time. Munagala, Shen, Wang and Wang (2022) extend this result by showing that the $\mathcal{O}(1)$ -sat approximate core is always non-empty, but it is unknown whether the corresponding budget allocation can be found in polynomial-time.

Theorem 3.3.27 (Munagala, Shen, Wang and Wang 2022). For every instance I and profile \mathbf{A} of cardinal ballots, the 9.27-sat approximate core is always non-empty.

This result is obtained by some rather intricate rounding of fractional budget allocations. Note that Munagala, Shen, Wang and Wang (2022) also obtain results for non-additive cardinal ballots. These results are out of the scope of this survey.

Let us now delve into a second type of approximation of the core that has been introduced: entitlement approximation. The idea here is that deviations of coalitions of voters would not be possible if we were to scale down their entitlement (which is equal to b/n in the definition of the core). We provide the definition of Jiang, Munagala and Wang (2020) below.

Definition 3.3.28 (The α -Entitlement Approximate Core with Cardinal Ballots). Given an instance $I = \langle \mathcal{P}, c, b \rangle$, a profile \mathbf{A} of cardinal ballots, and a scalar $\alpha \geq 1$, a budget allocation $\pi \in \text{FEAS}(I)$ is in the α -entitlement approximate core of I if for every group of voters $N \subseteq \mathcal{N}$ and subset of projects $P \subseteq \mathcal{P}$ such that $|N|/n \geq \alpha \cdot c(P)/b$, there exists a voter $i^* \in N$ with:

$$\sum_{p \in \pi} A_{i^*}(p) \geq \sum_{p \in P} A_{i^*}(p).$$

By suitable rounding of lotteries over budget allocation, Jiang, Munagala and Wang (2020) show that the $\mathcal{O}(1)$ -entitlement approximate core is always non-empty.

Theorem 3.3.29 (Jiang, Munagala and Wang 2020). For every instance I and profile \mathbf{A} of cardinal ballots, the 32-entitlement approximate core is always non-empty.

Using the above definition of approximate core, Munagala, Shen and Wang (2022) studied the problem of *core auditing* in PB. This is the computational problem that seeks, given an instance I , a profile \mathbf{A} of cardinal ballots and a budget allocation $\pi \in \text{FEAS}(I)$, what is the largest α such that π is *not* in the α -entitlement approximate core. For this problem, Munagala, Shen and Wang (2022) prove different hardness results, including hardness of approximation, and also provide an approximation algorithm.

The Core with Approval Ballots

When we turn to approval ballots, the picture is quite different: We do not know if the core is always non-empty or not, even for unit-cost instances. This is actually one of the main open problems in the literature on multi-winner voting (Lackner and Skowron, 2023).

For the sake of completeness, we provide below the definition of the core with approval ballots.

Definition 3.3.30 (The Core with Approval Ballots). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$, a profile \mathbf{A} of approval ballots and a satisfaction function sat , a budget allocation $\pi \in \text{FEAS}(I)$ is in the core[sat] of I for sat if for every group of voters $N \subseteq \mathcal{N}$ and subset of projects $P \subseteq \mathcal{P}$ such that $|N|/n \geq c(P)/b$, there exists a voter $i^* \in N$ with:*

$$sat_{i^*}(\pi) \geq sat_{i^*}(P).$$

The question of whether we can always find a budget allocation in the core[sat] is open, even for sat^{card} and sat^{cost} .

3.3.3 Priceability

The next property on our agenda is *priceability*. The idea is that voters have access to a virtual currency, and, if, by following simple rules, they can use their money to fund a given budget allocation, then the latter will be called *priceable*. All voters receive the same amount of virtual currency initially. In that sense, priceability is a proportionality requirement as the power to influence the outcome is shared equally among the voters. It can also be seen as an explainability requirement: a priceable budget allocation is an outcome that could have been obtained if the process had been run as a market.

The initial definition of priceability—in the context of multi-winner voting—is due to Peters and Skowron (2020). We present below the adaptation of this definition to the context of PB proposed by Peters, Pierczyński and Skowron (2021) for PB with cardinal ballots.²⁷

Definition 3.3.31 (Priceability for Cardinal Ballots). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile \mathbf{A} of cardinal ballots, a budget allocation π satisfies priceability, or is priceable, if there exists an entitlement $\alpha \in \mathbb{R}_{\geq 0}$ and a collection $(\gamma_i)_{i \in \mathcal{N}}$ of contribution functions $\gamma_i : \mathcal{P} \rightarrow [0, \alpha]$ such that all of the following conditions are satisfied.*

C1: *If $\gamma_i(p) > 0$ then $A_i(p) > 0$ for all $p \in \mathcal{P}$ and $i \in \mathcal{N}$: Agents only contribute to projects they derive satisfaction from.*

²⁷Note that we changed the terminology to avoid using the terms “budget” and “price”, which can be confused with the basic elements of an instance. This avoids sentences such as “ π is priceable for a budget $B \geq b$ ”.

- C2:** If $\gamma_i(p) > 0$ then $p \in \pi$ for all $p \in \mathcal{P}$ and $i \in \mathcal{N}$: Only projects in π receive contributions.
- C3:** $\sum_{p \in \mathcal{P}} \gamma_i(p) \leq \alpha$ for all $i \in \mathcal{N}$: The total contribution of an agent never exceeds their entitlement α .
- C4:** $\sum_{i \in \mathcal{N}} \gamma_i(p) = c(p)$ for all $p \in \pi$: The projects in π are receiving sufficient contributions to be funded.
- C5:** $\sum_{i \in \mathcal{N} \mid A_i(p) > 0} (\alpha - \sum_{p \in \mathcal{P}} \gamma_i(p)) \leq c(p)$ for all $p \in \mathcal{P} \setminus \pi$: No group of agents supporting an unselected project p is left with more than $c(p)$.

The pair $\langle \alpha, (\gamma_i)_{i \in \mathcal{N}} \rangle$ is called a price system.

Note that it would be more natural to have a strict inequality in (C5), *i.e.*, to guarantee that no group of agents has enough money left over to afford a project for which each member of the group has positive utility. Unfortunately, this would be impossible to satisfy as it is sometimes necessary to break ties between equally popular projects.

Moreover, in the definition of priceability we only distinguish between assigning a zero score to a project, or a strictly positive score. Therefore, the definition does not change whether cardinal or simply approval ballots are used. Note that this definition of priceability requires the underlying assumption that satisfaction is strictly monotonic.

Priceable Rules

Given the similarities between the definition of priceability and that of MES, it will not surprise anyone that its outcome is always priceable. Maybe more surprisingly, this also is the case for sequential Phragmén and the maximin support rules.

Proposition 3.3.32 (Peters, Pierczyński and Skowron 2021). *For every instance I and profile \mathbf{A} of cardinal ballots, $\text{MES}(I, \mathbf{A})$ is priceable.*

Proposition 3.3.33 (Los, Christoff and Grossi 2022). *For every instance I and profile \mathbf{A} of approval ballots, $\text{SEQPHRAG}(I, \mathbf{A})$ is priceable.*

Proposition 3.3.34 (Brill, Forster, Lackner, Maly and Peters 2023). *For every instance I and profile \mathbf{A} of approval ballots, $\text{MAXIMINSUPP}(I, \mathbf{A})$ is priceable.*

Priceability and PJR

In the context of multi-winner voting, links have been drawn between PJR and priceability (Peters and Skowron, 2020). Brill, Forster, Lackner, Maly and Peters (2023) extend this result for PB with approval ballots. They show that priceability implies $\text{PJR-X}[\text{sat}^{\text{cost}}]$. More importantly, they show that a stronger notion of priceability implies $\text{PJR-X}[\text{sat}]$ for all DNS functions sat (see Definition 3.3.16).

Theorem 3.3.35 (Brill, Forster, Lackner, Maly and Peters 2023). *For every instance $I = \langle \mathcal{P}, c, b \rangle$ and profile \mathbf{A} of approval ballots, consider a budget allocation $\pi \in \text{FEAS}(I)$ that is priceable for a price system $\langle \alpha, (\gamma_i)_{i \in \mathcal{N}} \rangle$ such that $\alpha > b$ and that also satisfies the following additional condition:*

C6: $\sum_{i \in \mathcal{N} \mid A_i(p) > 0} \gamma_i(p') \leq c(p)$ for all $p \in \mathcal{P} \setminus \pi$ and $p' \in \pi$: *No group of agents can save money by jointly shifting their contributions to a project that they all support.*

Then, π satisfies PJR-X[sat] for every DNS function sat .

In particular, Brill, Forster, Lackner, Maly and Peters (2023) show that $\text{MES}[sat^{card}]$, SEQPHRAG and MAXIMINSUPP provide budget allocations that are priceable for their extended notion of priceability, and thus satisfy PJR-X[sat] for every DNS sat .

3.3.4 Proportionality in Ordinal PB

Until now, we have focused on cardinal ballots. In the following we consider ordinal ballots and proportionality requirements for such ballots.

Aziz and Lee (2021) is the main reference here. In their work, they generalise proportionality concepts for multi-winner voting with strict ordinal ballots, to the setting of PB with weak ordinal ballots. These concepts are all based on the idea of solid coalitions, the counterpart of cohesive groups when ballots are ordinal.

Definition 3.3.36 (Solid Coalition). *Let $I = \langle \mathcal{P}, c, b \rangle$ be an instance and consider $\mathbf{A} = (\succsim_i)_{i \in \mathcal{N}}$ a profile of weak ordinal ballots. Given a subset of projects $P \subseteq \mathcal{P}$, a group of voters $N \subseteq \mathcal{N}$ is a P -solid coalition if for all voters $i \in N$ and projects $p \in P$, we have $p \succsim_i p'$ for all $p' \in \mathcal{P} \setminus P$.*

A group of voters N is thus a P -solid coalition if they all prefer the projects in P to the ones outside of P .

Equipped with solid coalitions, Aziz and Lee (2021) define two incomparable generalisations of the *proportionality for solid coalitions* (Dummert, 1984). Before defining them, we introduce new notation. Interpret a weak order \succsim over \mathcal{P} as a vector of indifference classes $\succsim = (P_1, P_2, \dots)$ such that all projects in P_j are preferred to the ones in $P_{j+1} \cup P_{j+2} \cup \dots$. Then, let $\text{top}(\succsim, k)$, for $k \in \mathbb{N}$ be defined as $\text{top}(\succsim, k) = P_1 \cup \dots \cup P_{j^*} \cup P_{j^*+1}$ where $j^* \in \mathbb{N}_{\geq 0}$ is the largest number such that $|\bigcup_{j=1}^{j^*} P_j| < k$.

Definition 3.3.37 (Comparative Proportionality for Solid Coalitions). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and profile $\mathbf{A} = (\succsim_i)_{i \in \mathcal{N}}$ of weak ordinal ballots, a budget allocation $\pi \in \text{FEAS}(I)$ is said to satisfy comparative proportionality for solid coalitions (CPSC) if for every $P \subseteq \mathcal{P}$, there is no P -solid coalition $N \subseteq \mathcal{N}$ for which there exists $P' \subseteq P$ such that:*

$$c(\{p \in \pi \mid \exists i \in N \text{ such that } p \in \text{top}(\succsim_i, |P|)\}) < c(P') \leq \frac{|N|}{n} \cdot b.$$

Definition 3.3.38 (Inclusion Proportionality for Solid Coalitions). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile $\mathbf{A} = (\succsim_i)_{i \in \mathcal{N}}$ of weak ordinal ballots, a budget allocation $\pi \in \text{FEAS}(I)$ is said to satisfy inclusion proportionality for solid coalitions (IPSC) if for every $P \subseteq \mathcal{P}$, there is no P -solid coalition $N \subseteq \mathcal{N}$ for which there exists $p^* \in P \setminus \{p \in \pi \mid \exists i \in N \text{ such that } p \in \text{top}(\succsim_i, |P|)\}$ such that:*

$$c(\{p \in \pi \mid \exists i \in N \text{ such that } p \in \text{top}(\succsim_i, |P|)\}) + c(p^*) \leq \frac{|N|}{n} \cdot b.$$

Aziz and Lee (2021) show that it is not always possible to find budget allocations satisfying CPSC, but that we can always find in polynomial time budget allocations satisfying IPSC.

Theorem 3.3.39 (Aziz and Lee 2021). *There exist an instance I and a profile \mathbf{A} of weak ordinal ballots such that no $\pi \in \text{FEAS}(I)$ satisfies CPSC.*

For every instance I and a profile \mathbf{A} of weak ordinal ballots there exists $\pi \in \text{FEAS}(I)$ that satisfies IPSC. Such a budget allocation can be found in polynomial time.

To conclude, note that Peters, Pierczyński and Skowron (2021) introduce a version of MES working with strict ordinal ballots, that they link to the framework of Aziz and Lee (2021). In particular, they show that it satisfies PSC, a weakening of the properties we defined above.

3.3.5 Other Fairness Requirements

In the following section, we go through other fairness requirements that have been introduced in the literature. Since these are properties that have received less attention, we will go a bit faster on them.

Full Justified Representation

The first axiom we discuss is *full justified representation*. Peters, Pierczyński and Skowron (2021) proposed this strengthening of EJR, which is the strongest axiom based on justified representation that we know can always be satisfied. It strengthens EJR by relaxing the cohesiveness requirement.

Definition 3.3.40 (Full Justified Representation for Cardinal Ballots). *Consider an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile of cardinal ballots \mathbf{A} . A group of voters $N \subseteq \mathcal{N}$ is weakly (β, P) -cohesive for a scalar $\beta \in \mathbb{R}$ and a subset of projects $P \subseteq \mathcal{P}$ if $|N|/n \cdot b \geq c(P)$ and $\sum_{p \in P} A_i(p) \geq \beta$ for every $i \in N$.*

Given I and \mathbf{A} , a budget allocation $\pi \in \text{FEAS}(I)$ satisfies full justified representation (FJR) if for all $P \subseteq \mathcal{P}$, all $\beta \in \mathbb{R}$ and all weakly (β, P) -cohesive groups N , there exists an agent $i \in N$ such that:

$$\sum_{p \in \pi} A_i(p) \geq \beta.$$

Using a greedy cohesive rule, [Peters, Pierczyński and Skowron \(2021\)](#) have been able to show that we can always find a budget allocation satisfying FJR. This rule is however rather artificial. It is an open problem whether there is a polynomial-time rule that satisfies FJR.

Proposition 3.3.41 ([Peters, Pierczyński and Skowron 2021](#)). *For any instance I and profile A of cardinal ballots, there exists a budget allocation $\pi \in \text{FEAS}(I)$ that satisfies FJR.*

Interestingly, this applies even for cardinal ballots over budget allocations, as long as they are monotone.

FJR can be adapted to the world of PB with approval ballots. The definition is provided below.

Definition 3.3.42 (Full Justified Representation for Approval Ballots). *Consider an instance $I = \langle \mathcal{P}, c, b \rangle$, a profile of approval ballots A and a satisfaction function sat . A budget allocation $\pi \in \text{FEAS}(I)$ satisfies full justified representation for sat (FJR[sat]) if for every group of voters $N \subseteq \mathcal{N}$ and subset of projects $P \subseteq \mathcal{P}$ such that $|N|/n \cdot b \geq c(P)$, there exists $i \in N$ for whom:*

$$\text{sat}_i(\pi) \geq \text{sat}_i(P).$$

Because [Peters, Pierczyński and Skowron \(2021\)](#) prove that FJR can be satisfied even for monotonic cardinal ballots over budget allocations, FJR[sat] can be satisfied for all sat .

Variants with Relative Budgets

Most of the proportionality requirements we introduced heavily rely on the budget limit b . This is particularly true for the axioms based on justified representation. [Aziz, Lee and Talmon \(2018\)](#) suggest to work on properties that are independent of the budget limit and only defined in terms of the cost of the budget allocation under consideration.

They revisit their adaptations of PJR for PB by changing the notion of cohesive group, making it dependent on the cost $c(\pi)$ of the budget allocation π under consideration instead of b . All of these new concepts are weaker than the standard ones. They also are all satisfiable, simply by using $\pi = \emptyset$ (note that because of how we organised the elements in our definition for cohesive groups—definitions 3.3.1 and 3.3.9—this does not lead to any division by 0).

Laminar Proportionality

The next property we want to mention is *laminar proportionality*. It is a proportionality requirement that only applies to specific instances, the *laminar* ones. These instances are very well-structured in a way that makes it obvious which outcomes

are proportional. Laminar proportionality requires the outcome to be proportional with respect to this structure.

This property was defined for PB by [Los, Christoff and Grossi \(2022\)](#). They show that rules that satisfy laminar proportionality in the multi-winner setting (namely MES and SEQPHRAG) cease to be so on PB instances.

Proportionality for Solid Coalitions

In Section 3.3.4 we have defined two axioms for proportionality with weak ordinal ballots. Approval ballots can be seen as a special case of weak ordinal ballots where all ballots have at most two indifference classes. Following this observation, [Aziz and Lee \(2021\)](#) provide definitions of IPSC and CPSC for approval ballots. We provide these definitions below. Observe that they are both closely related to PJR.

Definition 3.3.43 (CPSC with Approval Ballots). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile \mathbf{A} of approval ballots, a budget allocation $\pi \in \text{FEAS}(I)$ is said to satisfy CPSC if the following two conditions hold:*

- ▶ π satisfies $\text{PJR}[\text{sat}^{\text{cost}}]$;
- ▶ π is of maximal cost: $\pi \in \arg \max_{\pi' \in \text{FEAS}(I)} (\pi')$.

Definition 3.3.44 (IPSC with Approval Ballots). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile \mathbf{A} of approval ballots, a budget allocation $\pi \in \text{FEAS}(I)$ is said to satisfy IPSC if the following two conditions hold:*

- ▶ for all sets of voters $N \subseteq \mathcal{N}$ such that $c(\bigcup_{i \in N} \{p \in \pi \mid A_i(p) = 1\}) < |N|/n \cdot b$ and for all $p \in \bigcap_{i \in N} \{p \in \mathcal{P} \setminus \pi \mid A_i(p) = 1\}$ we have:

$$c(p) + c\left(\bigcup_{i \in N} \{p \in \pi \mid A_i(p) = 1\}\right) > |N|/n \cdot b,$$

- ▶ π is exhaustive.

The first bullet point of the above definition closely resembles $\text{PJR-X}[\text{sat}^{\text{cost}}]$. One can actually prove that IPSC implies $\text{PJR-X}[\text{sat}^{\text{cost}}]$. Indeed, if a budget allocation π fails $\text{PJR-X}[\text{sat}^{\text{cost}}]$, then the P -cohesive N witnessing this violation would also be a witness of the violation of the first bullet point of the definition of IPSC.

It should be quite clear from the definition that CPSC is still not satisfiable with approval ballots. IPSC is, since it already was with generic weak ordinal ballots.

Proportionality with Cumulative Ballots

Among the different types of cardinal ballots we defined, there is one for which we still have not discussed proportionality requirements: *cumulative ballots*. Now is the time to do so. The only study on cumulative ballots has been conducted by [Skowron, Slinko, Szufa and Talmon \(2020\)](#). Among others, they study proportionality axioms for this setting. We present here what they call *proportional representation*.

Definition 3.3.45 (Proportional Representation with Cumulative Ballots). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile A of cumulative ballots, a budget allocation $\pi \in \text{FEAS}(I)$ is said to satisfy proportional representation if for every $\ell \in \{1, \dots, b\}$, every group of agents $N \subseteq \mathcal{N}$ with $|N|/n \cdot b \geq \ell$ and every subset of projects $P \subseteq \mathcal{P}$ with $c(P) \leq \ell$, it holds that if for all $i \in N$ and $p \in P$, we have $A_i(p) > 0$, and for all $i \in \mathcal{N} \setminus N$ and $p \in \mathcal{P} \setminus P$ we have $A_i(p) = 0$, then we must have $P \subseteq \pi$.*

[Skowron, Slinko, Szufa and Talmon \(2020\)](#) also introduce a weaker and a stronger variant of the above. They prove that all of them are satisfiable.

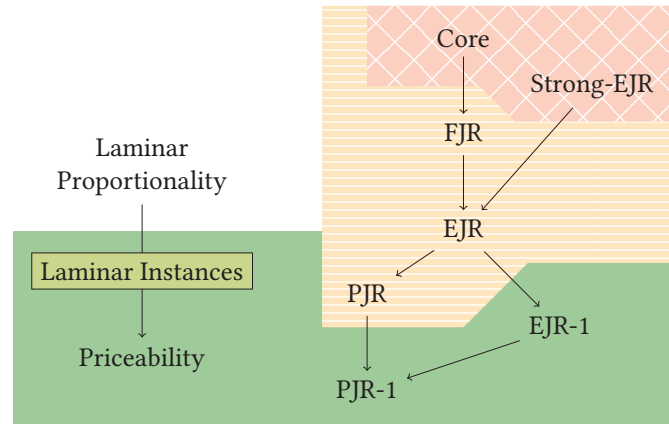
3.3.6 Fairness in Extended Settings

The core of this thesis is concerned with variations in the standard model of PB, we now mention some papers that have studied fairness in PB beyond the standard model. This section overlaps in some way with Section 3.6 though we only focus on fairness requirements here.

- ▶ In their study of PB with multiple resources, [Motamed, Soeteman, Rey and Endriss \(2022\)](#) introduced several proportionality axioms and studied whether they could be satisfied by some load-balancing mechanisms.
- ▶ When studying PB with uncertainty on the cost of the projects, [Baumeister, Boes and Laußmann \(2022\)](#) investigated the link between properties specific to their setting and justified representation axioms such as $\text{PJR}[\text{sat}^{\text{cost}}]$ (a.k.a. BPJR-L) and EJL.
- ▶ In a model in which the budget is endogenous, [Aziz and Ganguly \(2021\)](#) studied versions of the core and of a simple proportionality axiom, investigating which welfare-maximising rule satisfy them.

3.3.7 Taxonomies of Proportionality in PB

Throughout this section we have introduced a significant number of properties related to proportionality in PB. In an attempt to clarify the relationship between these properties, we draw several taxonomies. The taxonomy for cardinal ballots can be found in Figure 3.3.1. Figure 3.3.2 presents the taxonomy for approval ballots. All the details are available in the figures. We also summarise which rules satisfy which axioms, in Table 3.3.1 for cardinal ballots, and in Table 3.3.2 for approval ballots.



- These concepts cannot always be satisfied.
- These concepts can always be satisfied, but finding a suitable budget allocation cannot be done in polynomial time unless $P = NP$.
- These concepts can always be satisfied, and a suitable budget allocation can be found in polynomial time.
- Laminar Proportionality is always satisfiable, and the computational complexity of finding a budget allocation satisfying it is unknown.

Figure 3.3.1: Taxonomy of the proportionality requirements for PB with cardinal ballots. An arrow between two concepts means that any budget allocation satisfying one also satisfies the other. All missing arrows are known to be missing.

Most of this picture is based on [Los, Christoff and Grossi \(2022\)](#) who showed: the absence of arrows between either the core, EJR or PJR and priceability; the link between laminar proportionality and priceability (only for laminar instances); the absence of arrows between laminar proportionality and either PJR, EJR, or the core. The link between FJR and EJR is due to [Peters, Pierczyński and Skowron \(2021\)](#). For the satisfiability of the concepts, see [Table 3.3.1](#).

Cardinal Ballots		
Core	None	Peters, Pierczyński and Skowron (2021)
FJR	Greedy cohesive rule	Peters, Pierczyński and Skowron (2021)
Strong-EJR	None	
EJR	Greedy cohesive rule	Peters, Pierczyński and Skowron (2021)
EJR-1	MES	Peters, Pierczyński and Skowron (2021)
PJR	Greedy cohesive rule	
PJR-1	MES	Los, Christoff and Grossi (2022)
Laminar Proportionality	?	Los, Christoff and Grossi (2022)
Priceability	MES	Peters, Pierczyński and Skowron (2021)

Table 3.3.1: Rules satisfying the fairness criteria for cardinal ballots

Approval Ballots		
Core [sat]	?	
FJR [sat]	▶ for any sat , a greedy cohesive rule	[1]
Strong-EJR [sat]	▶ None	
EJR [sat]	▶ A greedy cohesive rule for any sat ▶ $MES[sat^{card}]$ for $sat = sat^{card}$	[1] [1]
EJR-X [sat]	▶ for any sat , a greedy cohesive rule ▶ for any DNS function sat , $MES[sat]$	[1] [2]
EJR-1 [sat]	▶ for any sat , a greedy cohesive rule ▶ for any additive sat , $MES[sat]$	[1] [1]
PJR [sat]	▶ for any sat , a greedy cohesive rule	[1]
PJR-X [sat]	▶ for any sat , a greedy cohesive rule ▶ for any DNS function sat , $MES[sat]$, $SEQPHRAG$, and $MAXIMINSUPP$	[1] [2]
CPSC	▶ None	[3]
IPSC	▶ The expanding approval rule	[3]
Local-BPJR-L [sat]	▶ $MES[sat]$, $SEQPHRAG$, and $MAXIMINSUPP$	[2, 4]
Strong-BPJR-L	▶ None	[4]
Priceability	▶ $MES[sat]$, $SEQPHRAG$, and $MAXIMINSUPP$	[1, 2]
Priceability with $\alpha > b$	▶ $MES[sat]$, $SEQPHRAG$, and $MAXIMINSUPP$	[1, 2]
Priceability with C6 and $\alpha > b$	▶ $MES[sat^{card}]$, $SEQPHRAG$, and $MAXIMINSUPP$	[2]

[1] [Peters, Pierczyński and Skowron \(2021\)](#)

[2] [Brill, Forster, Lackner, Maly and Peters \(2023\)](#)

[3] [Aziz and Lee \(2021\)](#)

[4] [Aziz, Lee and Talmon \(2018\)](#)

Table 3.3.2: Rules satisfying each of the fairness properties we introduced for approval ballots. In the above, sat is an arbitrary satisfaction function.

	Exhaustiveness	
	General Instances	Unit-Cost Instances
MAXCARD	✓	✓
GREEDCARD	✓	✓
MAXCOST	✓	✓
GREEDCOST	✓	✓
SEQPHRAG	✗	✓
MAXIMINSUPP	✗	✓
MES	✗	✗

Table 3.4.1: Satisfaction of exhaustiveness for different rules.

3.4 Axiomatic Analysis

Fairness requirements are the most studied properties in the literature on PB but are not the only ones. In the following, we review other axioms.

Our analysis will start with a discussion of exhaustiveness (Section 3.4.1) and a presentation of the monotonicity axioms for PB (Section 3.4.2). From there, we will move on to axioms relating to strategic behaviour of the agents (Section 3.4.3). We will conclude this section by our usual discussion of the concepts that exist in the literature but do not fit in earlier sections (Section 3.4.4).

3.4.1 Exhaustiveness

Let us start with *exhaustiveness*. Its definition has already been provided in Section 2.1. It is sometimes considered a standard requirement that should be enforced by default. However, as we will see, it is incompatible with some other axioms, notably priceability. Note that [Talmon and Faliszewski \(2019\)](#) introduced *budget monotonicity*, an axiom that is equivalent to exhaustiveness for resolute rules and very similar to it for irresolute rules; the name exhaustiveness is due to [Aziz, Lee and Talmon \(2018\)](#).

Table 3.4.1 summarises which of the usual rules satisfy exhaustiveness. The results for the welfare-maximising and greedy rules are straightforward. Interestingly, the fact that SEQPHRAG, MAXIMINSUPP and MES fail exhaustiveness is due the fact that they are priceable. Indeed, the two requirements are incompatible.

Proposition 3.4.1 ([Peters, Pierczyński and Skowron 2021](#)). *There exists an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile \mathbf{A} , such that there is no budget allocation $\pi \in \text{FEAS}(I)$ which is both priceable and exhaustive, even though there are feasible budget allocations that are priceable, and others that are exhaustive.*

Since exhaustiveness is sometimes considered to be a must, [Peters, Pierczyński and Skowron \(2021\)](#) proposed several ways to obtain exhaustive outcomes when using non-exhaustive rules.

- ▶ **Completion via Exhaustive Rule:** This technique consists of completing the original outcome of the rule by applying another rule, which is exhaustive, on the reduced instance where the selected projects have been removed and the budget reduced accordingly. Typically, one could use a greedy selection procedure or an exhaustive variant of SEQPHRAG.
- ▶ **Exhaustion by Variation of the Budget Limit:** Using this technique, the rule is run several times for different values of the budget limit until finding an outcome that is feasible and exhaustive for the initial budget. Typically, the budget limit is increased by one unit per voter at each iteration and the final outcome is the first exhaustive one that is found, or the first one for which increasing the budget one more time would lead to an outcome that is not feasible for the original budget limit.

Note that this technique does not guarantee that the outcome will be exhaustive (notably because when used with MES, the outcome would still be priceable). Moreover, this is not necessarily a “completion technique” since many rules are not limit monotonic (see Section 3.4.2), so the final outcome does not need to be a superset of the initial outcome.

- ▶ **Exhaustion by Perturbation of the Ballots:** This final technique modifies the profile slightly so that the outcome is guaranteed to be exhaustive. Which perturbation mechanism should be used depends on the rule. For instance, for MES with cardinal ballots, it is known that if every voter reports a *strictly* positive score for all the projects, then the outcome of MES is exhaustive. Therefore, one could apply MES on the modified profile in which all 0 scores have been replaced by an arbitrary small value.

3.4.2 Monotonicity Requirements

Talmon and Faliszewski (2019) introduced several monotonicity axioms for PB that represent to this date the largest corpus of axioms that has been proposed (if we disregard proportionality requirements). All of these axioms regard the behaviour of PB rules in dynamic environments: when the cost function changes, when the set of projects changes, *etc.* Hence, they can also be interpreted as robustness requirements: they enforce that the outcome does not change much with small variations of the instance. We will define and discuss these axioms in the following.

The first axiom we consider constrains the behaviour of the rule when the cost function changes.

Definition 3.4.2 (Discount Monotonicity). *A PB rule F is said to be discount-monotonic if, for any two PB instances $I = \langle \mathcal{P}, c, b \rangle$ and $I' = \langle \mathcal{P}, c', b \rangle$ such that for some distinguished project $p^* \in \mathcal{P}$, we have $c(p^*) > c'(p^*)$, and $c(p) = c'(p)$ for all $p \in \mathcal{P} \setminus \{p^*\}$, it is the case that $p^* \in F(I, \mathbf{A})$ implies $p^* \in F(I', \mathbf{A})$ for all profiles \mathbf{A} .*

Thus, a rule is discount monotonic if whenever the price of a selected project p decreases, the rule would still select project p .

The second axiom, inspired by *committee monotonicity* in the multi-winner voting literature (Lackner and Skowron, 2023), investigates the behaviour of the rule when the budget limit changes.

Definition 3.4.3 (Limit Monotonicity). *A PB rule F is said to be limit-monotonic if, for any two PB instances $I = \langle \mathcal{P}, c, b \rangle$ and $I' = \langle \mathcal{P}, c, b' \rangle$ with $b < b'$ and $c(p) \leq b$ for all projects $p \in \mathcal{P}$, it is the case that $F(I, \mathbf{A}) \subseteq F(I', \mathbf{A})$ for all profiles \mathbf{A} .*

Thus, a rule is limit monotonic if it selects a superset of the original set of selected projects when the budget limit increases.

The next two axioms concern cases where the set of projects changes, with some projects being either merged or split. Note that these axioms have only been considered for approval ballots. Since generalising them to arbitrary cardinal ballots is not straightforward, we focus on approval profiles here.

Given a PB instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile \mathbf{A} of approval ballots, we say that the instance $I' = \langle \mathcal{P}', c', b \rangle$ and the profile \mathbf{A}' of approval ballots are *the result of splitting project $p^* \in \mathcal{P}$ into $P^* \subseteq \mathcal{P}'$* (with $P^* \cap \mathcal{P} = \emptyset$, i.e., P^* is a set of new projects), if the following conditions are satisfied:

- ▶ The project p^* is replaced by P^* in the set of projects: $\mathcal{P}' = (\mathcal{P} \setminus \{p^*\}) \cup P^*$;
- ▶ The total cost of P^* is that of p^* , i.e., $c'(P^*) = c(p^*)$; and for all $p \in P^*$, it is the case that $c'(p) > 0$;
- ▶ The cost of every other project is as in c : $c'(p) = c(p)$ for all projects $p \in \mathcal{P}' \setminus P^*$;
- ▶ The project p^* is replaced by P^* in the approval ballots containing it: for every $i \in \mathcal{N}$ with $A_i(p^*) = 0$, we have $A'_i = A_i$, and for every $i \in \mathcal{N}$ with $A_i(p^*) = 1$, we have $A'_i(p) = 1$ for all $p \in P^*$, and $A'_i(p) = A_i(p)$ for all $p \in \mathcal{P}' \setminus P^*$.

We also say that I and \mathbf{A} are *the result of merging P^* into p^** given I' and \mathbf{A}' .

Definition 3.4.4 (Splitting Monotonicity). *A PB rule F is said to be splitting-monotonic if, for any two PB instances $I = \langle \mathcal{P}, c, b \rangle$ and $I' = \langle \mathcal{P}', c', b \rangle$ with corresponding profiles of approval ballots \mathbf{A} and \mathbf{A}' and any project $p \in F(I, \mathbf{A})$ such that I' and \mathbf{A}' are the result of splitting project p into a subset of projects P given I and \mathbf{A} , it is the case that $F(I', \mathbf{A}') \cap P \neq \emptyset$.*

Definition 3.4.5 (Merging Monotonicity). *A PB rule F is said to be merging-monotonic if, for any two PB instances $I = \langle \mathcal{P}, c, b \rangle$ and $I' = \langle \mathcal{P}', c', b \rangle$ with corresponding profiles of approval ballots \mathbf{A} and \mathbf{A}' , and any subset of projects $P \subseteq F(I, \mathbf{A})$ such that I' and \mathbf{A}' are the result of merging project set P into project p given I and \mathbf{A} , it is the case that $p \in F(I', \mathbf{A}')$.*

	Monotonicity			
	Limit	Discount	Splitting	Merging
MAXCARD	✗	✓	✓	✗
GREEDCARD	✗	✓	✓	✗
MAXCOST	✗	✗	✓	✓
GREEDCOST	✗	✗	✗	✓
MES	✗			

Table 3.4.2: Summary of the results concerning the monotonicity axioms for rules used with approval ballots.

The results for MAXCARD, GREEDCARD, MAXCOST and GREEDCOST are due to [Talmon and Faliszewski \(2019\)](#). Note that their proofs contained several mistakes, corrected in part by [Baumeister, Boes and Seeger \(2020\)](#). Specifically, the proof that GREEDCARD fails merging monotonicity is wrong, but the results still holds (though it is solely based on the use of tie-breaking rules that apply differently before and after merging projects). MES fails limit monotonicity as it already did on unit-cost instances ([Lackner and Skowron, 2023](#)).

These two axioms thus require the rule to also apply the splitting and merging operations on its outcome. Note that for splitting monotonicity, a stronger version of it would require *all* the smaller projects to be selected (instead of only one).

We present in [Table 3.4.2](#) what is known about the standard PB rules regarding those axioms. The relevant references are provided there. Observe that it is not known which monotonicity axioms are satisfied by SEQPHRAG, MAXIMINSUPP and MES. One exception is that we know that MES cannot satisfy limit monotonicity, as it does not satisfy committee monotonicity, the equivalent of limit monotonicity for unit-cost instances ([Lackner and Skowron, 2023](#)).

The definitions we provided above concern *resolute* PB rules, that is, rules that always output a single budget allocation. They have been extended to *irresolute* rules. [Baumeister, Boes and Seeger \(2020\)](#) (and subsequently [Sreedurga, Bhardwaj and Narahari, 2022](#)) extend the monotonicity axioms in an *existential* fashion: for a given instance I and profile \mathbf{A} , and for every budget allocation $\pi \in F(I, \mathbf{A})$ that satisfy a specific pre-condition, it must be the case that for every suitable I' and \mathbf{A}' , there exist a budget allocation $\pi' \in F(I', \mathbf{A}')$ satisfying the specific post-condition. We will see alternative definitions for irresolute rules in [Chapter 6](#).

3.4.3 Strategy-Proofness

The next class of requirements we consider is that of incentive compatibility axioms. These axioms are concerned with preventing agents from engaging in strategic behaviour.

Let us first discuss the concept of *strategy-proofness*. Intuitively speaking, it states

that no agent should be able to obtain a better outcome by reporting a ballot that is different from their true preferences. To define it, we thus need a way of comparing outcomes from the point of view of the agents. When using cardinal ballots, we will assume that the ballot represents the utility of the agents for the projects. For approval ballots, we will use the notion of satisfaction function as the measure of utility.²⁸

Definition 3.4.6 (Strategy-Proofness for Cardinal Ballots). *A PB rule F is said to be strategy-proof if for every instance I and profile \mathbf{A} of cardinal ballots, and for every agent $i \in \mathcal{N}$, there exists no cardinal ballot A'_i such that for the profile $\mathbf{A}' = (A_1, \dots, A_{i-1}, A'_i, A_{i+1}, \dots, A_n)$ we have:*

$$\sum_{p \in F(I, \mathbf{A}')} A_i(p) > \sum_{p \in F(I, \mathbf{A})} A_i(p).$$

Observe that the satisfaction of the manipulating agent i with the output under the new profile \mathbf{A}' is computed with regards to the initial ballot A_i .

Definition 3.4.7 (Strategy-Proofness for Approval Ballots). *Given a satisfaction function sat , a PB rule F is said to be strategy-proof for sat if for every instance I and profile \mathbf{A} of approval ballots, for every agent $i \in \mathcal{N}$, there exists no approval ballot A'_i such that for the profile $\mathbf{A}' = (A_1, \dots, A_{i-1}, A'_i, A_{i+1}, \dots, A_n)$ we have:*

$$sat(F(I, \mathbf{A}') \cap A_i) > sat(F(I, \mathbf{A}) \cap A_i).$$

It is already known from multi-winner voting, i.e., when instances have unit costs, that strategy-proofness is incompatible with very weak notions of proportionality (Peters, 2018, 2019). This result obviously also applies to general PB instances.

Theorem 3.4.8 (Peters 2018). *A PB rule F is said to be weakly proportional on unit-cost instances if for every unit-cost instance I and profile \mathbf{A} of cardinal ballots such that for all voters $i, i' \in \mathcal{N}$ either $\{p \in \mathcal{P} \mid A_i(p) > 0\} = \{p \in \mathcal{P} \mid A_{i'}(p) > 0\}$, or these two sets do not intersect (meaning that \mathbf{A} is a party-list profile), then for any project $p \in \mathcal{P}$ such that $|\{i \in \mathcal{N} \mid A_i(p) > 0\}| \geq n/b$ we have $p \in F(I, \mathbf{A})$.*

There is no rule that satisfies simultaneously weak proportionality on unit-cost instances and strategy-proofness.

Note that in the actual statement of Peters (2018, 2019), an additional efficiency requirement is needed. This is because in the multi-winner voting setting, one *has* to ensure that a rule selects the required number of candidates (i.e., the rule has to be exhaustive). Since this constraint is lifted in the PB setting, there is no need for such an additional axiom.

²⁸Note here that we are indeed discussing utilities and not satisfaction levels since we are considering behaviours that the agents engage into themselves, according to their private information.

It should also be noted that the proportionality axiom defined in the above statement is particularly weak and is known to be implied by all kinds of other requirements (Peters, 2018), including all the ones introduced in Section 3.3.1. In particular, this implies that rules such as SEQPHRAG, MAXIMINSUPP or MES are not strategy-proof.

This result has been replicated in the multi-resource PB case, for suitable adaptations of the axioms (Motamed, Soeteman, Rey and Endriss, 2022). Moreover, it also applies to irresolute rules (Kluiving, de Vries, Vrijbergen, Boixel and Endriss, 2020).

It is known that with unit-cost instances, welfare-maximising rules such as GREEDCOST (which is equivalent to GREEDCARD, MAXCARD and MAXCOST on unit-cost instances) are strategy-proof. When moving to general PB instances, we can show that GREEDCOST is only *approximately* strategy-proof. We provide below the definition of Goel, Krishnaswamy, Sakshuwong and Aitamurto (2019) that weakens strategy-proofness in a “up-to-one” fashion.

Definition 3.4.9 (Approximate Strategy-Proofness for Approval Ballots). *Given a satisfaction function sat , a PB rule F is said to be approximately strategy-proof for sat if for every instance I and profile \mathbf{A} of approval ballots, for every agent $i \in \mathcal{N}$, there exists no approval ballot A'_i such that for the profile $\mathbf{A}' = (A_1, \dots, A_{i-1}, A'_i, A_{i+1}, \dots, A_n)$, for all $p \in \mathcal{P}$ we have:*

$$sat(F(I, \mathbf{A}') \cap A_i) > sat((F(I, \mathbf{A}) \cap A_i) \cup \{p\}).$$

Proposition 3.4.10 (Goel, Krishnaswamy, Sakshuwong and Aitamurto 2019). *The GREEDCOST is approximately strategy-proof for sat^{cost} .*

Note that the result by Goel, Krishnaswamy, Sakshuwong and Aitamurto (2019) uses knapsack ballots. This is not required when projects are indivisible.²⁹ It is also worth noting that this result does not hold for sat^{card} .

Interestingly, exact welfare-maximising rules such as MAXCARD or MAXCOST fail even approximate strategy-proofness on PB instances, for large sets of satisfaction functions. This can come as a surprise since they are strategy-proof on unit-cost instances. Note that this also holds if ballots are knapsack ballots.

²⁹Let us sketch the proof, originally devised by Ulle Endriss. For any given $I = \langle \mathcal{P}, c, b \rangle$, consider $I' = \langle \mathcal{P}', c', b \rangle$, where projects in \mathcal{P} have been split into sets of subprojects, each of cost 1. I' is thus a unit-cost instance. We can transform any given profile \mathbf{A} of approval ballots in the same manner to obtain a profile \mathbf{A}' of approval ballots. Now, it is clear that the approval scores of the projects in \mathbf{A}' are equal to those of the projects in \mathcal{P} they come from in \mathbf{A} . Assume that the tie-breaking rule is extended in a consistent way from projects in \mathcal{P} to projects in \mathcal{P}' . Then we know that there exists at most one project $p \in \mathcal{P}$ such that $\text{GREEDCOST}(I', \mathbf{A}')$ contains a proper subset of its corresponding subprojects. Let $\pi' \subseteq \mathcal{P}$ be the budget allocation that includes any project in \mathcal{P} for which at least one corresponding subproject is in $\text{GREEDCOST}(I', \mathbf{A}')$. We thus have $\text{GREEDCOST}(I, \mathbf{A}) \cup \{p\} = \pi'$. Since GREEDCOST is strategy-proof over unit-cost instances (Peters, 2018), no agent can reach a better budget allocation than π' by strategising, when considering the satisfaction function sat^{cost} .

Example 3.4.11. Consider the instance $I = \langle \mathcal{P}, c, b \rangle$ with $\mathcal{P} = \{p_1, \dots, p_5\}$, the cost are such that $c(p_1) = 6$, $c(p_2) = 3$ and $c(p_3) = c(p_4) = c(p_5) = 3$, and the budget limit is $b = 6$.

Assume that three agents are involved in the process for whom the truthful ballots are to approve of p_1 for agent 1; p_2 for agent 2; and p_3, p_4 , and p_5 for agent 3. If ties are broken lexicographically, the outcome of both MAXCARD and MAXCOST would then be $\pi = \{p_1\}$. Note that agent 3 has satisfaction 0 for π . Now, if agent 3 were to approve of p_2, p_3, p_4 and p_5 instead, the outcome would be $\pi' = \{p_2, p_3, p_4, p_5\}$. Is it clear that for any satisfaction function that is strictly monotonic³⁰ and for every project $p \in \mathcal{P}$, agent 3 prefers π' over $\pi \cup \{p\}$. \triangle

3.4.4 Other Axioms

Let us conclude by mentioning some other axioms and axiomatic directions that have been followed in the context of PB.

In their study on maximin PB with approval ballots, [Sreedurga, Bhardwaj and Narahari \(2022\)](#) adapt several axioms from the multi-winner literature to the context of PB with irresolute rules. These axioms are the *narrow-top criterion* (an adaptation of unanimity) and *clone-proofness* (the outcome of a rule remains the same if projects are cloned). They also introduce a new axiom called *maximal coverage* stating that no redundant project should ever be selected unless it is not possible to cover more voters, where a voter is *covered* if at least one of their approved projects have been selected, and a project is *redundant* if removing it does not change the set of covered voters. Note that this axiom can be seen as a fairness requirement.

Following a more typical social choice route, [Ceron, Gonzalez and Navarro-Ramos \(2022\)](#) initiated the axiomatic characterisation of PB rules, focusing on GREEDCOST.

Finally, it is also worth mentioning that [Goel, Krishnaswamy, Sakshuwong and Aitamurto \(2019\)](#) provided the first analysis of PB rules in terms of epistemic criteria (being a maximum likelihood estimator) to date, another branch of the axiomatic approach ([Elkind and Slinko, 2016](#); [Pivato, 2019](#)), but for the divisible setting.

3.5 Algorithmic Considerations

Another large part of the literature focuses on the algorithmic aspects of PB. This usually concerns computing outcomes of PB rules and the exact complexity of welfare maximisation under different models.

We will discuss these different aspects, focusing first on outcome determination (Section 3.5.1), then on the complexity of welfare maximisation (Section 3.5.2), and finally on the other algorithmic problems that have been studied (Section 3.5.3).

³⁰A satisfaction function *sat* is strictly monotonic if for all $P \subseteq \mathcal{P}$ and $P' \subsetneq P$, we have $sat(P') < sat(P)$.

3.5.1 Outcome Determination of Standard PB Rules

The main focus of the algorithmic perspective on social choice is to assess the computational complexity of computing “good” outcomes. With all that has been presented so far, we already know a lot about the quality of the outcome of the standard PB rules. The last step is thus to assess how hard it is to compute said outcomes.

Formally speaking, this is the problem of computing the outcome of a given rule, the so-called *outcome determination problem*. We present below one version of this problem, for a given resolute PB rule F .

OUTCOMEDETERMINATION(F)	
Input:	An instance $I = \langle \mathcal{P}, c, b \rangle$, a profile \mathbf{A} , and a project $p \in \mathcal{P}$.
Question:	Is $p \in F(I, \mathbf{A})$?

Note that this definition only makes sense for *resolute* PB rules. Other formulations are also possible, for example as a function problem.

The complexity of the winner determination problem for irresolute PB rules has not been considered in the literature yet and it is not immediately clear how the outcome determination problem should be formulated. One natural idea would be to define the problem as checking whether a project is *always* selected, or whether it is *sometimes* selected.

It should be more or less clear that the outcome determination problem can be efficiently solved for most of the rules that we have focused on, at least in the resolute case. The definitions of GREEDCARD, GREEDCOST and SEQPHRAG should make it somewhat obvious that computing their outcomes can be done efficiently. For MAXIMINSUPP, Aziz, Lee and Talmon (2018) presents a linear program allowing to compute efficiently an optimum load distribution at each round. Finally, Peters, Pierczyński and Skowron (2021) discuss how to efficiently compute outcomes of MES.

The only rules whose outcomes cannot be computed efficiently are the ones that relate to exact welfare maximisation. Indeed, maximising the social welfare is usually hard, as we shall see next.

3.5.2 Maximising Social Welfare

Let us now turn to the computational problem of maximising measures of social welfare.

First, we introduce the different notions of social welfare that have been studied in the literature. Note that throughout this section, we will work with cardinal ballots. We also repeat the definition of UTIL-SW so that the reader does not need to go back to Section 3.2.1.

- **Utilitarian Social Welfare:** Given an instance I and a profile \mathbf{A} of cardinal ballots, the *utilitarian social welfare* achieved by a budget allocation π is defined

as:

$$\text{UTIL-SW}(I, \mathbf{A}, \pi) = \sum_{i \in \mathcal{N}} \sum_{p \in \pi} A_i(p).$$

This is the most standard definition of social welfare simply considering the sum of the satisfactions of the individuals. A budget allocation maximising UTIL-SW selects the items that are individually best, *i.e.*, it ignores any interactions between the projects.

- **Chamberlin-Courant Social Welfare:** Given an instance I and a profile \mathbf{A} of cardinal ballots, the *Chamberlin-Courant social welfare* achieved by a budget allocation π is defined as:

$$\text{CC-SW}(I, \mathbf{A}, \pi) = \sum_{i \in \mathcal{N}} \max_{p \in \pi} A_i(p).$$

The Chamberlin-Courant social welfare assumes that agents only consider one project from each budget allocation, namely the one that leads to the highest satisfaction. Maximising CC-SW corresponds thus to aiming for a budget allocation that represents as many voters as possible.

Note that CC-SW has been studied by [Laruelle \(2021\)](#) under the name *Rawlsian social welfare*.

- **Egalitarian Social Welfare:** Given an instance I and a profile \mathbf{A} of cardinal ballots, the *egalitarian social welfare* achieved by a budget allocation π is defined as:

$$\text{EGAL-SW}(I, \mathbf{A}, \pi) = \min_{i \in \mathcal{N}} \sum_{p \in \pi} A_i(p).$$

The egalitarian social welfare assumes that the welfare of a society is the satisfaction of its most dissatisfied member. Maximising EGAL-SW hence means maximising the satisfaction of the worst-off voter.

EGAL-SW is studied by [Sreedurga, Bhardwaj and Narahari \(2022\)](#) under the name *maximin PB*.

- **Nash Social Welfare:** Given an instance I and a profile \mathbf{A} of cardinal ballots, the *Nash social welfare* achieved by a budget allocation π is defined as:

$$\text{NASH-SW}(I, \mathbf{A}, \pi) = \prod_{i \in \mathcal{N}} \sum_{p \in \pi} A_i(p).$$

The Nash social welfare measure can be seen as a compromise between utilitarian and egalitarian social welfare. By maximising NASH-SW, one aims to find a fair budget allocation ([Fluschnik, Skowron, Triphaus and Wilker, 2019](#)).

Note that maximising NASH-SW is equivalent to maximising the sum of the logarithms of the satisfactions of the agents.

The typical computational problem is then to determine whether there is a budget allocation that provides at least a certain amount of satisfaction according to a specific measures of welfare. [Fluschnik, Skowron, Triphaus and Wilker \(2019\)](#) studied this problem for UTIL-SW, NASH-SW and CC-SW. [Sreedurga, Bhardwaj and Narahari \(2022\)](#) considered the case of EGAL-SW, in the context of approval ballots with sat^{cost} . [Talmon and Faliszewski \(2019\)](#) focused on UTIL-SW with approval ballots and several satisfaction functions. We summarise the main findings in [Table 3.5.1](#).

Welfare maximisation problems have also been studied for many of the variations of the standard model that have been introduced. We just mention them here and refer the reader to [Section 3.6](#) for more details. [Hershkowitz, Kahng, Peters and Procaccia \(2021\)](#) studied welfare maximisation in a model in which projects are grouped into districts. Similarly, [Jain, Sornat, Talmon and Zehavi \(2021\)](#) and [Patel, Khan and Louis \(2021\)](#) investigated different social welfare maximisation when projects are grouped in categories. [Jain, Sornat and Talmon \(2020\)](#) looked into social welfare for non-additive satisfaction functions. Social welfare has been studied in multi-resource PB ([Motamed, Soeteman, Rey and Endriss, 2022](#)), when the cost is dependent on the number of users of the projects ([Lu and Boutilier, 2011](#)), and when the budget is endogenous ([Aziz and Ganguly, 2021](#); [Aziz, Gujar, Padala, Suzuki and Vollen, 2022](#); [Chen, Lackner and Maly, 2022](#)).

3.5.3 Other Algorithmic Problems

Participatory budgeting offers other avenues for studies focusing on the computational complexity of related problems.

For instance, [Baumeister, Boes and Hillebrand \(2021\)](#) study the computational complexity of control in PB instances with approval ballots. Control problems are problems of the form “Can the decision maker achieve certain objectives by changing certain parameters of the instance?”. More specifically, [Baumeister, Boes and Hillebrand \(2021\)](#) studied two types of control for GREEDCARD, GREEDCOST, MAXCARD, and MAXCOST when the decision maker can decide on the price of a project, or on the budget limit. Under *constructive control*, the decision maker aims at forcing the selection of a given project, while under *destructive control*, they aim at preventing a given project from being selected.

3.6 Variations and Extensions of the Standard Model

The literature we reviewed so far studied the standard model of PB. Beyond that, a myriad of variations of the model have been introduced. Part [Three](#) of the thesis will present three such extensions. In the following we delve into the other variations that have been proposed.

	Weakly NP-complete	► Even with one voter
UTIL-SW	Pseudo-poly. solvable	
	Poly. solvable	► With approval ballots and sat^{card}
	Strongly NP-complete	<ul style="list-style-type: none"> ► Even with one voter ► Even with two voters and unit-cost instances ► Even with unit-cost instances and $A_i(p) \in \{0, 1\}$ for all $i \in \mathcal{N}$ and $p \in \mathcal{P}$
NASH-SW	W[1]-hard	<ul style="list-style-type: none"> ► Parameterised by the budget limit b, even with unit-cost instances and $A_i(p) \in \{0, 1\}$ for all $i \in \mathcal{N}$ and $p \in \mathcal{P}$ ► Parameterised by the budget limit b and the number of voters n, even with unit-cost instances and unary encoding ► Even with single-peaked or single crossing profiles
	XP	► Parameterised by the number of voters n
	FPT	► Parameterised by the number of voters n and $\max_{i \in \mathcal{N}} \{\sum_{p \in \pi} A_i(p) \mid \pi \in \text{FEAS}(I)\} $
	Pseudo-poly. solvable	► For single-peaked and single-crossing profiles
CC-SW	Strongly NP-complete	► Even for binary valuations, <i>i.e.</i> , ballots with only two different values, and unit-cost instances
	FPT	► Parameterised by the number of voters and $\sum_{i \in \mathcal{N}} \sum_{p \in \mathcal{P}} A_i(p)$
	W[2]-hard	► Parameterised by the budget limit b
	Strongly NP-complete	► Even if $A_i(p) \in \{0, c(p)\}$ for all $i \in \mathcal{N}$ and $p \in \mathcal{P}$
EGAL-SW	Pseudo-poly. solvable	► When $A_i(p) \in \{0, c(p)\}$ for all $i \in \mathcal{N}$ and $p \in \mathcal{P}$ and the number of distinct ballots is constant
	Poly. solvable	► When $A_i(p) \in \{0, c(p)\}$ for all $i \in \mathcal{N}$ and $p \in \mathcal{P}$, the number of distinct ballots is constant, and $\frac{\max_{p \in \mathcal{P}} c(p)}{\text{GCD}\{c(p) \mid p \in \mathcal{P}\}}$ is constant.

Table 3.5.1: Computational complexity of social welfare maximisation problems. For a given measure of welfare SW, the exact decision problem that is considered is the following: given an instance $I = \langle \mathcal{P}, c, b \rangle$, a profile \mathbf{A} of cardinal ballots, and a target $x \in \mathbb{Q}_{\geq 0}$, is there a budget allocation $\pi \in \text{FEAS}(I)$ such that $\text{SW}(I, \mathbf{A}, \pi) \geq x$?

Statements for UTIL-SW follow immediately from the literature on knapsack problems (Kellerer, Pferschy and Pisinger, 2004) as explained by Talmon and Faliszewski (2019). The results for NASH-SW and CC-SW are due to Fluschnik, Skowron, Triphaus and Wilker (2019). CC-SW with approval ballots was studied by Talmon and Faliszewski (2019). Sreedurga, Bhardwaj and Narahari (2022) studied EGAL-SW.

3.6.1 Local versus Global Processes

Real-life PB processes tend to be implemented at the scale of a municipality. It is very common for the municipality to actually implement several local PB processes, one for each district for instance, instead of one general process. This is the case in Amsterdam (City of Amsterdam, 2022) for instance. Motivated by this observation Hershkowitz, Kahng, Peters and Procaccia (2021) investigate the effect of the local versus global implementation of PB processes.

In their study, Hershkowitz, Kahng, Peters and Procaccia (2021) introduce a model of *district-based PB*. Each project belongs to a specific district and contributes a fixed additive amount to the welfare of its district. In addition, there is a budget limit for each district. A budget allocation is called *district-fair* if it provides each district at least as much social welfare as they could achieve with their share of the budget limit. The authors then consider the problem of selecting a global budget allocation that is district-fair.

Hershkowitz, Kahng, Peters and Procaccia (2021) show that it is computationally hard to maximise social welfare under district-fairness constraints. In addition, they show that one can, in polynomial time, find probabilistic outcomes that maximise the global social welfare while being almost district-fair in expectation. Finally, they show that by slightly overspending (by a factor $1.647 + \epsilon$, with $\epsilon > 0$), one can find in polynomial time budget allocations that maximise the global social welfare while providing “district-fairness up to one project” to each district.

3.6.2 Additional Distributional Constraints

An important part of the literature on extensions of the standard model focuses on incorporating additional constraints into the standard model. These constraints are usually distributional ones that affect which projects can be selected. They can model the fact that some projects are *incompatible*, or that some projects have *positive interactions* for instance. Let us present some examples.

The main type of additional constraints that have been studied are *categorical*. These constraints model the idea that projects are grouped into categories and that additional constraints apply regarding which of the projects can be selected within each category. More specifically:

- ▶ Jain, Sornat, Talmon and Zehavi (2021) study what they refer to as *PB with project groups* (which we call cost quota constraints in Chapter 6). In their setting, projects are grouped into categories and there are constraints on the total cost of the selected projects from each category. They focus on the computational aspects of finding a feasible budget allocation maximising the social welfare, and they provide an in-depth analysis of this extended PB setting: Parameterized complexity analysis, and approximability and inapproximability results. In particular, they provide efficient algorithms to maximise or to ap-

proximate the social welfare when the number of categories is small; while proving hardness for arbitrary number of categories.

- ▶ [Patel, Khan and Louis \(2021\)](#) investigate the computational complexity of selecting group-fair knapsacks. This problem is equivalent to selecting budget allocations maximising the utilitarian social welfare in PB instances with categories over the projects, and upper and lower quotas on the categories. The quotas are expressed either in terms of number of selected projects per category, or contribution to the social welfare per category. They prove hardness results, and provide intricate dynamic programming algorithms that compute approximate solutions.
- ▶ Quotas on the number of project selected per category have also been considered by [Chen, Lackner and Maly \(2022\)](#) in a model with endogenous funding.

Note that [Chen, Lackner and Maly \(2022\)](#) do not assume categories to be disjoint while [Jain, Sornat, Talmon and Zehavi \(2021\)](#) and [Patel, Khan and Louis \(2021\)](#) do.

Let us also mention that when studying PB with multidimensional costs, [Motamed, Soeteman, Rey and Endriss \(2022\)](#) show how to encode distributional constraints simply by using extra resources. They discuss dependency constraints, categorical constraints (upper quota on the cost of a category), and incompatibility constraints (categorical constraints with quotas on the upper number of projects selected in a category).

Further additional constraints will be detailed in Chapter 6.

3.6.3 Interaction Between Projects

One assumption that is almost always made is that projects are independent. We have seen above how to incorporate distributional constraint challenging that assumption at the level of which budget allocations are admissible or not. In a similar spirit, [Jain, Sornat and Talmon \(2020\)](#) challenge the independence assumption from the perspective of the voters, assuming that the satisfaction of the voters is not *additive*, *i.e.*, can be more, or less, than the sum of its parts.

Specifically, [Jain, Sornat and Talmon \(2020\)](#) assume that there is an *interaction structure* partitioning the projects into categories. The utility of the voters is defined as the sum of their satisfaction for each category, the latter being an increasing, but potentially non-linear, function of the number of approved and selected projects from within the category. This model enables the study of substitution or complementarity effects between the projects from the perspective of the voters.

On top of their conceptual contribution, [Jain, Sornat and Talmon \(2020\)](#) present a computational analysis of welfare-maximising problems in this setting. They provide a mixture of tractability and intractability results. They also identify restrictions of the ballots submitted by the voters—defined with respect to a specific interaction structure—for which the computational problems become tractable.

Note that in the work of [Jain, Sornat and Talmon \(2020\)](#), the interaction structure is given and fixed for all voters. In subsequent work, [Jain, Talmon and Bulteau \(2021\)](#) analysed how to obtain such an interaction structure based on several partitions of the projects submitted by the agents. The focus is computational there as well.

3.6.4 Enriched Cost Functions

Another typical assumption that is made is to assume that the cost of the projects is fixed and expressed in only one dimension. Both of these aspects of the cost function have been challenged by different authors.

In one of the first papers on a model at the time not yet called participatory budgeting, [Lu and Boutilier \(2011\)](#) consider the problem of selecting multiple costly alternatives under a given budget constraint. Their model is slightly different from the standard one for PB as they aim at modelling recommendation systems. In particular, selected alternatives are assigned to some agents. What is more interesting for us here is that they assume that the cost of a project is composed of a fixed part and of a variable part. Specifically, the cost of a project is an affine function of the number of agents assigned to that project.

The assumption that costs are unidimensional has also been lifted. [Motamed, Soeteman, Rey and Endriss \(2022\)](#) focus on analysing the effect of multidimensional costs. They extend the standard model for PB, assuming that the costs of the projects are expressed in terms of several resources. In this setting, they define and study proportionality requirements, incentive compatibility axioms, and their interactions. They also touch on the computational aspect of maximising social welfare.

3.6.5 Uncertainty in PB

In practice there is a lot of uncertainty around the actual implementation of the projects. It is for instance rarely possible to assess the cost of the projects exactly, let alone their completion time. [Baumeister, Boes and Laußmann \(2022\)](#) initiated the study of PB under uncertainty about the projects.

In their model, [Baumeister, Boes and Laußmann \(2022\)](#) assume that the costs of the projects are uncertain. For each project, its cost is described as a probability distribution over a specific interval. Projects are associated with a completion time and the actual cost of a project is revealed only once the project has been completed. They consider online mechanisms that select the projects to be funded in a dynamic fashion. Within this framework, they provide a series of impossibility results showing that no online mechanism can be at the same time punctual (finishes within the given time bound), not too risky (the probability of exceeding the budget is never too high, or the excess is never too high), and exhaustive (the budget is not underused). They also adapt the justified representation axioms to this setting, showing that an adaptation of MES provides interesting fairness guarantees.

3.6.6 PB with Endogenous Funding

The standard PB model assumes that the budget is provided by the organising entity. Several authors have proposed alternative models in which the voters can actually contribute their own funds to help implement some projects.

In a model in which voters submit cardinal ballots over the projects, [Chen, Lackner and Maly \(2022\)](#) introduce the idea that voters can also submit monetary contributions to specific projects. They investigate suitable aggregation methods for this framework. The risk with donation is that some voters could have too much influence on the final outcome. Therefore, they focus on devising rules for which the satisfaction of no voter decreases when taking into account donations, compared to the case where the donations are ignored. They provide several such rules, and study their merits regarding some donation-specific monotonicity requirements. They conclude their analysis by studying the computational complexity of outcome determination problems, and the problem of finding optimal donation policy for the voters.

Moving further away from PB, [Aziz and Ganguly \(2021\)](#) propose a setting in which there is no exogenous fund, instead, each agent joins the process with a given personal budget that will be used to fund the projects. Agents submit approval ballots and a rule in this setting determines, given an approval profile and the personal budget of the agents, a set of projects to be funded and the monetary contribution of each individual to the selected projects. This model is slightly different from PB in the sense that it is not about the allocation of public funds. It is nevertheless a framework studying aggregation problems when selecting costly alternatives. They introduce and study several axioms dealing with efficiency (Pareto-optimality), and fairness (core and proportionality). Finally, they investigate several welfare-maximisation rules—based on utilitarian, egalitarian, or Nash social welfare—in terms of these axioms.

[Aziz, Gujar, Padala, Suzuki and Vollen \(2022\)](#) study a similar model except that agents submit cardinal ballots instead of approval ones, and that they have quasi-linear utilities (that depend on the money they spend). They focus on the computational aspects of maximising the utilitarian social welfare subject to some participation requirements (that guarantees the agents not to contribute more than they receive), showing both computational hardness and inapproximability results.

3.6.7 Weighted PB

In their study about PB with ordinal ballots, [Aziz and Lee \(2021\)](#) make the assumption that the voters have different weights. Their analysis does not really focus on this assumption however, and little is known about what its impact is in general.

3.7 Related Frameworks

In this section, we present several frameworks that relate to PB. We do not provide much detail about them but merely give pointers for the interested reader.

Multi-winner voting. The most obviously related framework, as we have mentioned several times already, is multi-winner voting. It is a special case of PB—where instances have unit costs and budget allocations are required to be exhaustive—and has been extensively studied for many years, way before PB became a topic of interest. A recent book by [Lackner and Skowron \(2023\)](#) presents a large part of that literature for approval ballots and provides many relevant references. A good starting point for multi-winner voting beyond approval ballots is the chapter by [Faliszewski, Skowron, Slinko and Talmon \(2017\)](#). Other relevant pointers have already been included in the different sections above.

Collective optimisation problems. As we have seen already PB can be seen as a collective variant of the knapsack problem (see, e.g., [Fluschnik, Skowron, Triphaus and Wilker, 2019](#)). The idea of looking at collective variants of optimisation problems is a growing field into which PB fits nicely ([Boes, Colley, Grandi, Lang and Novaro, 2021](#)). Other optimisation problems for which their collective variants have been studied include finding spanning trees or scheduling jobs on machines ([Darmann, Klamler and Pferschy, 2009, 2011](#); [Brandt, Conitzer, Endriss, Lang and Procaccia, 2016b](#); [Pascual, Rzdca and Skowron, 2018](#)).

Divisible participatory budgeting. Throughout this chapter (and this thesis), we only focused on the case of indivisible PB where the projects are either fully funded or not at all. Relaxing this assumption by allowing projects to receive any amount of funding leads to the world of divisible PB. This framework has sometimes been called *portioning* where a given public resource has to be shared among different divisible projects. Its study dates back to [Bogomolnaia, Moulin and Stong \(2005\)](#) and has since then received substantial attention. Perspectives that have been considered include welfare maximisation ([Goel, Krishnaswamy, Sakshuwong and Aitamurto, 2019](#); [Michorzewski, Peters and Skowron, 2020](#)), fairness guarantees ([Fain, Goel and Munagala, 2016](#); [Caragiannis, Christodoulou and Protopapas, 2022](#); [Airiau, Aziz, Caragiannis, Kruger, Lang and Peters, 2023](#)), strategic behaviour ([Aziz, Bogomolnaia and Moulin, 2019](#); [Freeman, Pennock, Peters and Vaughan, 2021](#); [Brandl, Brandt, Peters and Stricker, 2021](#)). This setting is also closely related to *probabilistic social choice* ([Brandt, 2018](#)).

Fair allocation. PB also relates to the literature on fair allocation ([Rothe, 2015](#); [Brandt, Conitzer, Endriss, Lang and Procaccia, 2016a](#)) and more specifically on the *fair allocation of public goods* ([Conitzer, Freeman and Shah, 2017](#)) where the allocated items can impact several agents (they are not privately owned as is assumed in the typical fair division literature). This framework can be seen as an unconstrained version of PB as there need not be a budget constraint. Note that [Fain, Munagala and Shah \(2018\)](#) consider the same model but with constraints on the outcome, though not necessarily budget constraints.

Part Two

Variations on the Method

Chapter 4

Defining Fairness via Equality of Resources

The fact that budget allocations are composed of several elements—more than one project is typically selected—makes fairness a crucial topic in the investigation of PB. Indeed, when considering aggregation scenarios where the outcome consists of a single alternative, such as presidential elections, there is much less room for fairness as selecting only one alternative will necessarily make some voters unhappy (of course we could then turn to temporal fairness, as we will do in Chapter 7). Intuitively, by selecting more alternatives, we can ensure that more voters are, at least partly, satisfied with the outcome. Fairness deserves thus a special focus in the formal study of PB, and, as we have seen in Chapter 3, a large part of the literature indeed focuses on fairness. This chapter presents a different take on how to define fairness in PB.

The first fundamental question we have to answer concerns the actual definition of fairness for PB, and more precisely, what type of fairness we are interested in. PB is concerned with the allocation of common resources such as public money to projects benefiting members of the society. Fairness concepts related to such framework have usually been studied under the term *distributive justice*. Let us quote the Stanford encyclopedia of philosophy to understand what the latter encompasses (Lamont and Favor, 2017):

« The economic, political, and social frameworks that each society has—its laws, institutions, policies, etc.—result in different distributions of benefits and burdens across members of the society. [...] The structure of these frameworks is important because the distributions of benefits and burdens resulting from them fundamentally affect people's lives. Arguments about which frameworks and/or resulting distributions are morally preferable constitute the topic of distributive justice. »

From this quote, it should be clear that the normative approach dictating how shared resources should be allocated through a PB process relates to distributive justice. Overall, justice—or fairness—in PB has to be defined in terms of how the benefits of the selected budget allocation are distributed across the citizens.

We now know the general framework in which fairness in PB fits. Looking at it in more depth, we run into our second fundamental question: how can we define, and measure, the benefits of the PB process experienced by the citizen? Answering this question is of a primal importance as it will guide us in defining fairness criteria.

Looking into the literature on fairness in PB (see Section 3.3) it immediately appears that it is entirely focusing on *equality of welfare* (Dworkin, 1981a). Indeed, all fairness requirements that have been introduced discuss how to reach a fair distribution of *satisfaction* amongst the voters. This provides a first, partial, answer to our question: the benefits of the PB process experienced by the voters are measured in terms of their satisfaction. This answer is however only partial. Indeed, even though we have settled that satisfaction is the measure of interest, we have still not provided clear ways to measure it.

This brings us to, yet another, fundamental question: how can we measure the satisfaction of the voters in a PB process? This is highly linked to the ballot format that is used. Indeed, the only information about the voter that is available to the decision maker is the ballot they submitted. Since we are focusing on approval ballots here, the amount of information provided by the voter is actually rather limited, making it difficult to discuss their satisfaction. In Chapter 2, we introduced the concept of *satisfaction function* (Definition 2.2.1), which allowed us to still do so, despite the use of approval ballots. However, I do believe that this approach, though mathematically appealing, is not suitable for more applied purposes. This is mainly due to the lack of consensus regarding what should be *the* satisfaction function to consider. Thus, any result applying only to a specific satisfaction function would have somewhat limited applicability. Note that I also do not believe that a consensus on which satisfaction function to use can ever be found, as satisfaction functions are such simplistic representations of what could be the internal reasoning of a voter.³¹

In a nutshell, it seems—at least to me—that there is no satisfactory way to discuss satisfaction with approval ballots. Of course, one could then argue for the use of other ballot formats. However, by considering more sophisticated ballots, one runs into different issues regarding the actual usage of the ballots. Focusing on one question at a time, the conclusion of this discussion is that the approach based on equality of welfare is not fully suited to PB with approval ballots.

The fact that we cannot devise fairness criteria based on notions of satisfaction or welfare that are clearly motivated does not imply that it is not possible to discuss about fairness at all. First, one could still pursue the satisfaction-based approach,

³¹As already discussed in Chapter 3, an approach to mitigate this issue is to obtain results for classes of satisfaction functions (along the line of Brill, Forster, Lackner, Maly and Peters, 2023). Such results are however difficult to obtain, and usually only apply to weaker notions of fairness.

which after all is interesting at a theoretical level (and we will do so in Chapter 7). Moreover, there are other approaches to distributive justice that can be explored. The focus of this chapter will be on one of them, namely *equality of resources* as defined by Dworkin (1981b). The research question we aim to answer is thus the following:

*How can equality of resources be formalised and implemented
in PB with approval ballots?*

Let us briefly introduce equality of resources. The idea behind it is not to aim for a fair distribution of satisfaction, but instead to strive to invest the same effort into satisfying each voter. Interestingly, this approach does not suffer from the drawbacks we mentioned above since the amount of resources spent is a quantity we can measure objectively from the ballots.

This chapter presents a fairness theory for PB based on equality of resources. In the first section of this chapter, we will define the *share*, a measure of the amount of resources spent on each voter that we use to operationalise the concept of equality of resources (Section 4.1). From there, we will derive a first set of fairness criteria based on the idea that all agents deserve to enjoy a similar amount of share (Section 4.2). A second set of fairness criteria will then be introduced, all based on the idea that voters deserve a certain share by virtue of being part of cohesive groups (Section 4.3). At this point, we will investigate the relationships between the criteria we introduced, and other standard criteria from the literature (Section 4.4). We will conclude our study by presenting experimental results investigating how close to the share-based fairness criteria we can get in practice (Section 4.5). Finally, the main take-home messages from this chapter will be presented (Section 4.6).

4.1 The Share

We will focus on the standard model of PB in this chapter. We will thus use the notation introduced in Chapter 2.

The goal of this chapter is to explore fairness properties for PB that are defined in terms of equality of resources. We will propose several such fairness properties, all of them being based on the fundamental notion of the *share* of an agent.

Definition 4.1.1 (Share). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile \mathbf{A} , the share of a subset of projects $P \subseteq \mathcal{P}$ is defined as follows:*

$$\text{share}(I, \mathbf{A}, P) = \sum_{p \in P} \frac{c(p)}{|\{A \in \mathbf{A} \mid p \in A\}|}.$$

Moreover, the share of an agent $i \in \mathcal{N}$ is defined as:

$$\text{share}_i(I, \mathbf{A}, P) = \text{share}(I, \mathbf{A}, P \cap A_i).$$

When clear from context, we shall omit the arguments of I and \mathbf{A} in the notation of the share.

The share is thus a measure of the distribution of resources according to which all the supporters of a project p are allocated their share of p , defined as the cost of p divided by the number of its supporters. We interpret an agent's share as the amount of resources spent by the decision maker on trying to satisfy the needs of that agent. Note that it can also be interpreted as an influence index, measuring the impact of each agent on the final decision.

It is important to note that the share cannot be captured via independent cardinal utility functions as the share of an agent depends on the ballots submitted by the other agents.

This concept is so fundamental to this chapter that it deserves to also be presented in an example.

Example 4.1.2. Consider a PB instance I with three projects as described below, and a budget limit $b = 8$. The profile \mathbf{A} is composed of four ballots presented in the following table.

	p_1	p_2	p_3
Cost	6	2	2
A_1	✓	✓	✗
A_2	✓	✓	✗
A_3	✓	✗	✗
A_4	✗	✗	✓
$b = 8$			

If we select projects according to GREEDCOST, we would obtain the budget allocation $\pi = \{p_1, p_2\}$. The share of the agents would then be:

$$share_1(\pi) = share_2(\pi) = \frac{6}{3} + \frac{2}{2} = 3 \quad share_3(\pi) = \frac{6}{3} \quad share_4(\pi) = 0.$$

The agents thus have rather different shares. Somewhat anticipating the next section, this distribution of share can be deemed rather unfair. If the budget allocation $\pi' = \{p_1, p_3\}$ was to be selected instead, the share of all the agents would be of value 2. This definitely seems fairer. \triangle

4.2 The Fair Share

In this section we present our first set of fairness properties. They are all based on the idea a budget allocation can only be considered fair if every agent enjoys more share than a given threshold. This threshold, that will be introduced shortly after, is what we call the *fair share* of an agent.

4.2.1 Exact Fair Share

We first define the fair share property. It is based on the idea that each voter deserves $1/n$ of the budget—a fundamental idea familiar, for instance, from the classical fair division (“cake cutting”) literature (Robertson and Webb, 1998). So, a perfect allocation would give each voter a share of b/n . Since it could be that some voters do not approve enough projects for this to be possible, we correct for this corner case.

Definition 4.2.1 (Fair Share). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile \mathbf{A} , the fair share of agent $i \in \mathcal{N}$ is defined as:*

$$\text{fairshare}_i(I, \mathbf{A}) = \min\{b/n, \text{share}_i(I, \mathbf{A}, A_i)\}.$$

A budget allocation $\pi \in \text{FEAS}(I)$ is said to satisfy fair share (FS) if for every instance I , profile \mathbf{A} and agent $i \in \mathcal{N}$, we have:

$$\text{share}_i(I, \mathbf{A}, \pi) \geq \text{fairshare}_i(I, \mathbf{A}).$$

Note that if for a given agent $i \in \mathcal{N}$, it is the case that $\min\{b/n, \text{share}_i(I, \mathbf{A}, A_i)\} = \text{share}_i(I, \mathbf{A}, A_i)$, then π satisfies FS only if $A_i \subseteq \pi$. This ensures that FS is not trivially unsatisfiable because of the ballots not being large enough. As for the share, we omit the instance I and the profile \mathbf{A} from the fair share notation when they are clear from the context.

The notion of fair share is the fundamental fairness property studied in this chapter. A budget allocation satisfying FS indeed provides perfect equity of resources, *i.e.*, perfect fairness according to the perspective we adopt in this chapter.

Unfortunately, it is rather easy to see that for some instances no budget allocation would provide fair share. This notably implies that no rule can possibly satisfy FS.

Proposition 4.2.2. *There exists an instance $I = \langle \mathcal{P}, c, b \rangle$ with unit costs and a profile \mathbf{A} for which no budget allocation $\pi \in \text{FEAS}(I)$ provides FS.*

PROOF. Consider the instance I with two projects p_1 and p_2 , both of cost 1, and a budget limit $b = 1$. Let \mathbf{A} be a profile with agents named 1 and 2. Agent 1 approves only of p_1 and agent 2 only of p_2 . Then, for both agents, their fair share is $\min\{1/2, 1\} = 1/2$. However, at most one project can be selected, and whichever project is selected, the share of one of the agents would be 0. It is thus impossible to satisfy FS. \square

Still, as already witnessed in Example 4.1.2, there are some instances in which it is possible to provide fair share. It would thus be interesting to provide fair share when possible. However, we can show that there exists no polynomial-time computable rule that returns an FS allocation whenever one exists (unless $P = NP$). Indeed, checking the existence of a budget allocation satisfying FS is an NP-complete problem. The computational problem is formally defined as follows.

 FS SATISFIABILITY

Input: An instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile \mathbf{A} .
Question: Is there a budget allocation $\pi \in \text{FEAS}(I)$ that satisfies FS?

We show that this problem is NP-complete in the following.

Proposition 4.2.3. *The problem FS SATISFIABILITY is strongly NP-complete, even in the unit-cost setting.*

PROOF. It is clear that FS SATISFIABILITY is in NP, the certificate simply being the budget allocation itself. We show NP-hardness by a reduction from the 3-SET-COVER problem, known to be strongly NP-complete (Fürer and Yu, 2011). It is defined as follows.

 3-SET-COVER

Input: A universe $U = \{u_1, \dots, u_{|U|}\}$, a set S of 3-element subsets of U , and an integer $k \in \mathbb{N}$.
Question: Is there a subset S^* of S such that $\bigcup S^* = U$ and $|S^*| \leq k$?

Let (U, S, k) be an instance of 3-SET-COVER. We assume without loss of generality that $k \leq |U|$. Given (U, S, k) , we construct a PB instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile \mathbf{A} as follows:

- ▶ The set of projects is $\mathcal{P} = P \cup \{p^*\}$ with $P = \{p_1, \dots, p_{|S|}\}$, i.e., for every 3-element subsets of U denoted $s_j \in S$, there is a project p_j , and there is one auxiliary project p^* ;
- ▶ All projects have cost 1 and the budget limit is $b = k + 1$;
- ▶ The set of voters is $\mathcal{N} = \{1, \dots, 3|U| + 3\}$, i.e., there is a voter for every element of U and $2|U| + 3$ many auxiliary voters;
- ▶ The ballot of a “element” voter $i \in \{1, \dots, |U|\}$ is given by $A_i = \{p_j \in P \mid u_i \in s_j\}$, i.e., i approves the project representing the set s_j if and only if their corresponding element u_i is in s_j ;
- ▶ The ballot of an “auxiliary” voter $i \in \{|U| + 1, \dots, 3|U| + 3\}$ is $A_i = \{p^*\}$.

Given the instance I and the profile \mathbf{A} thus constructed, we claim that there is a budget allocation $\pi \in \text{FEAS}(I)$ that satisfies FS if and only if (U, S, k) is a positive instance of 3-SET-COVER.

Assume first that there is no set cover of size k . Thus, for any set $S^* \subseteq S$ of size k there is at least one element $u_i \in U$ that is not contained in any of

the sets in S^* . Then, by construction, for every budget allocation $\pi \in \text{FEAS}(I)$ containing k or fewer projects, there must be one “element” voter $i \in \{1, \dots, |U|\}$ with $\text{share}_i(\pi) = 0$. Moreover, for any budget allocation $\pi \in \text{FEAS}(I)$ such that $p^* \notin \pi$, all “auxiliary” voters $i \in \{|U| + 1, \dots, 3|U| + 3\}$ have share 0. Hence, no budget allocation containing at most $k + 1$ projects can satisfy FS. Finally, since the set of feasible budget allocations is $\text{FEAS}(I) = \{P \subseteq \mathcal{P} \mid |P| \leq k + 1\}$, this concludes the first part of the proof.

Now, assume that there exists $S^* \subseteq S$ that is a set cover of U of size k . We claim that the budget allocation $\pi = \{p_j \in P \mid s_j \in S^*\} \cup \{p^*\}$ satisfies FS. Given our assumption that $k \leq |U|$, we have:

$$\frac{b}{|\mathcal{N}|} = \frac{k + 1}{3|U| + 3} \leq \frac{|U| + 1}{3|U| + 3} = \frac{1}{3}.$$

Moreover, since $|s_j| = 3$ for all $s_j \in S$, exactly three voters approve of each project $p_j \in P$. Now, because S^* is a set cover, for each “element” voter $i \in \{1, \dots, |U|\}$, there is a project $p \in \pi$ such that $p \in A_i$. It follows that $\text{share}_i(\pi) \geq 1/3 \geq \text{fairshare}_i$, for all $i \in \{1, \dots, |U|\}$. In addition, for every $i \in \{|U| + 1, \dots, 3|U| + 3\}$ we have $A_i = \{p^*\}$, so we trivially have $\text{share}_i(\pi) = \text{share}_i(A_i) = \text{fairshare}_i$. It follows that π satisfies FS. \square

Because it is not always possible to find a budget allocation satisfying fair share, and as it is computationally hard to know whether we can, we weaken the fair share requirement in the hope of obtaining more positive results. Two weakenings will be presented in the following.

4.2.2 Fair Share up to One Project

The fact that FS cannot always be guaranteed is highly due to the indivisibility of the projects. It is indeed typical that to achieve strong fairness properties we would need to only partially select some projects. One standard way of relaxing properties that cannot be satisfied because of indivisibility issues is to consider their “up to one” variants, in which we require that selecting one extra project would allow us to satisfy the property. EJR-1 that we introduced in Definition 3.3.15 is an example of such property. This is also a standard approach in fair division (see, e.g., [Lipton, Markakis, Mossel and Saberi, 2004](#); [Budish, 2011](#)).

We follow the same approach here and introduce *fair share up to one project*.³²

³²Note that we diverge a bit from the standard way of defining “up to one” variants of fairness properties in PB here. Indeed we use a greater-or-equal instead of a strict inequality in the definition (contrary to EJR-1). One implication is that FS and FS-1 do not coincide on unit-cost instances. Our rationale is that adding one project guarantees that FS is satisfied, but no more than that.

Definition 4.2.4 (FS up to One Project). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile \mathbf{A} , a budget allocation $\pi \in \text{FEAS}(I)$ is said to satisfy fair share up to one project (FS-1) if, for every agent $i \in \mathcal{N}$, there is a project $p \in \mathcal{P}$ such that:*

$$\text{share}_i(I, \mathbf{A}, \pi \cup \{p\}) \geq \text{fairshare}_i(I, \mathbf{A}).$$

Thus, according to FS-1, a budget allocation can fail FS, but for each agent $i \in \mathcal{N}$, we should be able to select an extra project so that i exceeds their fair share.

Unfortunately, and surprisingly, FS-1, just as FS, cannot always be satisfied.

Proposition 4.2.5. *There exists an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile \mathbf{A} for which no budget allocation $\pi \in \text{FEAS}(I)$ provides FS-1.*

PROOF. Consider the following instance with three projects, all of cost 3 and a budget limit of $b = 5$. There are three agents, whose ballots are as displayed below.

	p_1	p_2	p_3
Cost	3	3	3
A_1	✓	✓	✗
A_2	✓	✗	✓
A_3	✗	✓	✓
$b = 5$			

Here the fair share of each agent is $5/3 \approx 1.67$. As a single project only yields a share of 1.5 to an agent who approves of it, for any agent to reach their fair share, two projects must be selected. However, a feasible budget allocation can select at most one project, meaning that for one agent none of the projects they approve of will be selected. So, even if we were to select an extra project for that agent, they would still not obtain their fair share. This is a violation of FS-1. \square

The difficulties keep on piling up for FS-1 as we can also show that it is computationally hard to check whether an FS-1 budget allocation exists, meaning that no PB rule running in polynomial time can satisfy FS-1 when possible (unless $\text{P} = \text{NP}$). Let us first introduce the formal computational problem.

FS-1 SATISFIABILITY

Input: An instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile \mathbf{A} .

Question: Is there a budget allocation $\pi \in \text{FEAS}(I)$ that satisfies FS-1?

Next, we show that FS-1 SATISFIABILITY is NP-complete.

Proposition 4.2.6. *The problem FS-1 SATISFIABILITY is strongly NP-complete, even in the unit-cost setting.*

PROOF. It is clear that FS-1 SATISFIABILITY is in NP. We show NP-hardness using the 3-SET-COVER problem (Fürer and Yu, 2011) once again. For the formal definition of the latter, see the proof of Proposition 4.2.3.

Let $\langle U, S, k \rangle$ be an instance of the 3-SET-COVER problem where $U = \{u_1, \dots, u_{|U|}\}$ is the universe, $S = \{s_1, \dots, s_{|S|}\}$ is a set of 3-element subsets of U , and $k \in \mathbb{N}$ is an integer. Without loss of generality, we make two assumptions: (i) every $u_i \in U$ appears in at least one $s_j \in S$; (ii) $|U|/3 \leq k < |U|$. Note that whenever one of these assumptions is violated, the answer of the 3-SET-COVER problem can easily be found in polynomial time.

We furthermore make the assumption that $k > |U|/3$. This assumption cannot be made without loss of generality. We justify it by the following reduction that shows that the problem is still NP-hard when $k > |U|/3$. Consider an instance $\langle U', S', k' \rangle$ of the 3-SET-COVER problem for which $k' = |U'|/3$. We construct a new instance $\langle U'', S'', k'' \rangle$ in which $k'' > |U''|/3$. It is such that:

- ▶ The universe is extended with four new elements: $U'' = U' \cup \{u''_1, u''_2, u''_3, u''_4\}$;
- ▶ Two new triplets are considered: $S'' = S' \cup \{\{u''_1, u''_2, u''_3\}, \{u''_2, u''_3, u''_4\}\}$;
- ▶ The solution size is increased by two: $k'' = k' + 2$.

Since $|U''| = |U'| + 4 = 3k' + 4$ and $k'' = k' + 2$, we clearly have $k'' > |U''|/3$. Moreover, a subset $S'_* \subseteq S'$ is a solution of the 3-SET-COVER problem for $\langle U', S', k' \rangle$ if and only if $S''_* = S'_* \cup \{\{u''_1, u''_2, u''_3\}, \{u''_2, u''_3, u''_4\}\}$ is a solution for $\langle U'', S'', k'' \rangle$. This shows that 3-SET-COVER remains NP-hard if we require that $k > |U|/3$.

We now get back to the original reduction and the instance $\langle U, S, k \rangle$. We distinguish between two cases based on the relative value of k and $|U|$, and construct different instances in each cases.

▷ Assume first that $1/3|U| < k \leq 2/3|U|$. We construct a PB instance I as follows. The set of voters is $\mathcal{N} = \{1, \dots, 2|U|\}$, i.e., there are two voters per element of U . The set of projects is $\mathcal{P} = \{p_j^1 \mid s_j \in S\} \cup \{p_j^2 \mid s_j \in S\}$, i.e., there are two projects per element of S . All projects have cost 1 and the budget limit is $b = k$. The profile \mathbf{A} is constructed with the following approval ballots:

$$\begin{aligned} A_i &= \{p_j^1 \mid u_i \in s_j\} \cup \{p_j^2 \mid u_i \in s_j\} && \text{for all } i \in \{1, \dots, |U|\}, \\ A_i &= \{p_j^1 \mid u_{i-|U|} \in s_j\} \cup \{p_j^2 \mid u_{i-|U|} \in s_j\} && \text{for all } i \in \{|U| + 1, \dots, 2|U|\}, \end{aligned}$$

i.e., agent i approves of the two projects representing the set s_j if and only if u_i (or $u_{i-|U|}$ if $i > |U|$) is in s_j . We now prove that there exists a suitable $S^* \subseteq S$

to answer the 3-SET-COVER problem, if and only if, there exists an FS-1 budget allocation in the instance I and profile \mathbf{A} previously described.

Observe that in I and \mathbf{A} , two projects need to be selected from each ballot so that all agents reach their fair share. Indeed, remember that we assumed $1/3|U| < k \leq 2/3|U|$. Since $b/n = k/2|U|$, we thus have $1/6 < b/n \leq 1/3$. Moreover, for every project $p \in \mathcal{P}$, there are exactly six agents approving of it, so for any agent $i \in \mathcal{N}$ approving of p , we have $share_i(\{p\}) = 1/6$. Given that every agent $i \in \mathcal{N}$ approves of at least 2 projects, we thus have $share_i(A_i) \geq 1/3$. Overall, we know that for every $i \in \mathcal{N}$, we have $1/6 < fairshare_i \leq 1/3$, so every agent needs two projects to reach their fair share.

Now, a budget allocation $\pi \in \text{FEAS}(I)$ satisfies FS-1 if and only if every agent has a non-zero share in π : According to the above, if an agent has a non-zero share then adding an extra project will always grant them their fair share; and if an agent has a share of zero, one would need to add two extra projects to π , which is not allowed in FS-1. All agents having a non-zero share is possible if and only if there exists a set $S^* \subseteq S$ of size at most k such that every elements of U appears in at least one element of S^* . This concludes the proof for this case.

▷ Assume now that $2/3|U| < k < |U|$. The PB instance I we construct is as follows. The set of voters is $\mathcal{N} = \{1, \dots, 2|U|\} \cup \{2|U|+1, \dots, 3k\}$, *i.e.*, there are two voters per element of U , together with a certain number of additional voters. Note that since $2/3|U| < k$, we do have $2|U|+1 \leq 3k$ and \mathcal{N} is thus well-defined. The set of projects is $\mathcal{P} = \{p_j^1 \mid s_j \in S\} \cup \{p_j^2 \mid s_j \in S\} \cup \{p^*\}$, *i.e.*, there are two projects per element of S , and an additional one p^* . All projects have cost 1 and the budget limit is $b = k$. The profile \mathbf{A} consists of the following approval ballot:

$$\begin{aligned} A_i &= \{p_j^1 \mid u_i \in s_j\} \cup \{p_j^2 \mid u_i \in s_j\} && \text{for all } i \in \{1, \dots, |U|\}, \\ A_i &= \{p_j^1 \mid u_{i-|U|} \in s_j\} \cup \{p_j^2 \mid u_{i-|U|} \in s_j\} && \text{for all } i \in \{|U|+1, \dots, 2|U|\}, \\ A_i &= \{p^*\} && \text{for all } i \in \{2|U|+1, \dots, 3k\}. \end{aligned}$$

Overall, agent i for $i \in \{1, \dots, 2|U|\}$ approves of the two projects representing the set s_j if and only if u_i (or $u_{i-|U|}$ if $i > |U|$) is in s_j ; while the additional agents all approve only of project p^* .

Let us discuss this construction. By definition we have $b/n = k/3k = 1/3$. As for the previous case, for every agent $i \in \{1, \dots, 2|U|\}$, the share of every project they approve of is $1/6$. They thus deserve a share of at least $1/3$ to get their fair share, which can only be done by selecting at least two projects. Observe in addition that every budget allocation $\pi \in \text{FEAS}(I)$ satisfies the FS-1 condition for agents $i \in \{2|U|+1, \dots, 3k\}$ as the project p^* (the only project they approve of) can always be added, if it is not already in π , by virtue of FS-1.

Overall, an allocation $\pi \in \text{FEAS}(I)$ satisfies FS-1 if and only if every agent $i \in \{1, \dots, 2|U|\}$ has a non-zero share in π . Such a π exists if and only if there is

a set $S^* \subseteq S$ of size at most k such that every element of U appears in at least one element of S^* . This concludes the proof. \square

Our investigation of FS-1 is now over. In a nutshell, the picture did not get brighter when weakening FS into FS-1. In the following section, we will study another weakening of FS that we will prove to be satisfiable, in polynomial time even!

4.2.3 Local Fair Share

Our second weakening of FS is inspired by the local variant of PJR introduced by [Aziz, Lee and Talmon \(2018\)](#), called *Local BPJR*. We introduced this concept in Definition 3.3.22. We now provide the definition of Local-FS; we will discuss it afterwards.

Definition 4.2.7 (Local-FS). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile \mathbf{A} , a budget allocation $\pi \in \text{FEAS}(I)$ is said to satisfy local fair share (Local-FS) if there is no project $p \in \mathcal{P} \setminus \pi$ such that, there is an agent $i \in \mathcal{N}$ with $p \in A_i$, and for all agents $i \in \mathcal{N}$ with $p \in A_i$, we have:*

$$\text{share}_i(\pi \cup \{p\}) < \text{fairshare}_i.$$

Intuitively, if there exists a project p that could be added to the budget allocation π without any of its supporters receiving at least their fair share, then every supporter of p receives strictly less than their fair share and one of the following holds:

- p can be selected without exceeding the budget limit b ;
- some voter i^* not approving of p receives more than their fair share.

In the first case, it is clear that p should be selected and thus π must be deemed unfair. In the second case, it might be considered fairer to exchange one project supported by i^* with project p . In this sense, the property can be seen as an “upper quota” property, as we have to add projects such that no voter receives more than their fair share as long as possible.

In contrast to FS-1 and FS, we can always find an allocation that satisfies Local-FS. Indeed, we can show that MES used with the share satisfies Local-FS.

Theorem 4.2.8. *MES[share] satisfies Local-FS.*

PROOF. Let us recall some elements of the definition of an MES rule (Definition 2.3.3) we provided in Chapter 2. For a given subset of projects $P \subseteq \mathcal{P}$ and an agent $i \in \mathcal{N}$, $\ell_i(P)$ represents the load of agent i , that is, the amount of virtual money they have spent. Moreover, for $P \subseteq \mathcal{P}$, $p \in P$ and $\alpha \in \mathbb{R}$, $\gamma_i(P, \alpha, p)$

denotes the contribution of agent $i \in \mathcal{N}$ for project p for the affordability factor α if projects P have already been selected. Let us introduce one additional concept: given a budget allocation π and a scalar $\alpha > 0$, we say that agent $i \in \mathcal{N}$ *contributes in full* to project $p \in A_i$ if we have $\gamma_i(P, \alpha, p) = \alpha \cdot \text{share}_i(\{p\})$.

During a run of $\text{MES}[\text{share}]$, all the supporters of a project $p \in \mathcal{P}$ contribute in full to p if and only if p is 1-affordable. In this case, for all supporters i of p , we have $\ell_i(\{p\}) = \text{share}(\{p\}, i)$. Moreover, if a project p is α -affordable but at least one voter cannot contribute in full to p , then $\alpha > 1$. $\text{MES}[\text{share}]$ only terminates when no project is α -affordable, for any α . Therefore, there is a round where no project p is 1-affordable. Let k be the first such round and let π_k be the budget allocation before round k . It follows that every project in π_k was 1-affordable and hence $\ell_i(\pi_k) = \text{share}_i(\pi_k)$ for all $i \in \mathcal{N}$. As no project p is 1-affordable in round k , there is no projects in $\mathcal{P} \setminus \pi_k$ for which all the supporters contribute in full to. Thus, for every $p \in \mathcal{P} \setminus \pi_k$, there is a voter $i \in \mathcal{N}$ such that $b/n - \ell_i(\pi_k) < \text{share}_i(\{p\})$. Since $\ell_i(\pi_k) = \text{share}_i(\pi_k)$ and the share is additive, it follows that $\text{share}(\pi_k \cup \{p\}, i) > b/n$. So π_k already satisfies Local-FS. As $\text{MES}[\text{share}]$ returns an allocation π with $\pi_k \subseteq \pi$, it satisfies Local-FS. \square

In fact, the proof of Theorem 4.2.8 establishes a slightly stronger statement than Local-FS: there is no project $p \in \mathcal{P} \setminus \pi$ such that for all agents $i \in \mathcal{N}$ with $p \in A_i$ we have $\text{share}_i(\pi \cup \{p\}) \leq b/n$. In other words, any project added to π gives at least one voter *more* than their fair share. This would correspond to having a strict inequality in the definition of Local-FS.

We have now found a relaxation of FS that can always be satisfied. Given that $\text{MES}[\text{share}]$ runs in polynomial time, we can even find a budget allocation satisfying Local-FS efficiently. In the coming section, we focus on different relaxations of FS.

4.3 The *Justified* Share

Local-FS and FS-1 require the outcome to be, in some sense, close to satisfying FS. Another idea for weakening FS is to spend on a voter only the resources they can claim to deserve. In Chapter 3, we already surveyed a long list of fairness properties based on this ideas, namely the axioms based on *justified representation* (see Section 3.3.1). We follow the same approach here.

Our blue-print for this section is the adaptation of EJR to the context of PB (see Definition 3.3.10). Cohesive groups will thus play an important role here. Remember that they have been defined in Definition 3.3.9. Ideally, we would want to satisfy the following property.

Definition 4.3.1 (Strong Extended Justified Share). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile \mathbf{A} , a budget allocation $\pi \in \text{FEAS}(I)$ is said to satisfy strong extended*

justified share (*Strong-EJS*) if for all $P \subseteq \mathcal{P}$, all P -cohesive groups N and all agents $i \in N$, we have:

$$\text{share}_i(\pi) \geq \text{share}_i(P).$$

The idea behind Strong-EJS is the following: since every P -cohesive group S controls enough budget to fund P , every agent in S deserves to enjoy at least as much share as what they would have gotten if P had been the outcome.

Intuitively, Strong-EJS is very similar to Strong-EJR (Definition 3.3.10), a property that is known not to be always satisfiable (see Proposition 3.3.11). The same holds for Strong-EJS: there exist instances for which no budget allocation satisfies this axiom.

Proposition 4.3.2. *There exists an instance $I = \langle \mathcal{P}, c, b \rangle$ with unit costs and a profile \mathbf{A} for which no budget allocation $\pi \in \text{FEAS}(I)$ satisfies Strong-EJS.*

PROOF. Consider the instance $I = \langle \mathcal{P}, c, b \rangle$ with four projects p_1, p_2, p_3 and p_4 , all of cost 1, and a budget limit $b = 2$. The profile \mathbf{A} is composed of the four ballots presented below.

	p_1	p_2	p_3	p_4
Cost	1	1	1	1
A_1	✓	✗	✗	✓
A_2	✓	✓	✗	✗
A_3	✓	✗	✓	✗
A_4	✗	✓	✓	✗
$b = 2$				

The group of agents $\{1, 2, 3\}$ is $\{p_1\}$ -cohesive, the group $\{2, 4\}$ is $\{p_2\}$ -cohesive, and the group $\{3, 4\}$ is $\{p_3\}$ -cohesive. Hence one needs to select all three projects to satisfy Strong-EJS, which is not possible within the budget limit. \square

Observe that in this scenario it is not even possible to guarantee each P -cohesive group the same *average* share as they receive from P . Moreover, it is interesting to note that in the above example, all ballots are feasible and exhaustive, so these ballot restrictions would also not help satisfying Strong-EJS.

All of the above motivates us to weaken Strong-EJS, and to introduce (simple) EJS.

Definition 4.3.3 (Extended Justified Share). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile \mathbf{A} , a budget allocation $\pi \in \text{FEAS}(I)$ is said to satisfy extended justified share (EJS), if for all $P \subseteq \mathcal{P}$ and all P -cohesive groups N , there is an agent $i \in N$ for whom $\text{share}_i(\pi) \geq \text{share}_i(P)$.*

The difference between Strong-EJS and EJS is the switch from a universal to an existential quantifier: for the former, we impose a lower bound on the share of *every* agent in a cohesive group, while for the latter this lower bound only applies to one agent of each cohesive group. Therefore, in the counterexample presented in the proof of Proposition 4.3.2 both $\pi = \{p_1, p_3\}$ and $\pi' = \{p_2, p_3\}$ satisfy EJS, as either agent 3 or agent 4 satisfies the share requirement for every cohesive group.

Before moving on to the satisfiability of EJS, we present an interesting example showing that EJS and $\text{EJR}[sat^{card}]$ (Definition 3.3.12), while similar in spirit, do not coincide, even in the unit-cost setting.

Example 4.3.4. Consider an instance I with six projects, all of cost 1, and a budget limit $b = 4$. The profile \mathbf{A} is composed of the following four approval ballots:

	p_1	p_2	p_3	p_4	p_5	p_6
Cost	1	1	1	1	1	1
A_1	✓	✓	✓	✗	✗	✗
A_2	✓	✓	✗	✓	✗	✗
A_3	✗	✗	✗	✓	✓	✓
A_4	✗	✗	✗	✓	✓	✓
$b = 4$						

Then, 1 and 2 together form a $\{p_1, p_2\}$ -cohesive group. Similarly, 3 and 4 form a $\{p_4, p_5\}$ cohesive group. We claim that $\pi = \{p_3, p_4, p_5, p_6\}$ satisfies EJS but not $\text{EJR}[sat^{card}]$ while $\pi' = \{p_1, p_4, p_5, p_6\}$ satisfies $\text{EJR}[sat^{card}]$ but not EJS.

Let us first consider $\pi = \{p_3, p_4, p_5, p_6\}$. There, we have $|A_1 \cap \pi| = |A_2 \cap \pi| = 1$, so $\text{EJR}[sat^{card}]$ is not satisfied. On the other hand, $share_1(\pi) = 1 = share_1(\{p_1, p_2\})$ and $share_3(\pi) > share_3(\{p_4, p_5\})$. Therefore, π satisfies EJS.

Now consider $\pi' = \{p_1, p_4, p_5, p_6\}$. We have:

$$share_1(\pi') = \frac{1}{2}, \quad share_2(\pi') = \frac{5}{6}, \quad share_1(\{p_1, p_2\}) = share_2(\{p_1, p_2\}) = 1.$$

Therefore, the $\{p_1, p_2\}$ -cohesive group $\{1, 2\}$ witnesses that π' violates EJS. On the other hand, since $|A_2 \cap \pi'| = 2$ and $|A_3 \cap \pi'| = 3$, $\text{EJR}[sat^{card}]$ is satisfied by π' . \triangle

We have now prepared everything for studying EJS. We will start with the question of the satisfiability of EJS.

4.3.1 Achieving EJS

We now turn to the satisfiability of EJS. We prove using a standard argument (through the use of a greedy cohesive procedure) that it is always satisfiable.

Algorithm 1: Greedy Cohesive Procedure for EJS**Input:** An instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile \mathbf{A} **Output:** A budget allocation $\pi \in \text{FEAS}(I)$ satisfying EJSInitialise π and N^* as the empty set: $\pi \leftarrow \emptyset, N^* \leftarrow \emptyset$ **while** there exists an $N \subseteq \mathcal{N} \setminus N^*$ with $N \neq \emptyset$ and a $P \subseteq \mathcal{P} \setminus \pi$ with $P \neq \emptyset$, such that N is P -cohesive **do** Select $N \subseteq \mathcal{N} \setminus N^*$ and $P \subseteq \mathcal{P} \setminus \pi$ such that:

$$(N, P) \in \underset{\substack{(N', P') \in 2^{\mathcal{N} \setminus N^*} \times 2^{\mathcal{P} \setminus \pi} \\ N' \text{ is } P' \text{-cohesive}}}{\arg \max} \max_{i \in N'} \text{share}(P', i)$$

 Select the projects in P : $\pi \leftarrow \pi \cup P$ Agents in N have been satisfied: $N^* \leftarrow N^* \cup N$ **return** the budget allocation π

Proposition 4.3.5. For every instance $I = \langle \mathcal{P}, c, b \rangle$ and every profile \mathbf{A} , there exists a budget allocation $\pi \in \text{FEAS}(I)$ that satisfies EJS.

PROOF. We show that Algorithm 1 computes a feasible budget allocation that satisfies EJS. Let us consider an arbitrary instance $I = \langle \mathcal{P}, c, b \rangle$ and profile \mathbf{A} . Informally, Algorithm 1—that we refer to as the *greedy cohesive procedure for EJS*—selects projects in a greedy fashion, each time selecting the ones that are involved in a cohesive group with the highest share requirement. The procedure goes on until all cohesive groups have been satisfied, never considering agents or projects more than once.

We first show that the budget allocation returned by the algorithm is feasible.

Claim 4.3.6. Algorithm 1 always returns a feasible budget allocation.

Proof: Consider the run of the algorithm on I and \mathbf{A} and denote by π the budget allocation returned. Assume that the while-loop is run k times. Let us call (N_j, P_j) the sets of agents and projects that are selected during the j -th run of the while-loop, for all $j \in \{1, \dots, k\}$. We then have:

$$c(\pi) = \sum_{j=1}^k c(P_j) \leq \sum_{j=1}^k \frac{|N_j| \cdot b}{n} = \frac{b}{n} \cdot \sum_{j=1}^k |N_j| \leq b.$$

The first equality comes from the fact that P_1, \dots, P_k is a partition of π . The first inequality is derived from the fact that N_j is a P_j -cohesive group, for all $j \in \{1, \dots, k\}$ (it is an inequality because for any of the projects $p \in P_j$, some agents outside of N_j may approve of it; $c(p)$ can thus be split among more than

$|N_j|$ agents). The final inequality is linked to the fact all the N_1, \dots, N_k are pairwise disjoint. Overall, the outcome of Algorithm 1 always is a feasible budget allocation. \diamond

Let us now prove that the algorithm does compute an EJS budget allocation.

Claim 4.3.7. *The budget allocation π returned by Algorithm 1 on I and A satisfies EJS.*

Proof: Assume towards a contradiction that π violates EJS. Then, there must exist some $N^\dagger \subseteq \mathcal{N}$ and $P^\dagger \subseteq \mathcal{P}$ such that N^\dagger is P^\dagger -cohesive but also such that, for all agents $i \in N^\dagger$, we have $share_i(\pi) < share_i(P^\dagger)$. Note that, if $P^\dagger \not\subseteq \pi$, this means that at the end of the algorithm either one agent $i \in N^\dagger$ has been satisfied ($i \in N^*$ when the algorithm returns) or that one project $p \in P^\dagger$ has been selected ($p \in \pi$ when the algorithm returns). We distinguish these two cases.

▷ First, consider the case where one agent has been satisfied by the end of the algorithm. Using the same notation as for the previous claim, there exists then a smallest $j \in \{1, \dots, k\}$ such that there exist $i^* \in N^\dagger \cap N_j$. Note that this implies that $P_j \subseteq A_{i^*}$ since i^* is part of the group of agents N_j , that is P_j -cohesive. Given that (N^\dagger, P) has not been selecting during that run of the while loop, it means that:

$$\max_{i' \in N_j} share_{i'}(P_j) \geq \max_{i \in N^\dagger} share_i(P^\dagger).$$

Since the cost of a project is split equality among its supporters, it is easy to observe that for any P -cohesive group N , and for every two agents $i, i' \in N$, we have $share_i(P) = share_{i'}(P)$. Moreover, for any $P' \subseteq P \subseteq \mathcal{P}$, we also have that $share_i(P') \leq share_i(P)$ for any agent $i \in \mathcal{N}$. Overall, for our distinguished agent $i^* \in N^\dagger \cap N_j$, we have:

$$\begin{aligned} share_{i^*}(\pi) &\geq \max_{i' \in N_j} share_{i'}(P_j) \\ &\geq \max_{i \in N^\dagger} share_i(P^\dagger) \\ &\geq share_{i^*}(P^\dagger), \end{aligned}$$

which contradicts the fact that π fails EJS.

▷ Let us now consider the second case, *i.e.*, when $P^\dagger \cap \pi \neq \emptyset$ but $P^\dagger \not\subseteq \pi$. In this case, it is important to see that if N^\dagger is P^\dagger -cohesive, then it is also P -cohesive for all $P \subseteq P^\dagger$. Then, we can run the same proof considering the $(P^\dagger \setminus \pi)$ -cohesive group N^\dagger . Iterating this argument, would either lead to the conclusion that $P^\dagger \subseteq \pi$, a contradiction, or to another contradiction due to the first case we considered (when some agent of N^\dagger is already satisfied). \diamond

To conclude the proof, note that Algorithm 1 always terminates. Indeed, after each run of the while-loop, at least one agent is added to the set N^* . Moreover, if $N^* = \mathcal{N}$, the condition of the while-loop would be violated and the algorithm would terminate. Overall at most n runs through the while-loop can occur. \square

The fact that EJS can always be satisfied is a good thing in general. However, the reader who has gone through the proof will have noticed that the procedure we devise may require exponential time to run. This drawback is unfortunately unavoidable, unless $P = NP$, as we show next.

Theorem 4.3.8. *There is no strongly-polynomial time algorithm that, given an instance I and a profile \mathbf{A} as input, always computes a budget allocation satisfying EJS, unless $P = NP$.*

PROOF. Assume, that there is an algorithm \mathbb{A} that always computes an allocation satisfying EJS in strongly-polynomial time. We will prove that we could use \mathbb{A} to solve the TARGETED SUBSET-SUM problem, known to be NP-hard (Karp, 1972; Garey and Johnson, 1979).

TARGETED SUBSET-SUM

Input: A finite set $Z \subseteq \mathbb{N}$ and a target t .

Question: Is there a non-empty $Z' \subseteq Z$ such that $\sum_{z \in Z'} z = t$?

Given an instance of the problem $Z = \{z_1, \dots, z_m\}$ and t , we construct $I = \langle \mathcal{P}, c, b \rangle$ and \mathbf{A} as follows. We have m projects $\mathcal{P} = \{p_1, \dots, p_m\}$ with the following cost function $c(p_j) = z_j$ for all $j \in \{1, \dots, m\}$ and a budget limit $b = t$. There is moreover only one agent, who approves of all the projects.

Now, (Z, t) is a positive instance of TARGETED SUBSET-SUM if and only if there is a budget allocation $\pi \in \text{FEAS}(I)$ that cost is exactly b . If such an allocation π exists, then the one voter is π -cohesive. Therefore, any allocation π' that satisfies EJS must give that voter $\text{share}_1(\pi') \geq \text{share}_1(\pi) = c(\pi)$. Hence, (Z, t) is a positive instance of SUBSET-SUM if and only if $c(\mathbb{A}(I, \mathbf{A})) = b$. We have thus presented a way to solve the TARGETED SUBSET-SUM problem in strongly-polynomial time using algorithm \mathbb{A} . This is only possible if $P = NP$. \square

Interestingly, we can compute EJS budget allocations in FPT-time, when parameterized by the number of projects. This was already observed in the unit-cost setting for the EJR[sat^{card}] (Aziz, Brill, Conitzer, Elkind, Freeman and Walsh, 2017).

Proposition 4.3.9. *For every instance $I = \langle \mathcal{P}, c, b \rangle$ and every profile \mathbf{A} , we can compute a budget allocation $\pi \in \text{FEAS}(I)$ that satisfies EJS in time in $\mathcal{O}(n \cdot 2^{|\mathcal{P}|})$.*

$$\begin{aligned}
& \mathbf{maximise} \quad \epsilon && (4.1) \\
& \mathbf{subject\ to:} \\
& x_i \in \{0, 1\} && \text{for all } i \in \mathcal{N} \quad (4.2) \\
& z_p \in \{0, 1\} && \text{for all } p \in \mathcal{P} \quad (4.3) \\
& z_p + x_i - 1 \leq \mathbb{1}_{p \in A_i} && \text{for all } i \in \mathcal{N} \text{ and } p \in \mathcal{P} \quad (4.4) \\
& \frac{1}{n} \sum_{i \in \mathcal{N}} x_i \geq \frac{1}{b} \sum_{p \in \mathcal{P}} z_p \cdot c(p) && (4.5) \\
& x_i \cdot share_i(\pi) + \epsilon \leq \sum_{p \in \mathcal{P}} z_p \cdot share(\{p\}) && \text{for all } i \in \mathcal{N} \quad (4.6)
\end{aligned}$$

Figure 4.3.1: An integer linear program for verifying whether a budget allocation π satisfies EJS.

PROOF. We prove that Algorithm 1 has the suitable running time when run on instance $I = \langle \mathcal{P}, c, b \rangle$ and profile \mathbf{A} . The first thing to note is that at least one agent is added to N^* during each run of the while-loop. Thus, there are at most n iterations of the while-loop.

Let us have a closer look at what happens inside the while-loop. The main computational task here is the maximisation that goes through all subsets of \mathcal{N} and \mathcal{P} . The trick resides in the fact that we can actually avoid going through all subsets of agents. Indeed, consider a subset of projects $P \subseteq \mathcal{P}$ and let $N \subseteq \mathcal{N}$ be the largest set of agents such that for all $i \in N$, we have $P \subseteq A_i$. Note that such a set N can be efficiently computed (by going through all the approval ballots). Now, if there exists a group of agents that is P -cohesive, then for sure N also is P -cohesive. Moreover, note that for any P -cohesive group N' , and for every two agents $i \in N$ and $i' \in N \cup N'$, we have $share_i(P) = share_{i'}(P)$. Overall, one can, without loss of generality, only consider the group of agents N when considering the subset of projects P . This implies that the maximisation step can be computed by going through all the subsets of projects and, for each of them, by only considering a single subset of agents (that is efficiently computable). \square

To conclude this section we investigate the problem of verifying whether a given budget allocation satisfies EJS. It is easy to prove that, as is the case for EJR[*sat*^{card}] (Aziz, Brill, Conitzer, Elkind, Freeman and Walsh, 2017), this problem is in coNP. However, we can define an ILP solving it. A suitable one is presented in Figure 4.3.1. It searches for a set of projects $P \subseteq \mathcal{P}$ and a set of agents $N \subseteq \mathcal{N}$ that certifies a violation of the EJS property, *i.e.*, N is P -cohesive and all voters receive a strictly larger share from P than from π . For any agent $i \in \mathcal{N}$, we use variable x_i to indicate

whether $i \in N$, and for any project $p \in \mathcal{P}$, variable z_p to indicate whether $p \in P$. Conditions (4.4) and (4.5) enforce that N is indeed P -cohesive. Condition (4.6) implies that $share_i(\pi) < share_i(P)$ for every agent $i \in N$. The inequality in Condition (4.6) is only strict for $\epsilon > 0$. Consequently, π fails EJS if and only if this ILP yields a solution with $\epsilon > 0$.

We have seen that EJS can always be satisfied. However, this is not entirely satisfactory as no tractable rule can satisfy it. Unfortunately, in many PB applications, the use of intractable rules is not practical due to the large instance sizes. Therefore, we try to find fairness notions that can be satisfied in polynomial time by relaxing EJS.

4.3.2 Relaxing EJS

We first consider relaxing EJS in the usual “up-to-one” way. It requires at least one agent in every cohesive group to be at most one project away from being satisfied.³³

Definition 4.3.10 (EJS-1). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile \mathbf{A} , a budget allocation $\pi \in \text{FEAS}(I)$ is said to satisfy extended justified share up to one project (EJS-1) if for all $P \subseteq \mathcal{P}$ and all P -cohesive groups N there is an agent $i \in N$ for which there exists a project $p \in \mathcal{P}$ such that $share_i(\pi \cup \{p\}) \geq share_i(P)$.*

Because the rule $\text{MES}[share]$ is a variant of MES used with a satisfaction function that behaves in an additive manner, the proof of Peters, Pierczyński and Skowron (2021) that MES satisfies EJR-1 can trivially be adapted to our setting. This shows that $\text{MES}[share]$ satisfies EJS-1.

Proposition 4.3.11. *$\text{MES}[share]$ satisfies EJS-1.*

Interestingly, this implies, together with Example 4.3.4, that already in the unit-cost case $\text{MES}[sat^{card}]$ and $\text{MES}[share]$ are different rules.

In an attempt to go further, we also introduce a local variant of EJS, based on a similar motivation as Local-FS.

Definition 4.3.12 (Local-EJS). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile \mathbf{A} , a budget allocation $\pi \in \text{FEAS}(I)$ is said to satisfy local extended justified share (Local-EJS), if there is no P -cohesive group N , where $P \subseteq \mathcal{P}$, for which there exists a project $p \in P \setminus \pi$ such that for all agents $i \in N$ we have $share_i(\pi \cup \{p\}) < share_i(P)$.*

The idea behind Local-EJS is that there is no P -cohesive group N that can claim that they could “afford” another project p without a single voter in N receiving more share than they deserve due to their P -cohesiveness. In this sense, any allocation that satisfies Local-EJS is a local optimum for any P -cohesive group.

³³As was the case for FS-1, the definition of EJS-1 is slightly non-standard as we require that $share_i(\pi \cup \{p\}) \geq share_i(P)$ and not a strict inequality, as used for EJR-1 (see Definition 3.3.15).

We claimed that studying Local-EJS would allow us to go further than EJS-1, this is because in our setting Local-EJS is equivalent to a notion that could be called “EJS up to any project”, mimicking EJR-X (see Definition 3.3.15).

Proposition 4.3.13. *Let $I = \langle \mathcal{P}, c, b \rangle$ be an instance and \mathbf{A} a profile. An allocation π satisfies Local-EJS if and only if for every $P \subseteq \mathcal{P}$ and P -cohesive group N there exists an agent i such that for all projects $p \in P \setminus \pi$ we have $share_i(\pi \cup \{p\}) \geq share_i(P)$.*

PROOF. It is clear that the statement above implies Local-EJS. Now, let π be an allocation that satisfies Local-EJS, let $P \subseteq \mathcal{P}$ be a set of projects and N a P -cohesive group. Let $i^* \in N$ be an agent with maximal share from π in N . Consider any project $p \in P \setminus \pi$. By Local-EJS there is an agent i_p such that $share_{i_p}(\pi \cup \{p\}) > share_{i_p}(P)$. By the choice of i^* , we also have $share_{i^*}(\pi) \geq share_{i_p}(\pi)$. By the additivity of the share, it follows that $share_{i^*}(\pi \cup \{p\}) > share_{i^*}(P)$. This proves the statement. \square

Local-EJS can thus also be interpreted as being the strengthening of EJS-1 with a universal quantifier on the project to add, rather than an existential one. From this equivalence, it is easy to see that EJS implies Local-EJS that, in turn, implies EJS-1.

Since it satisfies EJS-1, $MES[share]$ is the most promising candidate for a rule that could satisfy Local-EJS. Unfortunately, it fails it, as the next example shows.

Example 4.3.14. Consider an instance I and a profile \mathbf{A} with five projects, a budget limit $b = 20$, and the costs and approval ballots as described below.

	p_1	p_2	p_3	p_4	p_5
Cost	8	5	2	2	10
A_1	✓	✓	✓	✓	✗
A_2	✓	✓	✓	✓	✗
A_3	✗	✗	✓	✓	✓
A_4	✗	✗	✓	✓	✓
$b = 8$					

Here, and with a suitable tie-breaking rule, $MES[share]$ will return the budget allocation $\pi = \{p_2, p_3, p_5\}$. Note that voters 1 and 2 are $\{p_1, p_4\}$ -cohesive and would thus deserve to enjoy a share of 4.5. However, if we add p_4 to π , voters 1 and 2 would only have a share of 3.5, showing that π fails Local-EJS. \triangle

Whether Local-EJS can always be satisfied in polynomial time remains an important open question.

Getting back to the comparison between $MES[sat^{card}]$ and $MES[share]$, we observe a crucial difference in the unit-cost setting: $MES[share]$ does not satisfy EJS, while $MES[sat^{card}]$ satisfies EJR[sat^{card}] (Peters and Skowron, 2020).

Example 4.3.15. Assume that there are two voters 1 and 2, and three projects p_1, p_2 and p_3 , all of cost 1. The budget limit is $b = 2$. Voter 1 approves of p_1 and p_3 and voter 2 of p_2 and p_3 . Then voter 1 is $\{p_1\}$ -cohesive and hence deserves a share of 1, the same applies to voter 2 and $\{p_2\}$. Nevertheless, with a suitable tie-breaking rule, $\text{MES}[\text{share}]$ would first select p_3 . In that case, neither $\{p_1, p_3\}$, nor $\{p_2, p_3\}$ would satisfy EJS, as at least one voter will have a share of only $1/2$. \triangle

Still focusing on the unit-cost setting, we can show that $\text{MES}[\text{share}]$ satisfies Local-EJS there (even though it fails EJS).

Theorem 4.3.16. $\text{MES}[\text{share}]$ satisfies Local-EJS in the unit-cost setting.

PROOF. Consider an instance $I = \langle \mathcal{P}, c, b \rangle$ in which all projects have cost 1, and a profile \mathbf{A} for I . Let $\pi = (p_1, \dots, p_k)$ be the budget allocation returned by $\text{MES}[\text{share}]$ on I and \mathbf{A} . We assume that during the run of $\text{MES}[\text{share}]$ project p_1 was selected first, project p_2 second, etc. For any round $j \in \{1, \dots, k\}$, we denote by $\pi_j = \{p_1, \dots, p_j\}$ the set of projects selected at the end of round j . Finally, consider $N \subseteq \mathcal{N}$, an arbitrary P -cohesive group, for some $P \subseteq \mathcal{P}$. We show that π satisfies the condition of Local-EJS for N and P . If $P \subseteq \pi$ then Local-EJS is satisfied for N and P by definition. We will thus assume that $P \not\subseteq \pi$.

Let k^* be the first round after which there exists a voter $i^* \in N$ whose load is larger than $b/n - 1/|N|$. Such a round must exist as otherwise the voters in N could afford another project from P . As we assumed $P \not\subseteq \pi$, this would mean that $\text{MES}[\text{share}]$ would not have stopped. Let $\pi^* = \pi_{k^*}$ and consider an arbitrary project $p^* \in P \setminus \pi^*$. Our goal is to prove that π^* satisfies Local-EJS for N , that is:

$$\begin{aligned}
& \text{share}_{i^*}(\pi^* \cup \{p^*\}) > \text{share}_{i^*}(P) \\
\Leftrightarrow & \text{share}_{i^*}(\pi^*) > \text{share}_{i^*}(P \setminus \{p^*\}) \\
\Leftrightarrow & \text{share}_{i^*}(\pi^* \cap P) + \text{share}_{i^*}(\pi^* \setminus P) > \\
& \text{share}_{i^*}(P \cap \pi^*) + \text{share}_{i^*}(P \setminus (\pi^* \cup \{p^*\})) \\
\Leftrightarrow & \text{share}_{i^*}(\pi^* \setminus P) > \text{share}_{i^*}(P \setminus (\pi^* \cup \{p^*\})). \tag{4.7}
\end{aligned}$$

We will work on each side of inequality (4.7) to prove that it indeed holds.

We start by the left-hand side of (4.7). Let us first introduce some notation that will allow us to reason in terms of share per unit of load. For a project $p \in \pi$, we denote by $\alpha(p)$ the smallest $\alpha \in \mathbb{R}_{>0}$ such that p was α -affordable when $\text{MES}[\text{share}]$ selected it. Moreover, we define $q(p)$ —the share that a voter that contributes fully to p gets per unit of load—as $q(p) = 1/\alpha(p)$.

Since before round k^* , agent i^* contributed in full for all projects in π^* (as $\ell_{i^*} < b/|N|$ after each round $1, \dots, k^*$), we know that $\alpha(p) \cdot \text{share}_{i^*}(\{p\})$ equals

the contribution of i^* for p , and that for any $p \in \pi^*$. We thus have:

$$\begin{aligned} \text{share}_{i^*}(\pi^* \setminus P) &= \sum_{p \in \pi^* \setminus P} \text{share}_{i^*}(\{p\}) \\ &= \sum_{p \in \pi^* \setminus P} \alpha(p) \cdot \frac{\text{share}_{i^*}(\{p\})}{\alpha(p)} \\ &= \sum_{p \in \pi^* \setminus P} \gamma_{i^*}(p) \cdot q(p), \end{aligned} \quad (4.8)$$

where $\gamma_{i^*}(p)$ denotes the contribution of i^* to any $p \in \pi$, defined such that if p has been selected at round j , i.e., $p = p_j$, then $\gamma_{i^*}(p) = \gamma_{i^*}(\pi_j, \alpha(p_j), p_j)$ (see Definition 2.3.3 for the definition of $\gamma(\cdot, \cdot, \cdot)$).

Now, let us denote by q_{\min} the smallest $q(p)$ for any $p \in \pi^* \setminus P$. From (4.8), we get:

$$\text{share}_{i^*}(\pi^* \setminus P) \geq q_{\min} \sum_{p \in \pi^* \setminus P} \gamma_{i^*}(p). \quad (4.9)$$

We now turn to the right-hand side of (4.7). We introduce some additional notation for that. For every project $p \in P$, we denote by $q^*(p)$ the share per load that a voter in N receives if only voters in N contribute to p , and they all contribute in full to p , defined as:

$$q^*(p) = \frac{\text{share}(\{p\})}{1/|N|} = \frac{|N|}{|\{A \in \mathbf{A} \mid p \in A\}|}.$$

We have then:

$$\begin{aligned} \text{share}_{i^*}(P \setminus (\pi^* \cup \{p^*\})) &= \sum_{p \in P \setminus (\pi^* \cup \{p^*\})} \text{share}_{i^*}(\{p\}) \\ &= \sum_{p \in P \setminus (\pi^* \cup \{p^*\})} \frac{\text{share}_{i^*}(\{p\})}{1/|N|} \cdot \frac{1}{|N|} \\ &= \sum_{p \in P \setminus (\pi^* \cup \{p^*\})} q^*(p) \cdot \frac{1}{|N|}. \end{aligned} \quad (4.10)$$

Defining q_{\max}^* as the largest $q^*(p)$ for all $p \in P \setminus (\pi^* \cup \{p^*\})$, (4.10) yields:

$$\text{share}_{i^*}(P \setminus (\pi^* \cup \{p^*\})) \leq q_{\max}^* \cdot \frac{|P \setminus (\pi^* \cup \{p^*\})|}{|N|}. \quad (4.11)$$

With the aim of proving inequality (4.7), we want to show that

$$q_{\min} \cdot \sum_{p \in \pi^* \setminus P} \gamma_{i^*}(p) > q_{\max}^* \cdot \frac{|P \setminus (\pi^* \cup \{p^*\})|}{|N|}. \quad (4.12)$$

Note that proving that this inequality holds, would in turn prove (4.7) thanks to (4.9) and (4.11). We divide the proof of (4.12) into two claims.

Claim 4.3.17. $q_{\min} \geq q_{\max}^*$.

Proof: Consider any project $p' \in P \setminus (\pi^* \cup \{p^*\})$. It must be the case that p' was at least $1/q^*(p)$ -affordable in round $1, \dots, k^*$, for all $p \in \pi^*$, as all voters in N could have fully contributed to it based on how we defined k^* .

As no $p' \in P \setminus (\pi^* \cup \{p^*\})$ was selected by $\text{MES}[\text{share}]$, we know that all projects that have been selected must have been at least as affordable, *i.e.*, for all $p \in \pi^*$ and $p' \in P \setminus (\pi^* \cup \{p^*\})$ we have:

$$\alpha(p) \leq \frac{1}{q^*(p')} \iff q(p) \geq q^*(p') \iff q_{\min} \geq q_{\max}^*.$$

This concludes the proof of our first claim. \diamond

Claim 4.3.18. $\sum_{p \in \pi^* \setminus P} \gamma_{i^*}(p) > \frac{|P \setminus (\pi^* \cup \{p^*\})|}{|N|}$.

Proof: From the choice of k^* , the load of agent i^* at round k^* must satisfy:

$$\ell_{i^*}(\pi^*) + \frac{1}{|N|} > \frac{b}{n}.$$

On the other hand, since N is a P -cohesive group, we know that:

$$\frac{|P|}{|N|} = \frac{|P \setminus \{p^*\}|}{|N|} + \frac{1}{|N|} \leq \frac{b}{n}.$$

Linking these two facts together, we get:

$$\ell_{i^*}(\pi^*) > \frac{|P \setminus \{p^*\}|}{|N|}.$$

By the definition of the load, we thus have:

$$\ell_{i^*}(\pi^*) = \sum_{p_j \in \pi^*} \gamma_{i^*}(p_j) > \frac{|P \setminus \{p^*\}|}{|N|}.$$

This is equivalent to:

$$\sum_{p_j \in P \cap \pi^*} \gamma_{i^*}(p_j) + \sum_{p_j \in P \setminus \pi^*} \gamma_{i^*}(p_j) > \frac{|P \cap \pi^*|}{|N|} + \frac{|P \setminus (\pi^* \cup \{p^*\})|}{|N|} \quad (4.13)$$

Now, we observe that every voter in N contributed in full for every project in π^* . It follows that the contribution of every voter in N for a project $p_j \in P \cap \pi^*$ is smaller or equal the contribution needed if the voters in N would fund the project by themselves. In other words for all $p \in P \cap \pi^*$ we have:

$$\gamma_{i^*}(p) \leq \frac{1}{|N|}.$$

It follows then that:

$$\sum_{p_j \in P \cap \pi^*} \gamma_{i^*}(p_j) \leq \frac{|P \cap \pi^*|}{|N|}.$$

For (4.13) to be satisfied, we must have that:

$$\sum_{p_j \in \pi^* \setminus P} \gamma_{i^*}(p_j) > \frac{|P \setminus (\pi^* \cup \{p^*\})|}{|N|}.$$

This concludes the proof of our second claim. \diamond

Putting together these two claims shows that inequality (4.12) is satisfied, which in turn shows that (4.7) also is. Since P , N and p^* were chosen arbitrarily, this shows that $\text{MES}[\text{share}]$ satisfied Local-EJS in the unit-cost setting. \square

4.4 Relationships between Criteria

As the reader may have noticed already, there is a significant inflation of the number of fairness properties introduced for PB. This chapter makes no exception in that regard. In order to clarify the criteria we have introduced, we now analyse their relationships. We start with links within the space of criteria we introduced earlier, then compare them to the notion of priceability and other representation-based criteria, and finally discuss efficiency requirements.

4.4.1 Share-Based Fairness Criteria

The following theorem establishes the relations between share-based fairness concepts. These relations are visualised in Figure 4.4.1.

Theorem 4.4.1. *Given an instance I and a profile \mathbf{A} , for every budget allocation $\pi \in \text{FEAS}(I)$ the following statements hold:*

- (i) *If π satisfies FS, then it also satisfies FS-1, Local-FS, and Strong-EJS;*
- (ii) *If π satisfies FS-1, then it also satisfies EJS-1;*

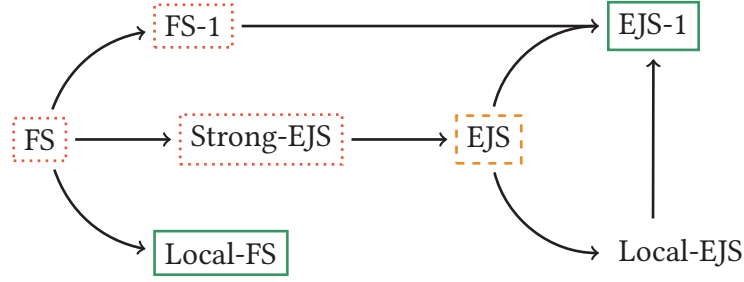


Figure 4.4.1: Taxonomy of criteria introduced in this chapter. An arrow from one criterion to another indicates that any budget allocation satisfying the former also satisfies the latter. $\text{MES}[share]$ satisfies the criteria boxed in green solid lines. For the criterion boxed in orange dashed lines, no efficient algorithms computing them exist (unless $P = NP$). Criteria boxed in red dotted lines are not always satisfiable. The status of Local-EJS is unknown.

(iii) If π satisfies Strong-EJS, then it also satisfies EJS;

(iv) If π satisfies EJS, then it also satisfies Local-EJS;

(v) If π satisfies Local-EJS, then it also satisfies EJS-1.

This list of implications is exhaustive when closed under transitivity.

PROOF. We prove the different claims consecutively.

(i) It is easy to verify from the definitions that every budget allocation satisfying FS also satisfy FS-1 and Local-FS. So let us show that FS also implies Strong-EJS. Let $i \in \mathcal{N}$ be an arbitrary agent. We distinguish two cases.

First, assume $share_i(A_i) < b/n$. For FS to be satisfied, we must have $share_i(\pi) \geq share_i(A_i)$. This entails that $A_i \subseteq \pi$. Hence, the conditions for Strong-EJS are trivially satisfied for agent i .

Second, assume that $share(A_i, i) \geq b/n$ holds. Since π satisfies FS, we know that $share(\pi, i) \geq b/n$. Let $N \subseteq \mathcal{N}$ be a P -cohesive group, for some $P \subseteq \mathcal{P}$, such that $i \in N$. By definition of a cohesive group, we know that $c(P) \leq b/n \cdot |N|$. Hence, $share(P, i) \leq b/n$. Overall, we have $share_i(\pi) \geq b/n \geq share_i(P)$ and thus π satisfies Strong-EJS.

(ii) Let π be a budget allocation that satisfies FS-1. First, consider an agent $i \in \mathcal{N}$ such that $share_i(A_i) < b/n$. For FS-1 to be satisfied, there must be a project $p \in \mathcal{P}$ such that $share_i(\pi \cup \{p\}) \geq share_i(A_i)$. This entails that $|A_i \setminus \pi| \leq 1$ should be the case. Hence, the conditions for EJS-1 are trivially satisfied for agent i .

Consider now an agent $i \in \mathcal{N}$ such that $share_i(A_i) \geq b/n$. Since π satisfies FS-1, we know that there must be a project $p \in \mathcal{P}$ such that $share_i(\pi \cup \{p\}) \geq b/n$. Let $N \subseteq \mathcal{N}$ be a P -cohesive group, for some $P \subseteq \mathcal{P}$, such that $i \in N$. By definition

of a cohesive group, we know that $c(P) \leq b/n \cdot |N|$. Hence, $share_i(P) \leq b/n$. Overall, we get $share_i(\pi \cup \{p\}) \geq b/n \geq share_i(P)$ and π satisfies EJS-1.

(iii) and (iv) These two claims are directly derived from the definitions.

(v) The fact that every budget allocation satisfying Local-EJS also satisfies EJS-1 is a direct consequence of Proposition 4.3.13 that states that Local-EJS can be interpreted as a universal strengthening of EJS-1. \square

In the proof of Theorem 4.4.1 we only showed “positive” implications. We now proceed to prove the absence of further ones through a series of counterexamples.

We start by studying Local-FS, and show that FS-1 does not imply Local-FS.

Example 4.4.2 (FS-1 does not imply Local-FS). Consider the following instance and profile with three agents, four projects and a budget limit of $b = 6$.

	p_1	p_2	p_3	p_4
Cost	3	3	6	1
A_1	✓	✓	✗	✗
A_2	✗	✗	✓	✓
A_3	✗	✗	✓	✓
$b = 6$				

Here, $\pi = \{p_1, p_2\}$ satisfies FS-1 as agent 1 already receives (more than) their fair share, while 2 and 3 receive their fair share from $\pi \cup \{p_3\}$. However, no supporter of p_4 receives their fair share from $\pi \cup \{p_4\}$. Therefore, Local-FS is violated. \triangle

The other direction—Local-FS does not imply FS-1—follows from the fact that a Local-FS allocation always exists (as a consequence of Theorem 4.2.8) while FS-1 is not always satisfiable (Proposition 4.2.5). We can still present a counterexample when both a Local-FS *and* an FS-1 budget allocation exist.

Example 4.4.3 (Local-FS does not imply FS-1). Consider the following instance and profile with three projects, a budget limit of $b = 10$, and four agents.

	p_1	p_2	p_3
Cost	3	4	7
A_1	✗	✗	✓
A_2	✗	✗	✓
A_3	✓	✓	✗
A_4	✓	✓	✓
$b = 10$			

First, note that the $fairshare_1 = fairshare_2 = 7/3$ and $fairshare_3 = fairshare_4 = 5/2$. The budget allocation $\pi = \{p_1, p_2\}$ satisfies FS-1, since adding project p_3 gives each agent at least their fair share. The budget allocation $\pi' = \{p_3\}$ does not satisfy FS-1, as voter 3 does not achieve their fair share by adding either only p_1 or only p_2 . It does, however, satisfy Local-FS, since adding p_1 or p_2 leads to too large a share for voter 4. Concretely, we have:

$$share_4(\pi' \cup \{p_1\}) = 3/2 + 7/3 > 5/2 \quad \text{and} \quad share_4(\pi' \cup \{p_2\}) = 2 + 7/3 > 5/2.$$

Consequently, Local-FS does not imply FS-1, even if an FS-1 allocation exists. \triangle

Note, that in the above example, the budget allocation $\{p_1, p_2\}$ satisfies both Local-FS and FS-1. So we have not ruled out the possibility that the existence of an FS-1 allocation implies the existence of an allocation satisfying both FS-1 and Local-FS.

Next, we show Strong-EJS also does not imply Local-FS.

Example 4.4.4 (Strong-EJS does not imply Local-FS). Consider the instance and the profile presented below, with four projects, a budget limit of $b = 16$, and two agents.

	p_1	p_2	p_3	p_4
Cost	12	12	1	4
A_1	✓	✓	✓	✗
A_2	✓	✗	✗	✓
$b = 16$				

In this instance, the cohesive groups are: $\{1, 2\}$ which is $\{p_1\}$ -cohesive; $\{1\}$ which is $\{p_3\}$ -cohesive, and $\{2\}$ which is $\{p_4\}$ -cohesive. Overall, to satisfy Strong-EJS, a budget allocation should provide a share of at least 6 to agents 1 and 2. The budget allocation $\pi = \{p_1, p_4\}$ thus satisfies Strong-EJS (note that it is exhaustive). However, one can easily check that π does not satisfy the conditions of Local-FS as adding p_3 to π only provides a share of 7 to agent 1 while their fair share is 8. \triangle

Example 4.4.4 also shows that neither EJS, Local-EJS, nor EJS-1 imply Local-FS.

Finally, we show that Local-FS does not imply EJS-1.

Example 4.4.5 (Local-FS does not imply EJS-1). Consider the instance and profile presented below, with three projects, a budget limit of $b = 6$, and two agents.

	p_1	p_2	p_3
Cost	6	1	1
A_1	✓	✓	✓
A_2	✗	✓	✓
$b = 6$			

Here, the budget allocation $\pi = \{p_1\}$ satisfies Local-FS since agent 1 already receive more than their fair share and approves of both p_2 and p_3 . However, π does not satisfy EJS-1: $\{2\}$ is a $\{p_2, p_3\}$ -cohesive group but neither p_2 nor p_3 are selected. \triangle

Note that due to the implications shown in Theorem 4.4.1, Local-FS also does not imply any of Local-EJS, EJS, Strong-EJS, and FS.

We have now proved that there are no additional implications involving Local-FS than the link between FS and Local-FS.

We now consider EJS-1 and show that it does not imply Local-EJS.

Example 4.4.6 (EJS-1 does not imply Local-EJS). Consider the instance and profile presented below, with two agents, six projects and a budget limit of $b = 4$.

	p_1	p_2	p_3	p_4	p_5	p_6
Cost	1	1	1	1	1	1
A_1	✓	✓	✓	✓	✓	✗
A_2	✗	✗	✗	✓	✓	✓
$b = 4$						

We claim that $\pi = \{p_1, p_2, p_3, p_4\}$ satisfies EJS-1 but not Local-EJS. The share of agent 1 in π is 3.5 so every cohesive group containing them will satisfy the conditions for EJS-1 and Local-EJS. Consider now voter 2. Their share in π is $1/2$. Note that they are $\{p_5, p_6\}$ -cohesive and deserve thus a share of $3/2$. Since $\pi \cup \{p_6\}$ would provide them a share of $3/2$, π satisfies EJS-1. However, $\pi \cup \{p_5\}$ would only provide agent 2 a share of 1, showing that π fails Local-EJS. \triangle

Interestingly, in the example above the budget allocation $\{p_3, p_4, p_5, p_6\}$ provides both agents with their fair share, showing that EJS-1 does not imply any of FS, FS-1, Strong-EJS, and EJS, and that even when they can be satisfied.

We continue with FS-1. We know that it implies EJS-1, we now show that it does not imply Local-EJS, the criteria one step above EJS-1 in our hierarchy. From Theorem 4.4.1, this entails that FS-1 also does not imply EJS.

Example 4.4.7 (FS-1 does not imply Local-EJS). Consider the instance and profile presented below, with four projects, a budget limit of $b = 12$, and three agents.

	p_1	p_2	p_3	p_4
Cost	4	2	5	7
A_1	✗	✗	✗	✓
A_2	✓	✓	✓	✗
A_3	✓	✓	✓	✗
$b = 12$				

Here, the budget allocation $\pi = \{p_1, p_4\}$ satisfies FS-1 but fails Local-EJS. FS-1 is satisfied since agent 1 already receive their fair share, and adding p_3 would make agents 2 and 3 receive theirs. Local-EJS is violated since the $\{p_2, p_3\}$ -cohesive group $\{2, 3\}$ deserves a share of 3.5, but adding p_2 to π would not only provide agents 2 and 3 with a share of 3. \triangle

Since the budget allocation $\{p_1, p_2, p_3\}$ satisfies Strong-EJS in the example above, we also know that FS-1 does not imply Strong-EJS. FS-1 also trivially does not imply FS.

The only implications we are still missing are the “backward” arrows: Local-EJS not implying EJS, EJS not implying Strong-EJS, and Strong-EJS not implying FS. Note that the latter is immediately derived from Example 4.4.4 and Theorem 4.4.1. The other two implication are immediate given the definitions.

The picture is now complete and we have derived an exhaustive taxonomy of the share-based criteria. We now turn to the comparison with priceability.

4.4.2 Comparison with Priceability

Priceability is a fairness criterion requiring that the budget allocation could in principle be obtained through a market-based approach. It is similar in spirit to share-based criteria as it also measures the amount of money spent on each agent. However, priceability does not require the cost of a project to be equally distributed between its supporters. Instead it requires the existence of a *price system*, *i.e.*, a distribution of the costs of the selected projects to their supporters that satisfies the following:

- C1: Agents only contribute to project they approve of;
- C2: Only selected projects receive contributions;
- C3: No agent contributes more than the entitlement;
- C4: The selected projects are receiving sufficient contributions to be funded;
- C5: The supporters of a non-selected project p have $c(p)$ or less money left.

The formal statement for each of these conditions can be found in Definition 3.3.31. Remember that we denote a price system by $\langle \alpha, (\gamma_i)_{i \in \mathcal{N}} \rangle$ where $\alpha \in \mathbb{R}_{\geq 0}$ is the entitlement and $(\gamma_i)_{i \in \mathcal{N}}$ a collection of contribution functions, with $\gamma_i(p)$ being the contribution of agent $i \in \mathcal{N}$ to project $p \in \mathcal{P}$. We will moreover denote by α_i^* the leftover money of agent i , that is, $\alpha_i^* = \alpha - \sum_{p \in \mathcal{P}} \gamma_i(p)$.

Due to the similar motivation of share-based fairness concepts and priceability, it is interesting to understand the relationships between them. We start with the fair share and show that the intuitive connection between FS and priceability does not hold.

Proposition 4.4.8. *There exists an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile \mathbf{A} such that there is $\pi \in \text{FEAS}(I)$ that satisfies FS, but such that no FS budget allocation π is priceable.*

PROOF. Consider the following instance and profile with four projects, a budget limit of $b = 9$, and three agents.

	p_1	p_2	p_3	p_4
Cost	1	5	3	1
A_1	✓	✗	✗	✗
A_2	✗	✓	✗	✗
A_3	✗	✗	✓	✓
$b = 9$				

Here, the only budget allocation satisfying FS is $\pi = \{p_1, p_2, p_3\}$. Assume towards a contradiction that π is priceable with entitlement $\alpha \in \mathbb{R}_{\geq 0}$ and contribution functions γ_1, γ_2 and γ_3 . Since only agent 2 approves of p_2 , from conditions **C1** and **C4** of priceability, we must have $\gamma_2(p_2) = 5$. Condition **C3** then implies that $\alpha \geq 5$. For similar reasons we should have $\gamma_3(p_3) = 3$ and $\gamma_3(p_1) = \gamma_3(p_2) = 0$. Condition **C2** also imposes $\gamma_3(p_4) = 0$. Overall this means that $\alpha_3^* = \alpha - \gamma_3(p_3) \geq 2$. This is a violation of condition **C5** for agent 3 and project p_4 as $\alpha_3^* > c(p_4)$. \square

Interestingly, the expected connection between fair share and priceability does hold when agents approve of sufficiently many alternatives.

Proposition 4.4.9. *For every instance $I = \langle \mathcal{P}, c, b \rangle$ and profile \mathbf{A} such that for every agent $i \in \mathcal{N}$ we have $\text{fairshare}_i = b/n$, it is the case that every budget allocation $\pi \in \text{FEAS}(I)$ that satisfies FS is also priceable.*

PROOF. Consider an instance $I = \langle \mathcal{P}, c, b \rangle$ and a suitable profile \mathbf{A} . Assume that there is a budget allocation that satisfies FS in I , and let $\pi \in \text{FEAS}(I)$ satisfy FS. We claim that π is priceable for the entitlement $\alpha = b/n$ and the contribution functions $(\gamma_i)_{i \in \mathcal{N}}$ defined for every agent $i \in \mathcal{N}$ and project $p \in \mathcal{P}$ as:

$$\gamma_i(p) = \begin{cases} \text{share}_i(\{p\}) & \text{if } p \in A_i \cap \pi, \\ 0 & \text{otherwise.} \end{cases}$$

First note that conditions **C1** and **C2** of priceability are trivially satisfied here.

Now, since π satisfies FS, we know that $\text{share}_i(\pi) \geq \text{fairshare}_i = b/n$ holds for every agent $i \in \mathcal{N}$. Since $\sum_{i \in \mathcal{N}} \text{share}_i(\pi) = c(\pi)$ and π is feasible, we must have $\text{share}_i(\pi) = b/n$ for all $i \in \mathcal{N}$. Overall, we have $\sum_{p \in \mathcal{P}} \gamma_i(p) = \text{share}(\pi, i) = b/n \leq \alpha$, so condition **C3** also is satisfied.

In addition, we have $\sum_{i \in \mathcal{N}} \gamma_i(p) = \sum_{i \in \mathcal{N}} \text{share}_i(\{p\}) = c(p)$. Condition **C4** is thus immediately satisfied.

Finally, as we have for every agent $\sum_{p \in \mathcal{P}} \gamma_i(p) = \text{share}_i(\pi) = b/n = \alpha$, no agent has any unspent allowance and condition **C5** is vacuously satisfied. \square

Next, we consider the relation between the weaker share-based notions and priceability. We actually show the absence of a relation, and that, even if we assume that agents approve of enough projects.

Proposition 4.4.10. *Local-FS, FS-1, and Strong-EJS do not imply priceability, even for instances and profiles for which $\text{fairshare}_i = b/n$ for every agent $i \in \mathcal{N}$. Reciprocally, priceability does not imply Local-FS or EJS-1, even if $\text{fairshare}_i = b/n$ for every agent $i \in \mathcal{N}$ and the agents have an entitlement of at least b/n .*

PROOF. Consider an instance I and a profile \mathbf{A} with two projects with $c(p_1) = 3$ and $c(p_2) = 2$, a budget limit $b = 3$, and two agents such that $A_1 = \{p_1\}$ and $A_2 = \{p_2\}$. Then the budget allocation $\{p_1\}$ satisfies FS-1 and Local-FS as we have $\text{share}_i(\{p_1, p_2\}) > \text{fairshare}_i$ for $i \in \{1, 2\}$. Moreover, Strong-EJS is trivially satisfied, as there are no cohesive groups. On the other hand, for $\{p_1\}$ to be priceable, each agent must receive an entitlement of 3. In this case, the fact that p_2 is not selected is a contradiction to condition **C5** as $p_2 \in A_2$ and agent 2 has more than $c(p_2)$ unspent entitlement.

Now, consider the following instance and profile with four projects, a budget limit $b = 20$, and two agents.

	p_1	p_2	p_3	p_4
Cost	8	8	5	5
A_1	✓	✓	✗	✗
A_2	✗	✓	✓	✓
$b = 20$				

Here, the budget allocation $\pi = \{p_1, p_2\}$ is priceable for the following contribution functions with an entitlement of $\alpha = 10$ per agent:

$$\begin{array}{cccc} \gamma_1(p_1) = 8 & \gamma_1(p_2) = 0 & \gamma_1(p_3) = 0 & \gamma_1(p_4) = 0, \\ \gamma_2(p_2) = 8 & \gamma_2(p_2) = 0 & \gamma_2(p_3) = 0 & \gamma_2(p_4) = 0. \end{array}$$

However, π fails Local-FS as $\text{share}_2(\pi \cup \{p_3\}) = 9 < 10 = \text{fairshare}_2$. Moreover, $\{2\}$ is $\{p_3, p_4\}$ -cohesive but $\text{share}_2(\pi \cup \{p\}) = 9 < 10 = \text{share}_2(\{p_3, p_4\})$ for any $p \in \{p_3, p_4\}$. Hence, π also does not satisfy EJS-1. \square

However, Local-FS, EJS-1, and priceability are compatible in the sense that for every instance I and profile \mathbf{A} , there always exists a budget allocation that satisfies all

three, namely the output of $\text{MES}[share]$ on I and \mathbf{A} . This follows directly from Theorem 4.2.8, Proposition 4.3.11, and the fact that MES is priceable for any satisfaction function (see Proposition 3.3.32). It remains open whether FS-1, EJS, and Local-EJS are compatible with priceability in this sense.

4.4.3 Representation-Based Fairness Criteria

Priceability is known to have links with representation-based requirements, and in particular, to imply $\text{PJR}[sat^{cost}]$ in the unit-cost setting (Lackner and Skowron, 2023). Under the assumption that $fairshare_i = b/n$ for all agents $i \in \mathcal{N}$, Proposition 4.4.9 entails that $\text{PJR}[sat^{cost}]$ and FS are compatible in the unit-cost setting.

It is unclear whether this link can be extended to the PB setting since the link between priceability and $\text{PJR-X}[sat]$ requires stronger notion of priceability (see Theorem 3.3.35), that are not offered by Proposition 4.4.9.

In the following we investigate other such relationships between representation-based and share-based fairness criteria.

The first question we investigate is whether the aforementioned link between FS and PJR can be extended to EJR. We answer in the negative, showing that FS and EJR are incompatible, already in the unit-cost case.

Proposition 4.4.11. *Both FS and FS-1 are incompatible with $\text{EJR}[sat^{card}]$, even in the unit-cost setting.*

PROOF. Consider the instance I and the profile \mathbf{A} as described next. There are fifteen projects, all of cost 1, a budget limit of $b = 8$, and 32 agents. To simplify the exposition, let $\mathcal{P} = \{p_1, \dots, p_8\} \cup P$ where P is defined as $P = \{p'_1, \dots, p'_7\}$. The approval ballots are as follows:

$$\begin{aligned} A_1 &= \dots = A_7 = \{p_1, p_2\} \cup P, \\ A_8 &= \dots = A_{14} = \{p_3, p_4\} \cup P, \\ A_{15} &= \dots = A_{21} = \{p_5, p_6\} \cup P, \\ A_{22} &= \dots = A_{28} = \{p_7, p_8\} \cup P, \\ A_{29} &= \{p_1, p_2\}, \\ A_{30} &= \{p_3, p_4\}, \\ A_{31} &= \{p_5, p_6\}, \\ A_{32} &= \{p_7, p_8\}. \end{aligned}$$

Consider a budget allocation $\pi \in \text{FEAS}(I)$ such that π satisfies $\text{EJR}[sat^{card}]$. Note that such a π always exists since $\text{MES}[sat^{card}]$ satisfies $\text{EJR}[sat^{card}]$ in the unit-cost setting. Since the group of voters $\{1, \dots, 28\}$ is P -cohesive group, there

is an agent $i \in \{1, \dots, 28\}$ for whom $|A_i \cap \pi| \geq 7$. For this to be true, it must contain at least 5 projects from P . Hence, at most 3 projects from $\{p_1, \dots, p_8\}$ can be in π . In particular, this entails that at least one voter from $29, \dots, 32$ has no approved project in π . Without loss of generality, assume that it is the case for voter 29, *i.e.*, that neither p_1 nor p_2 is in π . We show that this is incompatible with π satisfying FS-1. For voter 29, we have:

$$share_{29}(\pi) = 0 < fairshare_{29} = \frac{1}{4}.$$

Adding any project to π can increase 29's share by at most $\frac{1}{8}$. Hence π fails FS-1.

FS, and thus also FS-1, is however satisfiable in this instance: $\{p_1, \dots, p_8\}$ satisfies FS, but not $EJR[sat^{card}]$ (by our previous argument). \square

We now show that Local-FS is not related to any representation-based axiom. First, we prove that EJR does not imply Local-FS.

Proposition 4.4.12. *$EJR[sat^{card}]$ does not imply Local-FS, even in the unit-cost setting.*

PROOF. Consider the following instance and profile with eight projects, all of cost 1, a budget limit $b = 6$, and three agents.

	p_1	p_2	p_3	p_4	p_5	p_6	p_7	p_8
Cost	1	1	1	1	1	1	1	1
A_1	✓	✓	✓	✓	✓	✓	✗	✗
A_2	✓	✓	✓	✓	✗	✗	✓	✓
A_3	✓	✓	✓	✓	✗	✗	✓	✓
$b = 6$								

The budget allocation $\pi = \{p_1, p_2, p_3, p_4, p_5, p_6\}$ satisfies $EJR[sat^{card}]$. This holds, in particular, because the group of voter $\{1\}$ is $\{p_5, p_6\}$ -cohesive group, and the group $\{2, 3\}$ is $\{p_1, p_2, p_3, p_4\}$ -cohesive.

On the other hand, π does not satisfy Local-FS. First, note that $share_2(\pi) = share_3(\pi) = 4/3$. Now, consider project p_7 , for agent 2 we have:

$$share_2(\pi \cup \{p_7\}) = 4/3 + 1/2 = 11/6 < 2 = b/n = fairshare_2.$$

Hence, project p_7 is a witness that π fails Local-FS. \square

Then, we show that Local-FS does not even imply *justified representation* (JR), the fairness requirements defined exactly as EJR (or PJR), but only for P -cohesive groups such that $|P| = 1$. It is particularly weak and it trivially implied by both EJR and PJR.

Proposition 4.4.13. *Local-FS does not imply $\text{JR}[sat^{card}]$, even in the unit-cost setting.*

PROOF. Consider the following instance and profile with four projects, all of cost 1, a budget limit $b = 3$, and three agents.

	p_1	p_2	p_3	p_4
Cost	1	1	1	1
A_1	✓	✗	✗	✗
A_2	✓	✓	✓	✓
A_3	✗	✓	✓	✓
$b = 3$				

The budget allocation $\pi = \{p_2, p_3, p_4\}$ trivially satisfies Local-FS since only one project has not been selected. However, π fails $\text{JR}[sat^{card}]$ since voter 1 forms a $\{p_1\}$ -cohesive group and thus deserves at least one approved project in π . \square

This concludes our study of the relationship between representation-based and share-based guarantees. An important open question is whether there exists a rule that satisfies both Local-FS and $\text{EJR}[sat]$, or $\text{EJR-1}[sat]$, for some sat .

4.4.4 Share and Efficiency Requirements

It is well-known that enforcing fairness criteria often comes at a cost with respect to social welfare (Theorem 3.3.18; Lackner and Skowron, 2020; Elkind, Faliszewski, Igarashi, Manurangsi, Schmidt-Kraepelin and Suksompong, 2022). Exploring the connection between share-based concepts and social welfare is thus particularly interesting. We do not delve into that line of research here but present an interesting connection between EJS and an efficiency requirement known as project-wise unanimity.

Definition 4.4.14 (Project-Wise Unanimity). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile \mathbf{A} , a budget allocation $\pi \in \text{FEAS}(I)$ satisfies project-wise unanimity if for every $p, p' \in \mathcal{P}$ such that $\text{app}(p, \mathbf{A}) < \text{app}(p', \mathbf{A}) = n$, it holds that $p \in \pi$ implies $p' \in \pi$.*

Proposition 4.4.15. *There exists an instances I and a profile \mathbf{A} such that there is $\pi \in \text{FEAS}(I)$ that satisfies FS, but such that no budget allocation $\pi' \in \text{FEAS}(I)$ that satisfies either FS or EJS also satisfies project-wise unanimity, even in the unit-cost setting.*

PROOF. Consider an instance I with four projects p_1, \dots, p_4 , all of cost 1, and a budget limit $b = 3$. The profile \mathbf{A} is composed of nine ballots such that six agents approve of p_1, p_2 and p_3 , while the other three agents approve of p_3 and

p_4 . Overall, $\{1, \dots, 6\}$ is $\{p_1, p_2\}$ -cohesive for a share requirement of $1/3$; and $\{7, 8, 9\}$ is $\{p_4\}$ cohesive for a share requirement of $1/3$. Note that $1/3$ is also the fair share of any agent.

In this instance the only budget allocation satisfying EJS, and also FS, is $\pi = \{p_1, p_2, p_4\}$. It is however easy to see that p_3 is unanimously approved by the agents but not in π , showing that π violates project-wise unanimity. \square

This result indicates that enforcing share-based fairness criteria probably comes at a high cost in terms of efficiency. This connection definitely deserves further attention.

4.5 Approaching Fair Share in Practice

As we saw in Section 4.2, there exist PB instances for which it is impossible to give every agent their fair share. In this section we conduct an experimental study aimed at understanding how serious a problem this is. Our study is twofold. We first investigate how close to fair share we can get, for different notion of closeness. In a second experiment, we quantify how close to this optimal value certain known PB rules get.

Let us first describe the setting. For these experiments we use data from Pabulib (Stolicki, Szufa and Talmon, 2020), an online collection of real-world PB datasets. To be more precise, we used all instances from Pabulib with up to 65 projects (available in October 2022), except for trivial instances, where either no project or the set of all projects are affordable. Three instances have been additionally omitted for the first experiment due to very high compute time. A total of 353 PB instances are covered by our analysis. Our experiments are implemented on a Debian machine with 16 cores and 16GB RAM, using Gurobi 9.5.1 for solving the mixed integer linear programs.

4.5.1 Optimal Distance to Fair Share

Our first experiment explores how close to FS we can get in practice. We first introduce two ways of measuring how close to FS a given budget allocation is. These measures represent the foundation of our analysis. For each of the considered instances and profiles, we will compute the optimal value of these two measures.

The first measure of distance to fair share we consider is the *average capped fair share ratio*: For every agent i with approval ballot A_i we divide their share in the budget allocation π by their fair share, capped at 1 in case they get more than their fair share, and take the average of this ratio over all agents:

$$\frac{1}{n} \cdot \sum_{i \in \mathcal{N}} \min \left(\frac{\text{share}_i(\pi)}{\text{fairshare}_i}, 1 \right).$$

This measure is a maximisation objective: we seek budget allocations that are getting as close as possible to 1, its maximum. Note that when the average capped fair share

$$\begin{array}{ll}
\mathbf{maximise} & \sum_{i=1}^n s_i \\
\mathbf{subject\ to:} & \\
s_i \in \mathbb{R} & \text{for all } i \in \mathcal{N} \\
x_p \in \{0, 1\} & \text{for all } p \in \mathcal{P} \\
s_i \leq \frac{\sum_{p \in A_i} x_p \cdot \text{share}(\{p\})}{\text{fairshare}_i} & \text{for all } i \in \mathcal{N} \\
s_i \leq 1 & \text{for all } i \in \mathcal{N} \\
\sum_{p \in \mathcal{P}} x_p \cdot c(p) \leq b &
\end{array}$$

Figure 4.5.1: A mixed integer linear program for maximizing the average capped fair share ratio for a given instance $I = \langle \mathcal{P}, c, b \rangle$ and profile \mathbf{A} . For any project $p \in \mathcal{P}$, variable x_p indicates whether p is selected ($x_p = 1$) or not ($x_p = 0$). For any agent $i \in \mathcal{N}$, variable s_i stores i 's fair share ratio.

ratio is 1, the budget allocation indeed satisfies FS. Figure 4.5.1 presents how to find the optimum value of this measure through a mixed integer linear program.

Our second measure is the *average L_1 distance to FS*, measuring, for every agent, the absolute value of the difference between their actual share in π and their fair share:

$$\frac{1}{n} \cdot \sum_{i \in \mathcal{N}} |\text{share}_i(\pi) - \text{fairshare}_i|.$$

This measure is a minimisation objective: we seek budget allocations that are getting as close as possible to 0, its minimum. Note that when the average L_1 distance to FS is 0, the budget allocation indeed satisfies FS. We present in Figure 4.5.2 a mixed integer linear program that we use to compute an optimal budget allocation.

When presenting our results with the L_1 distance to FS, we will actually plot the following measure:

$$1 - \frac{1}{n} \sum_{i \in \mathcal{N}} \frac{|\text{share}_i(\pi) - \text{fairshare}_i|}{\text{fairshare}_i}.$$

This allows us to obtain a normalised value and for which the higher the better, with 1 being the maximum value.³⁴ That way, the behaviour is similar to the average capped fair share ratio and the results are thus more easily comparable.

³⁴Note that the empty budget allocation provides an L_1 distance to FS of $\text{fairshare}(i)$ for all $i \in \mathcal{N}$. Normalising the L_1 distance with $\text{fairshare}(i)$, thus ensures that we display the optimal L_1 distance to FS achieved with respect to the worst case.

$$\begin{array}{ll}
\text{minimise} & \sum_{i=1}^n s_i \\
\text{subject to:} & \\
s_i \in \mathbb{R} & \text{for all } i \in \mathcal{N} \\
x_p \in \{0, 1\} & \text{for all } p \in \mathcal{P} \\
s_i \geq \sum_{p \in A_i} x_p \cdot \text{share}(\{p\}) - \text{fairshare}_i & \text{for all } i \in \mathcal{N} \\
s_i \geq \text{fairshare}_i - \sum_{p \in A_i} x_p \cdot \text{share}(\{p\}) & \text{for all } i \in \mathcal{N} \\
\sum_{p \in \mathcal{P}} x_p \cdot c(p) \leq b &
\end{array}$$

Figure 4.5.2: A mixed integer linear program for minimising the average L_1 distance to FS for a given instance $I = \langle \mathcal{P}, c, b \rangle$ and profile \mathbf{A} . For any project $p \in \mathcal{P}$, variable x_p indicates whether p is selected ($x_p = 1$) or not ($x_p = 0$). For any agent $i \in \mathcal{N}$, variable s_i stores i 's L_1 distance to FS.

Let us say few words of comparison between these two measures. Intuitively, the average capped fair share ratio penalises budget allocation that do not provide agent with sufficiently large share to reach their fair share. The average L_1 distance to FS does the same, but also penalises giving agents more share than their fair share.

To better understand what might cause an instance not to admit a good solution, we also considered different ways of preprocessing the instances by removing “problematic” projects. These preprocessing methods are described below.

- **Threshold Preprocessing:** any project that is not approved by at least $x\%$ of the voters are removed from the instance. We considered three threshold values: 1%, 5%, and 10%.
- **Cohesiveness Preprocessing:** any project such that its supporters do not control enough money to buy the project are removed. Formally, we remove any project p such that:

$$\frac{|\{i \in \mathcal{N} \mid p \in A_i\}|}{n} \cdot b < c(p).$$

This removes any project that is not involved in any cohesive group.

Threshold preprocessing removes under 10% of projects for a threshold of 1%, around 10–20% for a threshold of 5%, and around 20–30% for a threshold of 10%. Cohesiveness preprocessing removes between 30% (for the largest instances) and 70% of projects (for the smallest instances). The specific values are presented in Figure 4.5.3.

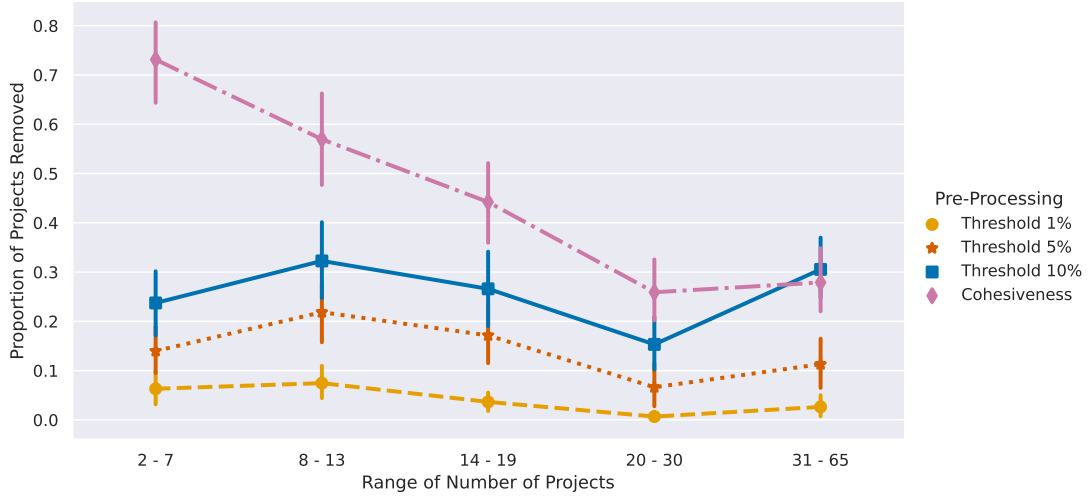


Figure 4.5.3: Proportion of projects removed during the preprocessing stage.

Note that we do not wish to advocate preprocessing as a method to make budget decisions in practice. Rather, we use it as a way of checking whether the failure to guarantee fair share is due to the specific structure of real-life PB instances and whether similar instances “nearby” might be significantly better behaved.

Let us now turn to the main findings of this experiment. Everything is presented in Figure 4.5.4. We draw the following conclusions.

- ▶ Without preprocessing, we can provide agents on average between 45% (for small instances) and 75% (for larger instances) of their fair share, albeit with a lot of variation.
- ▶ We can typically guarantee an L_1 distance to FS of 50% of the worst case.
- ▶ Preprocessing helps when using the cohesiveness condition, but not with the threshold condition.

Overall, our experimental findings suggest that guaranteeing fair share simply is very hard across a wide range of instances. This is witnessed by the low value of the average capped fair share ratio and the average L_1 distance to FS that we observed; together with the fact that preprocessing almost does not help. Note that we were able to satisfy FS for only one instance (that has 3 projects and 198 voters).

To corroborate this conclusion, we also investigate approximations of the average capped fair share ratio. Specifically, for a number of different given approximation ratios $\alpha \in (0, 1]$, we replaced the fair share by $\alpha \cdot \text{fairshare}_i$ in the definition, *i.e.*, for a given approximation ratio $\alpha \in (0, 1]$, the measure of interest is:

$$\frac{1}{n} \cdot \sum_{i \in \mathcal{N}} \min \left(\frac{\text{share}_i(\pi)}{\text{fairshare}_i \cdot \alpha}, 1 \right).$$

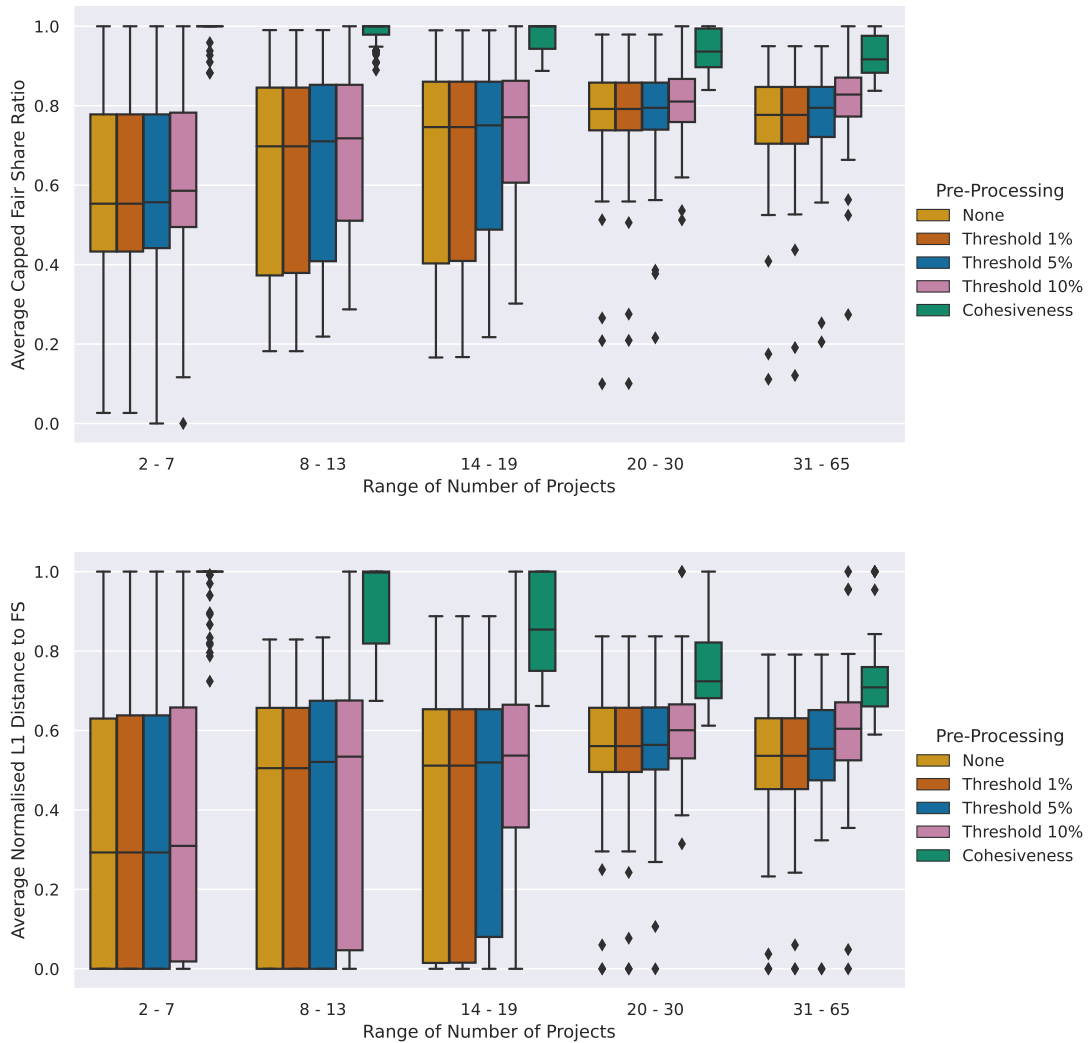


Figure 4.5.4: Average capped fair share ratio (top) and L_1 distance to FS (bottom) for Pabulib instances. Each range (for a number of projects) shown on the x -axis contains between 60 and 80 elections (instance and profile).




Figure 4.5.5: Average capped fair share ratio for different approximation ratios.

Results indicate that moving from $\alpha = 1$ to $\alpha = 0.2$, has a very small effect on the optimum value (around 10% better for $\alpha = 0.2$). Figure 4.5.5 presents the exact values. We also interpret this result as stating that FS is structurally hard to satisfy.

4.5.2 Distance to Fair Share of Common PB Rules

We now turn to our second experiment: how close to fair share are the outputs of known PB rules in practice. We will consider the following rules: $MES[share]$, $MES[sat^{cost}]$, $MES[sat^{card}]$, $SEQPHRAG$, and $GREEDCOST$. We only consider resolute versions of the rules, breaking ties in favour of the projects with the highest number of supporters (and then lexicographically if ties persist). All these rules have been formally introduced in Section 2.3.

For all instances and profiles, we compute the outcome returned by all of the rules, and assess how close to the optimal value they are in terms of both the average capped fair share ratio and average L_1 distance to FS. Results are presented in Figure 4.5.6.

The first striking observation is that $GREEDCOST$ is performing extremely well under the capped fair share ratio measure. This is particularly surprising given how oblivious to the structure of the profile it is. We postulate that this result is due to the high difference in the percentage of the budget used by the different rules: MES rules use around 40% of the budget on average, while $GREEDCOST$ and $SEQPHRAG$ use around 90% of the budget. See Figure 4.5.7 for more details. Since using more budget can only improve the average capped fair share ratio, this is the most likely explanation for the good performance of greedy approval compared to MES . It is thus hard to compare rules based on the average capped fair share ratio they achieve.

Interestingly, the average L_1 distance to FS does not suffer this drawback. Indeed,




Figure 4.5.6: Average capped fair share ratio (top) and average L_1 distance to FS (bottom) for different rules on Pabulib instances. Results are normalised by the optimum value achievable in each instance, giving a score between 0 and 1 where 1 is the best.

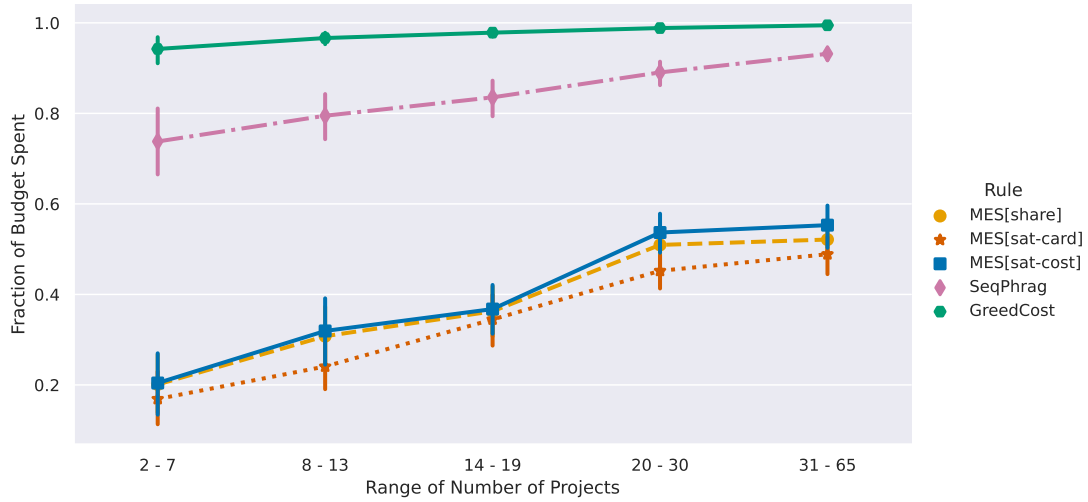


Figure 4.5.7: Fraction of the budget spent by the different rules.

since it also penalises rules that provide agents more than their fair share, spending more is not always better. Interpreting the results of Figure 4.5.6 in this light, we conclude that MES rules perform better than SEQPHRAG in terms of equality of resources. It may come as a surprise to the reader, but we also observe that $MES[sat^{cost}]$ performs slightly better than $MES[share]$. $MES[sat^{cost}]$ thus provides simultaneously good experimental results in terms of equality of resources, and strong representation guarantees (see theorems 3.3.17 and 3.3.21).

4.6 Summary

In this chapter, we have introduced a fairness theory for PB grounded in the concept of equality of resources. We have presented a large set of related fairness measures, all stemming from the *fair share* criterion, the idea that all agents deserve to be allocated the same amount of resources. It did not come as a surprise that fair share cannot always be satisfied. However, we managed to find related concepts that can always be satisfied, namely Local-FS and EJS-1. A full taxonomy can be found in Figure 4.4.1, in the middle section of this chapter. Interestingly both Local-FS and EJS-1 are satisfied by the same rule, $MES[share]$. Our experimental analysis also highlighted that a similar rule, namely $MES[sat^{cost}]$, performed particularly well in terms of approaching FS on real-life instances (even better than $MES[share]$). The latter can then be considered a good candidate for a rule providing strong guarantees both in terms of share-based and in terms of representation-based fairness (since it also satisfies EJR-1). The main take-home message of this chapter is that it is possible to study fairness in terms of equality of resources, and that doing so provides interesting insights on how to devise mechanisms for PB.

Chapter 5

Viewing Participatory Budgeting Rules through the Epistemic Lens

The ultimate goal of the research in social choice on PB, and of this thesis, is to devise appealing PB rules. However, the “appeal” of a given rule is always difficult to assess. In the typical case an appealing rule is one that produces outcomes that are good compromises between the preferences of the voters. This is for instance the approach we adopted in Chapter 4 where the most appealing rules were the ones providing strong fairness guarantees to the voters. This typical case is not the only one that deserves attention. In some other cases, the goal is not to reach good compromises but simply to selected the “best” alternatives. This chapter focuses on this case.

Assessing the merits of a rule based on its ability to retrieve the best outcome for society is the concern of *epistemic* social choice theory. This branch of the literature dates back to the Condorcet Jury Theorem (Condorcet, 1785; Dietrich and Spiekermann, 2019) and has received sustained attention since then (Elkind and Slinko, 2016; Pivato, 2019). It is based on the crucial assumption that there exists an outcome—the ground truth—that is *objectively* better than all the other outcomes. In the epistemic approach, rules are assessed on their ability to output the ground truth. The ground truth is not known by the agents and the aim of an elections is thus to retrieve it.

How can rules retrieve the ground truth if it is not a piece of information that is known by anyone? In the epistemic approach, it is assumed that the voters have some sense of what the ground truth might be, and thus that their ballots will reflect, at least partially, the ground truth. More formally, the ballots of the voters are assumed to be drawn from a *noise model*, a probability distribution over the set of all the admissible ballots, parametrised by a ground truth. The noise model can be viewed as the “black box” through which the voters determine their ballots given the information they have access to. Depending on the noise model under consideration the ballots will be more or less representative of the ground truth, but still connected to it. A rule

with high epistemic efficiency will thus be able to approach the ground truth from the ballots.

The aim of this chapter is to study PB rule via the epistemic approach we described above. The first step in this research agenda is to make sure that this approach is suitable for studying PB. In particular, one needs to motivate the idea that there would be a ground truth in PB. Let us present some arguments.

First, many PB projects come with a clear measures of efficiency that can only be measured *after* it has been realised: Will local residents actually use the compost bins? Will the number of mosquitoes go down? Will the new speed camera reduce the number of accidents?³⁵ In a pure determinist view, an entity—such as Laplace’s demon³⁶—who has perfect knowledge of the state of the world and of its rules would be able to objectively assess whether a project will be successful or not. In that view, the ground truth would be the set of the most successful projects.

In a more general view, PB can be seen as a selection process of costly alternatives. For some other such processes—that are, at least mathematically, equivalent to PB—the existence of a ground truth may be more obvious. Consider, for instance, the case of the [Eterna platform](#).³⁷ On this collaborative platform, users can submit different ways of folding a given protein. A subset of the proposed configurations is then synthesised in a laboratory to determine which are most stable. One can think of this as a PB process: the projects are the different protein foldings; their cost is the cost of synthesising them; the budget limit is the amount of money that is allocated to this process; and finally, the protein foldings submitted by a user constitute their approval ballot. Mathematically speaking, this is thus a well-defined PB process. Moreover, there is a clear ground truth here: a set of objectively most stable protein configurations. Interestingly, this is also the motivating example of the first epistemic analysis of multi-winner voting rules, though in a setting without costs ([Procaccia, Reddi and Shah, 2012](#)).

A last example, similar to the introductory example of this thesis, is that of a selection committee for research grant proposals. In such a committee, the members have to decide which of the grants should be funded. The grants typically have different costs, and there is a maximum amount of money that can be allocated. It is not hard to argue that the proposals have a “ground truth” probability of success that

³⁵All these examples are taken from projects that were brought to the vote in Toulouse 2019, see jeparticipe.metropole.toulouse.fr/processes/bp2019 for more details, and in particular the file “Catalogue des 30 idées soumises au vote”. The data is also hosted on the website pabulib.org ([Stolicki, Szufa and Talmon, 2020](#)), though the description of the projects is not available there.

³⁶Laplace’s demon is an entity that Laplace introduces (though does not call it a demon) to explain his view on determinism. He describes it as “An intellect which at a certain moment would know all forces that set nature in motion, and all positions of all items of which nature is composed, if this intellect were also vast enough to submit these data to analysis, it would embrace in a single formula the movements of the greatest bodies of the universe and those of the tiniest atom; for such an intellect nothing would be uncertain and the future just like the past would be present before its eyes” (see wikipedia.org/wiki/Laplace%27s_demon).

³⁷See eternagame.org and wikipedia.org/wiki/EteRNA for more information.

can only be observed *a posteriori*. At selection time, the committee members observe noisy signals regarding this ground truth through the reports of reviewers, and they need to make a decision based on this information.

Overall, the epistemic approach can be applied to many PB models where the existence of a ground truth is clear.

Until now, we have been rather evasive as to what it means exactly for a rule to retrieve the ground truth. Remember that the underlying machinery we consider here is that ballots are randomly generated given the ground truth through a noise model. In that sense, they represent noisy estimates of the ground truth. If we have a well-defined noise model, we, in principle, are able to design a voting rule that maximises the likelihood of returning the ground truth. The concept of *maximum likelihood estimators* (MLE) is well-studied in statistics and was designed exactly to answer such questions. Of course, in practice we do not have access to the exact noise model. Still, if a PB rule turns out to be the MLE of a natural noise model, then we can interpret this as an argument for using that rule. Similarly, if we can prove that for a given rule there does not exist *any* noise model that would make that rule an MLE, then we should interpret this as an argument against using that rule. Overall, the research question we try to answer in this chapter is the following:

Which PB rules can be interpreted in an epistemic fashion, that is, are maximum likelihood estimators for some noise models?

Answering this question will offer us another way to compare PB rules, based on their epistemic merits.

To answer this question, we will first introduce formally the epistemic approach for the standard model of PB (Section 5.1). We will then apply the approach to different sets of rules. We will start with the rules designed to provide proportional representation (Section 5.2). Next, we will consider what we call *monotonic argmax rules*, a class of rules that notably include many welfare-maximising rules (Section 5.3). We will focus on two specific sets of rules, related either to the Nash social welfare (Section 5.3.1) or to the utilitarian social welfare (Section 5.3.2). A summary of our findings will then be presented (Section 5.4).

Additional Related Work. The epistemic approach has been first applied to the standard ordinal voting model (Condorcet, 1785; Young, 1995; Conitzer and Sandholm, 2005; Caragiannis, Procaccia and Shah, 2014). Later on, other social choice scenarios have been investigated through the epistemic lens, notably multi-winner elections (Procaccia, Reddi and Shah, 2012; Caragiannis, Procaccia and Shah, 2013), and judgment aggregation (Bovens and Rabinowicz, 2006; Bozbay, Dietrich and Peters, 2014; Terzopoulou and Endriss, 2019). To the best of our knowledge, the only epistemic study of PB is the one section dedicated to the topic by Goel, Krishnaswamy, Sakshuwong and Aitamurto (2019), though in the context of divisible PB. Interestingly, they show that *knapsack voting* (a PB rule that resembles GREEDCOST in the divisible

setting) is an MLE, while we will see that GREEDCOST cannot be interpreted as an MLE in the context of *indivisible* PB (see Corollary 5.3.11 on page 134).

5.1 The Truth-Tracking Perspective

In this section, we present the truth-tracking perspective for PB. We will consider here the standard model of PB (Chapter 2) in which, when facing an instance $I = \langle \mathcal{P}, c, b \rangle$, agents submit approval ballots, that form altogether a profile \mathbf{A} . We will consider the PB rules in their irresolute versions, *i.e.*, when they return non-empty sets of feasible budget allocations: $F(I, \mathbf{A}) \subseteq \text{FEAS}(I)$.

In the truth-tracking perspective, there exists an objectively best feasible budget allocation for every instance that is the outcome that every reasonable rule should select. Such a budget allocation is called the *ground truth* and is denoted by π^* . The ground truth is not known, neither by the agents, nor by the decision maker. We will thus assess the quality of PB rules based on their ability to retrieve the ground truth given noisy votes.

Formally, a *noise model* \mathcal{M} is a generative model that produces random approval ballots for a given instance and ground truth. We represent it as a probability distribution over all approval ballots. For a given instance $I = \langle \mathcal{P}, c, b \rangle$, ground truth $\pi^* \in \text{FEAS}(I)$, and approval ballot $A \subseteq \mathcal{P}$, we denote by $\mathbb{P}_{\mathcal{M}}(A \mid \pi^*, I)$ the probability for the noise model \mathcal{M} to generate ballot A given I and π^* . Profiles are generated from \mathcal{M} by drawing the ballots identically and independently from \mathcal{M} . Specific examples of noise models will be presented later on.

Suppose the noise model \mathcal{M} indicates how the voters form their preferences. Then, a good rule should select the outcome that most likely would have generated the observed profile if it were the ground truth plugged into \mathcal{M} . This is the *maximum likelihood estimator* (MLE) of \mathcal{M} .

Definition 5.1.1 (Maximum Likelihood Estimators). *For a noise model \mathcal{M} , the likelihood of a profile \mathbf{A} for instance I and a budget allocation $\pi \in \text{FEAS}(I)$ is defined as:*

$$L_{\mathcal{M}}(\mathbf{A}, \pi, I) = \prod_{A \in \mathbf{A}} \mathbb{P}_{\mathcal{M}}(A \mid \pi, I).$$

A PB rule F is said to be the maximum likelihood estimator (MLE) of \mathcal{M} , if for every instance I and every profile \mathbf{A} we have:

$$F(I, \mathbf{A}) = \arg \max_{\pi^* \in \text{FEAS}(I)} L_{\mathcal{M}}(\mathbf{A}, \pi^*, I).$$

In the context of the standard model of voting theory, [Conitzer and Sandholm \(2005\)](#) identified a necessary condition for a voting rule to be interpretable as an MLE: it should satisfy what we are going to call *weak reinforcement*. This result straightforwardly carries over to the PB setting.

Definition 5.1.2 (Weak Reinforcement). *A PB rule F is said to satisfy weak reinforcement if and only if, for every instance I and every two profiles \mathbf{A} and \mathbf{A}' , we have:*

$$F(I, \mathbf{A}) = F(I, \mathbf{A}') \implies F(I, \mathbf{A} + \mathbf{A}') = F(I, \mathbf{A}) = F(I, \mathbf{A}').$$

Remember that $+$ is the concatenation operation between profiles.

Proposition 5.1.3 (Conitzer and Sandholm, 2005). *If a PB rule F does not satisfy weak reinforcement, then there exists no noise model \mathcal{M} for which F is the MLE.*

Note that this result applies for any set of possible ground truths, so also if we assume the ground truth to be exhaustive.

In this chapter, we will consider a large set of PB rules and investigate whether they can be interpreted as MLE.

5.2 Proportional PB Rules

We have already discussed at length the idea of fairness and proportionality in PB (see chapters 3 and 4). As we have seen, most of the prominent rules in the literature are ones that relate to the idea of proportionality, *i.e.*, rules that treat groups of agents fairly. We will thus start our analysis with these rules. We will study the rule SEQPHRAG and MES rules, and show that they cannot be interpreted as MLEs.

We start with the sequential Phragmén rule, SEQPHRAG. Its definition was provided in Definition 2.3.2 of Chapter 2.

Proposition 5.2.1. *There exists no noise model \mathcal{M} such that SEQPHRAG is the MLE of \mathcal{M} , not even on unit-cost instances with the additional assumption that the ground truth is exhaustive.*

PROOF. Consider an instance I with four projects denoted by p_1, p_2, p_3 and p_4 , all of cost 1, and budget limit $b = 3$. We consider two profiles, \mathbf{A}^1 and \mathbf{A}^2 , as presented below, where \mathbf{A}^1 is on the left and \mathbf{A}^2 is on the right.

	p_1	p_2	p_3	p_4
Cost	1	1	1	1
A_1^1	✓	✗	✗	✗
A_2^1	✓	✗	✓	✓
A_3^1	✗	✓	✓	✓
A_4^1	✗	✓	✓	✓
A_5^1	✗	✓	✓	✓
$b = 3$				

	p_1	p_2	p_3	p_4
Cost	1	1	1	1
A_1^2	✗	✓	✗	✗
A_2^2	✓	✗	✓	✗
A_3^2	✓	✗	✓	✗
A_4^2	✓	✗	✓	✓
$b = 3$				

We claim that on both profile \mathbf{A}^1 and profile \mathbf{A}^2 , SEQPHRAG outputs a single budget allocation $\pi = \{p_1, p_3, p_4\}$. Remember that, in the unit-cost setting, SEQPHRAG is exhaustive. The budget allocation π is thus exhaustive.

For \mathbf{A}^1 , after $1/2$ units of money have been distributed both p_3 and p_4 are bought at a price of $1/4$. Once an additional $1/4$ of money has been injected, project p_1 is bought, making π the unique budget allocation returned by SEQPHRAG.

For \mathbf{A}^2 , first p_1 is bought at a price of $1/4$, then p_3 at a price of $1/3$. Finally, after $1/3$ additional units of money are distributed, project p_4 is bought (at that time, the supporters of project p_2 have collected $1/4 + 2/3 < 1$ money). The final outcome is thus indeed $\{\pi\}$.

Now, consider the joint profile $\mathbf{A}^3 = \mathbf{A}^1 + \mathbf{A}^2$, and let us detail the run of SEQPHRAG on I and \mathbf{A}^3 . The first project to be bought is p_3 at price $1/7$. Then, once an extra $5/42$ amount of money has been distributed, project p_1 is bought. At that time, the supporters of p_4 who do not approve of p_2 have no money since they approve of p_1 . On the other hand, the only supporter of p_2 that does not approve of p_4 has a strictly positive amount of money. Project p_2 will then be the last project selected (after another $2/21$ money has been injected). Overall the outcome would be $\{\{p_1, p_2, p_3\}\} \neq \{\pi\}$.

We have thus proven that Sequential Phragmén fails weak reinforcement. Proposition 5.1.3 then concludes the proof. Note that all budget allocations we considered are exhaustive, the result thus also applies if we only focus on exhaustive ground truth. \square

Moving on to MES rules, we show that $\text{MES}[sat]$ can be never be interpreted as an MLE, regardless of the satisfaction function sat considered. If needed, see Definition 2.3.3 for $\text{MES}[sat]$.

Proposition 5.2.2. *For any given satisfaction function sat , there exists no noise model \mathcal{M} such that $\text{MES}[sat]$ is the MLE of \mathcal{M} , not even on unit-cost instances.*

PROOF. Consider an instance I with two projects denoted by p_1 and p_2 , both of cost 1, and a budget limit $b = 2$. Let sat be an arbitrary satisfaction function.

Consider the two profiles $\mathbf{A}^1 = (\{p_1\}, \{p_2\})$ and $\mathbf{A}^2 = (\{p_1, p_2\}, \{p_1, p_2\})$. We claim that on both of these profiles, $\text{MES}[sat]$ would output a single budget allocation, $\pi = \{p_1, p_2\}$. Indeed, on \mathbf{A}^1 both agents receive 1 unit of money and can both afford the single project they approve of. The outcome is thus trivially $\{\pi\}$. On \mathbf{A}^2 both agents approve of all the projects and can afford them. The outcome is thus also $\{\pi\}$.

Now, for the joint profile $\mathbf{A}^3 = \mathbf{A}^1 + \mathbf{A}^2 = (\{p_1\}, \{p_2\}, \{p_1, p_2\}, \{p_1, p_2\})$, we claim that the outcome of $\text{MES}[sat]$ is such that:

$$\text{MES}[sat](I, \mathbf{A}^3) \in \{\{\{p_1\}\}, \{\{p_2\}\}, \{\{p_1\}, \{p_2\}\}\}.$$

Importantly, the above states that $\pi \notin \text{MES}[\text{sat}](I, \mathbf{A}^3)$. Let us go through the run of $\text{MES}[\text{sat}]$ on I and \mathbf{A}^3 . Here, the initial budget is $1/2$ for each agent. Thus, the supporters of either p_1 and p_2 have collectively $3/2$ units of money. Since for both p_1 and p_2 , their supporters have strictly positive satisfaction for them (by definition of a satisfaction function), either p_1 or p_2 will be selected in the first round of $\text{MES}[\text{sat}]$. Denote by p^* the selected project, and by p the other project. To buy p^* , all its supporters paid $1/3$. The supporters of p are thus now left with $1/2 + 2 \cdot (1/2 - 1/3) = 1/2 + 1/3 < 1$ units of money, not enough to afford p . There is thus no way for $\text{MES}[\text{sat}]$ to return $\{\{p_1, p_2\}\}$ on \mathbf{A}^3 . Which of the three possible outcomes will be selected in the end depends on the satisfaction function sat .

$\text{MES}[\text{sat}]$ thus fails weak reinforcement and Proposition 5.1.3 concludes. \square

Interestingly, for both of the proofs we presented above, the outcomes on the individual profiles and the ones on the joint profiles never intersect. This implies that even resolute versions of the rules (obtained by introducing some form of tie-breaking) would fail weak reinforcement, *i.e.*, would not be interpretable as MLEs.

5.3 Monotonic Argmax Rules

As we have seen, our first obstacle to finding rules that are MLEs, is that none of the proportional rules we considered satisfy weak reinforcement. Aiming for more positive results, we will now follow a different approach: instead of checking whether known PB rules do satisfy weak reinforcement, we will start from rules we know satisfy it, and investigate their epistemic abilities. For the remainder of this chapter, we will thus focus on *monotonic argmax rules*, a large class of rules, all of which satisfy weak reinforcement.

We start by defining argmax rules and then define what a *monotonic* one is.

Definition 5.3.1 (Monotonic Argmax Rules). *A PB rule F is called an argmax rule if there exists a function f , taking as input an instance I , a profile \mathbf{A} , and a budget allocation π , and returning a number $f(I, \mathbf{A}, \pi) \in \mathbb{R}$, such that for all instances I and all profiles \mathbf{A} , we have:*

$$F(I, \mathbf{A}) = \arg \max_{\pi \in \text{FEAS}(I)} f(I, \mathbf{A}, \pi).$$

An argmax rule defined via the function f is called monotonic if for every two profiles \mathbf{A} and \mathbf{A}' and every two budget allocations π and π' , the following two conditions hold:

- (i) $\left. \begin{array}{l} f(I, \mathbf{A}, \pi) < f(I, \mathbf{A}, \pi') \\ f(I, \mathbf{A}', \pi) < f(I, \mathbf{A}', \pi') \end{array} \right\} \implies f(I, \mathbf{A} + \mathbf{A}', \pi) < f(I, \mathbf{A} + \mathbf{A}', \pi'),$
- (ii) $\left. \begin{array}{l} f(I, \mathbf{A}, \pi) = f(I, \mathbf{A}, \pi') \\ f(I, \mathbf{A}', \pi) = f(I, \mathbf{A}', \pi') \end{array} \right\} \implies f(I, \mathbf{A} + \mathbf{A}', \pi) = f(I, \mathbf{A} + \mathbf{A}', \pi').$

Note that *every* rule is an argmax rule— f can simply be the indicator function on the outcome of the rule for a given instance and profile—but not all are monotonic.

Are the monotonic argmax rules good candidates for being MLEs? Yes, they are, as we can show that monotonic argmax rules all satisfy weak reinforcement.

Proposition 5.3.2. *Every monotonic argmax rule satisfies weak reinforcement.*

PROOF. Consider the monotonic argmax rule F defined by the function f . Let I be an instance, and \mathbf{A} and \mathbf{A}' two profiles over I such that $F(I, \mathbf{A}) = F(I, \mathbf{A}')$. We show that we also have $F(I, \mathbf{A} + \mathbf{A}') = F(I, \mathbf{A})$.

Consider a budget allocation $\pi \in F(I, \mathbf{A})$. Since F is an argmax rule, for all $\pi' \in \text{FEAS}(I) \setminus F(I, \mathbf{A})$, we know that $f(I, \mathbf{A}, \pi) > f(I, \mathbf{A}, \pi')$. This also holds for \mathbf{A}' . By the definition of a monotonic argmax rule, we immediately get that $f(I, \mathbf{A} + \mathbf{A}', \pi) > f(I, \mathbf{A} + \mathbf{A}', \pi')$.

Moreover, for any two budget allocations $\pi, \pi' \in F(I, \mathbf{A})$, we have $f(I, \mathbf{A}, \pi) = f(I, \mathbf{A}, \pi')$. The same also holds for \mathbf{A}' . Since F is monotonic, we thus immediately get that $f(I, \mathbf{A} + \mathbf{A}', \pi) = f(I, \mathbf{A} + \mathbf{A}', \pi')$.

Overall, we proved that (i) no budget allocation that was winning under \mathbf{A} or \mathbf{A}' is dominated under $\mathbf{A} + \mathbf{A}'$, and (ii) that all budget allocations that were winning under \mathbf{A} and \mathbf{A}' all have the same score. It is then immediate that $F(I, \mathbf{A} + \mathbf{A}') = F(I, \mathbf{A})$. \square

In the following we will introduce and study several concrete examples of monotonic argmax rules, based either on the Nash or on the utilitarian social welfare.

5.3.1 Nash Social Welfare

We first study rules that are based on the Nash social welfare. Remember from Section 3.5.2 that according to the Nash social welfare, the score of a budget allocation π is measured as the product of the satisfactions of the agents in π . It thus tries to reach outcomes in which the distribution of the satisfactions of the agents is balanced. Formally, given an instance I , a profile \mathbf{A} , and a satisfaction function sat , we define the PB rule that selects budget allocations maximising the Nash social welfare as:³⁸

$$\text{NASH-SW}[sat](I, \mathbf{A}) = \arg \max_{\pi \in \text{FEAS}(I)} \prod_{i \in \mathcal{N}} sat_i(\pi).$$

It can easily be checked that $\text{NASH-SW}[sat]$ is the monotonic argmax rule defined by the function f_{sat}^{NASH} where for any I, \mathbf{A} and $\pi \in \text{FEAS}(I)$, we have:

$$f_{sat}^{\text{NASH}}(I, \mathbf{A}, \pi) = \begin{cases} \sum_{i \in \mathcal{N}} \log(sat_i(\pi)) & \text{if } sat_i(\pi) \neq 0 \text{ for all } i \in \mathcal{N}, \\ 0 & \text{otherwise.} \end{cases}$$

³⁸Note that we slightly abuse the notation here as $\text{NASH-SW}[sat]$ is a PB rule and not a welfare measure as we introduced it in Section 3.5.2.

From this, and Proposition 5.3.2, it is clear that $\text{NASH-SW}[sat]$ satisfies weak reinforcement for all sat .

We will consider $\text{NASH-SW}[sat]$ for different satisfaction functions, starting with the cardinality and cost ones.

Cardinality and Cost Satisfaction

We first study $\text{NASH-SW}[sat^{card}]$ and $\text{NASH-SW}[sat^{cost}]$. Observe that under the usual assumption that all projects are approved by at least one agent, these two rules are exhaustive (in the sense that at least one of the returned budget allocations will be exhaustive).

We start by introducing \mathcal{M}_{Ncost} , a noise model for which $\text{NASH-SW}[sat^{cost}]$ could, at a first glance, be the MLE. It is defined such that for all I , A and π^* , we have:

$$\mathbb{P}_{\mathcal{M}_{Ncost}}(A \mid \pi^*, I) = \frac{1}{Z_{\pi^*}^{Ncost}} c(A \cap \pi^*),$$

where, $Z_{\pi^*}^{Ncost}$ is a suitable normalisation factor ensuring that \mathcal{M}_{Ncost} is a well-defined probability distribution.

Under this noise model, the probability of generating a given ballot A increases with the cost of the ground-truth projects in A . The intuition here is that voters may reflect more carefully on expensive projects and thus are more likely to make correct choices for them. Moreover, the probability of generating A increases linearly in the “quality” of A . How realistic this is, is open to debate. On the one hand, this avoids having to assume extremely high probabilities for correctly identifying particularly expensive projects (in the case where the relationship would not be linear). On the other hand, the probability of generating a ballot that is completely wrong (a ballot not including even a single ground-truth project) is zero.

Under \mathcal{M}_{Ncost} , maximising the likelihood would be similar to maximising the Nash social welfare of a budget allocation, when using the cost satisfaction function. However, for this intuitive connection to hold, it should be that the normalisation factor $Z_{\pi^*}^{Ncost}$ is independent of π^* . Let us look at it in more detail.

Lemma 5.3.3. *For the noise model \mathcal{M}_{Ncost} to be a well-defined probability distribution, it must be the case that:*

$$Z_{\pi^*}^{Ncost} = 2^{|\mathcal{P}|-1} c(\pi^*).$$

PROOF. Consider an instance $I = \langle \mathcal{P}, c, b \rangle$, an approval ballot $A \subseteq \mathcal{P}$, and a ground truth $\pi^* \in \text{FEAS}(I)$. For \mathcal{M}_{Ncost} to be a probability distribution, it should be the case that:

$$\sum_{A \subseteq \mathcal{P}} \mathbb{P}_{\mathcal{M}_{Ncost}}(A \mid \pi^*, I) = 1 \iff Z_{\pi^*}^{Ncost} = \sum_{A \subseteq \mathcal{P}} c(|A \cap \pi^*|).$$

Remember that there are $2^{|\mathcal{P}|}$ subsets of projects and that any project $p \in \mathcal{P}$ appears in exactly half of them. Each time a project $p \in \pi^*$ appears in a subset $A \subseteq \mathcal{P}$, its contribution to the value of $Z_{\pi^*}^{Ncost}$ is exactly $c(p)$. We thus have:

$$Z_{\pi^*}^{Ncost} = \sum_{A \subseteq \mathcal{P}} c(|A \cap \pi^*|) = \sum_{p \in \pi^*} 2^{|\mathcal{P}|-1} c(p) = 2^{|\mathcal{P}|-1} c(\pi^*).$$

This proves the claim. □

This result tells us that the normalisation factor of the noise model \mathcal{M}_{app} depends on the ground truth, thus the value of the likelihood is impacted by the ground truth one is considering when computing the MLE. In particular, we cannot conclude that the rule $\text{NASH-SW}[sat^{cost}]$ is the MLE of this noise model since not all feasible budget allocations have the same cost.

Are there specific cases for which the normalisation factor is independent of the ground truth? Yes, namely for unit-cost instances, as then all exhaustive allocations have the same cost.

Proposition 5.3.4. *Under the assumption that the ground truth is exhaustive, both $\text{NASH-SW}[sat^{card}]$ and $\text{NASH-SW}[sat^{cost}]$ are the MLE of the noise model \mathcal{M}_{Ncost} for unit-cost instances.*

PROOF. Let I be a unit-cost instance I . Consider any two exhaustive budget allocations $\pi, \pi' \in \text{FEAS}_{\text{EX}}(I)$. Since we have $|\pi| = |\pi'| = c(\pi) = c(\pi')$, Lemma 5.3.3 entails that $Z_{\pi}^{Ncost} = Z_{\pi'}^{Ncost}$. For any profile \mathbf{A} , we have then:

$$\begin{aligned} \arg \max_{\pi \in \text{FEAS}_{\text{EX}}(I)} L_{\mathcal{M}_{Ncost}}(\mathbf{A}, \pi, I) &= \arg \max_{\pi \in \text{FEAS}_{\text{EX}}(I)} \prod_{A \in \mathbf{A}} \frac{c(A \cap \pi)}{Z_{\pi}^{Ncost}} \\ &= \arg \max_{\pi \in \text{FEAS}_{\text{EX}}(I)} \prod_{A \in \mathbf{A}} c(A \cap \pi) \\ &= \text{NASH-SW}[sat^{cost}](I, \mathbf{A}). \end{aligned}$$

The last line follows from the fact that $\text{NASH-SW}[sat^{cost}]$ is exhaustive.

Given that on unit-cost instances sat^{card} and sat^{cost} coincide, the result also applies to $\text{NASH-SW}[sat^{card}]$. □

The fact that $\text{NASH-SW}[sat^{card}]$ and $\text{NASH-SW}[sat^{cost}]$ are MLEs for \mathcal{M}_{app} only under some restricted hypothesis is the first hint of a general impossibility result. Indeed, we can show that there are no noise models of which these rule are MLEs.

Theorem 5.3.5. *There is no noise model \mathcal{M} such that either $\text{NASH-SW}[sat^{card}]$ or $\text{NASH-SW}[sat^{cost}]$ is the MLE of \mathcal{M} , not even on unit-cost instances.*

PROOF. Consider an instance I with two projects p_1 and p_2 of cost 1, and a budget limit $b = 2$. Let \mathcal{M} be a generic noise model, and denote by \mathbb{P}_A^π the value of $\mathbb{P}_{\mathcal{M}}(A \mid \pi, I)$ for any A and π . To simplify notation, we omit braces around sets.

For the noise model \mathcal{M} to be a well-defined probability distribution, the following two equalities should be satisfied:

$$\sum_{A \subseteq \mathcal{P}} \mathbb{P}_A^{p_1} = 1 \quad \Leftrightarrow \quad \mathbb{P}_\emptyset^{p_1} + \mathbb{P}_{p_1}^{p_1} + \mathbb{P}_{p_2}^{p_1} + \mathbb{P}_{p_1, p_2}^{p_1} = 1, \quad (5.1)$$

$$\sum_{A \subseteq \mathcal{P}} \mathbb{P}_A^{p_1, p_2} = 1 \quad \Leftrightarrow \quad \mathbb{P}_\emptyset^{p_1, p_2} + \mathbb{P}_{p_1}^{p_1, p_2} + \mathbb{P}_{p_2}^{p_1, p_2} + \mathbb{P}_{p_1, p_2}^{p_1, p_2} = 1. \quad (5.2)$$

Now, on the single-agent profile $\mathbf{A} = (\emptyset)$, $\text{NASH-SW}[sat^{card}]$ returns $\text{FEAS}(I)$. So for $\text{NASH-SW}[sat^{card}]$ to be the MLE of \mathcal{M} , we must have $\mathbb{P}_\emptyset^{p_1} = \mathbb{P}_\emptyset^{p_1, p_2}$. Moreover, on $\mathbf{A} = (\{p_1\})$, we have $\text{NASH-SW}[sat^{card}](I, \mathbf{A}) = \{\{p_1\}, \{p_1, p_2\}\}$, so $\mathbb{P}_{p_1}^{p_1} = \mathbb{P}_{p_1}^{p_1, p_2}$. Using these two equalities and by subtracting (5.2) from (5.1), we get:

$$(\mathbb{P}_{p_2}^{p_1} - \mathbb{P}_{p_2}^{p_1, p_2}) + (\mathbb{P}_{p_1, p_2}^{p_1} - \mathbb{P}_{p_1, p_2}^{p_1, p_2}) = 0. \quad (5.3)$$

Now, since $\text{NASH-SW}[sat^{card}](I, (\{p_2\})) = \{\{p_2\}, \{p_1, p_2\}\}$, we must have $\mathbb{P}_{p_2}^{p_1, p_2} > \mathbb{P}_{p_2}^{p_1}$. For $\mathbf{A} = (\{p_1, p_2\})$, we have $\text{NASH-SW}[sat^{card}](I, \mathbf{A}) = \{\{p_1, p_2\}\}$. We can then derive $\mathbb{P}_{p_1, p_2}^{p_1, p_2} > \mathbb{P}_{p_1, p_2}^{p_1}$. These two last inequalities contradict (5.3). It is then impossible for $\text{NASH-SW}[sat^{card}]$ to be the MLE of \mathcal{M} on I . From the unit-cost assumption, it is clear that this also applies to $\text{NASH-SW}[sat^{cost}]$. \square

This impossibility result concludes our analysis of $\text{NASH-SW}[sat^{card}]$ and $\text{NASH-SW}[sat^{cost}]$. We will now consider “normalised” satisfaction functions.

Normalised Satisfaction

In the hope of overcoming Theorem 5.3.5, we also consider normalised variants of the two satisfaction functions sat^{card} and sat^{cost} . In these variants, the satisfaction of an agent is expressed in terms of the proportion of the outcome that satisfies them. We denote them $\overline{sat^{card}}$ and $\overline{sat^{cost}}$ respectively. Formally, for agent $i \in \mathcal{N}$ with ballot A_i and any subset of projects $P \subseteq \mathcal{P}$, we have:

$$\overline{sat^{card}}_i(P) = \frac{|A_i \cap P|}{|P|} \quad \overline{sat^{cost}}_i(P) = \frac{c(A_i \cap P)}{c(P)}.$$

Note that strictly speaking $\overline{sat^{card}}$ and $\overline{sat^{cost}}$ are *not* satisfaction function in the sense of Definition 2.2.1 as they depend on the projects that have been selected, but not

approved by the agent. Slightly abusing the notation here, we will still use them as if they were satisfaction functions, and for instance use the notation $\text{NASH-SW}[\overline{\text{sat}^{\text{card}}}]$.

It is also important to observe that these normalised satisfaction functions are similar in spirit to the concept of *relative satisfaction* that will be introduced and studied in Chapter 7 (see Definition 7.2.4). However, while the denominator is defined here with respect to the budget allocation, the denominator in the relative satisfaction will be defined with respect to the ballot of the agent (so $|A_i|$ and $c(A_i)$ instead of $|P|$ and $c(P)$). They are also similar to the proportionality requirements defined relative to the budget allocation that Aziz, Lee and Talmon (2018) introduced (discussed in Section 3.3.5).

It is worth noting that the rules $\text{NASH-SW}[\overline{\text{sat}^{\text{card}}}]$ and $\text{NASH-SW}[\overline{\text{sat}^{\text{cost}}}]$ can lead to extreme behaviours. For example, consider an instance with budget limit b , that we assume to be even, and a set of projects $\mathcal{P} = \{p^*\} \cup \{p_1, \dots, p_b\}$ of arbitrary cost lower than b . Consider the two-agent profile \mathbf{A} such that:

$$A_1 = \{p^*\} \cup \{p_1, p_3, \dots, p_{b-1}\} \quad A_2 = \{p^*\} \cup \{p_2, p_4, \dots, p_b\}.$$

Then, according to both rules, selecting just p^* is better than anything else. Even if this can seem extreme, these rules can still be justified when considering voters who would rather save public money than use it on projects they do not approve. This would correspond to associating a strong rejection, rather than indifference, with the action of not approving a project (we will come back to that in Chapter 6).

Note that this example also implies that the rules are not exhaustive, even on unit-cost instances.

Let us first investigate the rule $\text{NASH-SW}[\overline{\text{sat}^{\text{cost}}}]$. We will continue using the noise model $\mathcal{M}_{N_{\text{cost}}}$ introduced earlier.

Recall the expression we found for the normalisation factor $Z_{\pi^*}^{N_{\text{cost}}}$ in Lemma 5.3.3. Plugging it into the definition of $\mathcal{M}_{N_{\text{cost}}}$, we obtain the following expression for any instance I , approval ballot A , and ground truth π^* :

$$\mathbb{P}_{\mathcal{M}_{N_{\text{cost}}}}(A \mid \pi^*, I) = \frac{1}{2^{|\mathcal{P}|-1}} \frac{c(A \cap \pi^*)}{c(\pi^*)}.$$

From this, it should be clear that $\text{NASH-SW}[\overline{\text{sat}^{\text{cost}}}]$ is the MLE of $\mathcal{M}_{N_{\text{cost}}}$.

Theorem 5.3.6. *The rule $\text{NASH-SW}[\overline{\text{sat}^{\text{cost}}}]$ is the MLE of the noise model $\mathcal{M}_{N_{\text{cost}}}$.*

PROOF. Let $I = \langle \mathcal{P}, c, b \rangle$ be an instance. The likelihood of a profile \mathbf{A} and a budget allocation $\pi \in \text{FEAS}(I)$ under the noise model $\mathcal{M}_{N_{\text{cost}}}$ is:

$$L_{\mathcal{M}_{N_{\text{cost}}}}(\mathbf{A}, \pi, I) = \prod_{A \in \mathbf{A}} \frac{1}{2^{|\mathcal{P}|-1}} \frac{c(A \cap \pi)}{c(\pi)} = \left(\frac{1}{2^{|\mathcal{P}|-1}} \right)^{|\mathbf{A}|} \prod_{A \in \mathbf{A}} \frac{c(A \cap \pi)}{c(\pi)}.$$

Since the first multiplicative factor in the above expression is constant over all budget allocations, we have:

$$\arg \max_{\pi \in \text{FEAS}(I)} L_{\mathcal{M}_{N_{\text{cost}}}}(\mathbf{A}, \pi, I) = \arg \max_{\pi \in \text{FEAS}(I)} \prod_{A \in \mathbf{A}} \frac{c(A \cap \pi)}{c(\pi)} = \text{NASH-SW}[\overline{\text{sat}^{\text{cost}}}] (I, \mathbf{A}).$$

Thus, maximising the likelihood under $\mathcal{M}_{N_{\text{cost}}}$ is the same as maximising the social welfare in the sense of $\text{NASH-SW}[\overline{\text{sat}^{\text{cost}}}]$. \square

We have finally been able to find a PB rule that can be interpreted as an MLE. In the following we will show a similar result for $\text{NASH-SW}[\overline{\text{sat}^{\text{card}}}]$. For this rule we introduce a new noise model denoted by $\mathcal{M}_{N_{\text{app}}}$. It is such that for any instance I , approval ballot A , and ground truth π^* , we have:

$$\mathbb{P}_{\mathcal{M}_{N_{\text{app}}}}(A \mid \pi^*, I) = \frac{1}{Z_{\pi^*}^{N_{\text{app}}}} |A \cap \pi^*|,$$

where $Z_{\pi^*}^{N_{\text{app}}}$ is a suitable normalisation factor.

Using similar proof technique as used above, we show that $\text{NASH-SW}[\overline{\text{sat}^{\text{card}}}]$ is the MLE of $\mathcal{M}_{N_{\text{app}}}$.

Theorem 5.3.7. *The rule $\text{NASH-SW}[\overline{\text{sat}^{\text{card}}}]$ is the MLE of the noise model $\mathcal{M}_{N_{\text{app}}}$.*

PROOF. Let us first compute the exact value of the normalisation factor $Z_{\pi^*}^{N_{\text{app}}}$. For the noise model $\mathcal{M}_{N_{\text{app}}}$ to be a well-defined probability distribution, the following must hold:

$$\sum_{A \subseteq \mathcal{P}} \mathbb{P}_{\mathcal{M}_{N_{\text{app}}}}(A \mid \pi^*, I) = 1 \Leftrightarrow Z_{\pi^*}^{N_{\text{app}}} = \sum_{A \subseteq \mathcal{P}} |A \cap \pi^*| \Leftrightarrow Z_{\pi^*}^{N_{\text{app}}} = 2^{|\mathcal{P}|-1} |\pi^*|.$$

Hence, given a profile \mathbf{A} , we have:

$$\begin{aligned} \arg \max_{\pi \in \text{FEAS}(I)} L_{\mathcal{M}_{N_{\text{app}}}}(\mathbf{A}, \pi) &= \arg \max_{\pi \in \text{FEAS}(I)} \prod_{A \in \mathbf{A}} \frac{1}{2^{|\mathcal{P}|-1}} \frac{|A \cap \pi|}{|\pi|} \\ &= \arg \max_{\pi \in \text{FEAS}(I)} \prod_{A \in \mathbf{A}} \frac{|A \cap \pi|}{|\pi|} \\ &= \text{NASH-SW}[\overline{\text{sat}^{\text{card}}}] (I, \mathbf{A}). \end{aligned}$$

This immediately implies that $\text{NASH-SW}[\overline{\text{sat}^{\text{card}}}]$ is the MLE of $\mathcal{M}_{N_{\text{app}}}$. \square

This concludes our study of the Nash social welfare. We now turn to the more standard measure of the social welfare: the utilitarian social welfare.

5.3.2 Utilitarian Social Welfare

Let us now turn to the analysis of monotonic argmax rules defined in terms of utilitarian social welfare. As we did before, we will consider different satisfaction functions.

Cardinality and Cost Satisfaction

We start with the usual cardinality and cost satisfaction functions, that correspond to the rules MAXCARD and MAXCOST as introduced in Chapter 2. Remember that these rules return the budget allocations maximising the sum of the satisfaction of the agents. It is thus clear that they are monotonic argmax rule (since the score they use is additive) and thus both satisfy weak reinforcement.

Following the idea developed by Conitzer and Sandholm (2005) for scoring rules (in the standard voting framework), we introduce the noise model \mathcal{M}_{app} . It is defined such that for any instance $I = \langle \mathcal{P}, c, b \rangle$, and approval ballot $A \subseteq \mathcal{P}$, and ground truth $\pi^* \in \text{FEAS}(I)$:

$$\mathbb{P}_{\mathcal{M}_{app}}(A \mid \pi^*, I) = \frac{1}{Z_{\pi^*}^{app}} \prod_{p \in \mathcal{P}} 2^{\mathbb{1}_{p \in A \cap \pi^*}} = \frac{1}{Z_{\pi^*}^{app}} 2^{|A \cap \pi^*|},$$

where $Z_{\pi^*}^{app}$ is a suitable normalisation factor.

\mathcal{M}_{app} is a particularly simple manifestation of what we would expect to see in a noise model: any possible ballot might be generated in principle, but the probability of generating ballot A increases exponentially with the size of the intersection between A and the ground truth.

With this noise model, maximising the likelihood may appear to have the same effect as maximising the approval score of a budget allocation. It could then be that the approval maximising rule is the MLE of \mathcal{M}_{app} . However, for this to hold, one has to have a closer look at the normalisation factor.

Lemma 5.3.8. *For the noise model \mathcal{M}_{app} to be a well-defined probability distribution, it must be the case that:*

$$Z_{\pi^*}^{app} = 2^{|\mathcal{P}|} \left(\frac{3}{2} \right)^{|\pi^*|}.$$

PROOF. Consider any instance $I = \langle \mathcal{P}, c, b \rangle$. Let $A \subseteq \mathcal{P}$ be an approval ballot and $\pi^* \in \text{FEAS}(I)$ a ground truth. For \mathcal{M}_{app} to be a probability distribution, it should be the case that:

$$\sum_{A \subseteq \mathcal{P}} \mathbb{P}_{\mathcal{M}_{app}}(A \mid \pi^*, I) = 1 \quad \iff \quad Z_{\pi^*}^{app} = \sum_{A \subseteq \mathcal{P}} 2^{|A \cap \pi^*|}.$$

Let's do some combinatorics. For $k \in \{0, \dots, |\pi^*|\}$, how many subsets of \mathcal{P} will intersect with π^* on exactly k projects? A suitable subset will consists of k

projects from π^* that make up the intersection and any number $j \in \{0, \dots, |\mathcal{P}| - |\pi^*|\}$ of projects from $\mathcal{P} \setminus \pi^*$ that do not have any impact on the intersection. Each such subset of projects contributes 2^k to the value of $Z_{\pi^*}^{app}$. We thus have:

$$\begin{aligned} Z_{\pi^*}^{app} &= \sum_{k=0}^{|\pi^*|} 2^k \sum_{j=0}^{|\mathcal{P}|-|\pi^*|} \binom{|\pi^*|}{k} \binom{|\mathcal{P}|-|\pi^*|}{j} \\ &= \sum_{k=0}^{|\pi^*|} \binom{|\pi^*|}{k} 2^k \sum_{j=0}^{|\mathcal{P}|-|\pi^*|} \binom{|\mathcal{P}|-|\pi^*|}{j} \\ &= 2^{|\mathcal{P}|-|\pi^*|} \sum_{k=0}^{|\pi^*|} \binom{|\pi^*|}{k} 2^k \\ &= 2^{|\mathcal{P}|} \left(\frac{3}{2}\right)^{|\pi^*|} \end{aligned}$$

where the last two lines are derived from the binomial expansion. \square

The normalisation factor of \mathcal{M}_{app} thus depends on the ground truth, since not all feasible budget allocations have the same cardinality. We thus cannot conclude that the approval maximising rule is the MLE of this noise model.

Interestingly, this is not the case on unit-cost instances when considering exhaustive budget allocations.

Proposition 5.3.9. *Under the assumption that the ground truth is exhaustive, both MAXCARD and MAXCOST are the MLE of the noise model \mathcal{M}_{app} for unit-cost instances.*

PROOF. Let I be a unit-cost instance. For any two exhaustive budget allocations π and $\pi' \in \text{FEAS}_{\text{EX}}(I)$, by virtue of Lemma 5.3.8, we have $Z_{\pi}^{app} = Z_{\pi'}^{app}$. So, for any profile \mathbf{A} , we have:

$$\begin{aligned} \arg \max_{\pi \in \text{FEAS}_{\text{EX}}(I)} L_{\mathcal{M}_{app}}(\mathbf{A}, \pi, I) &= \arg \max_{\pi \in \text{FEAS}_{\text{EX}}(I)} \prod_{A \in \mathbf{A}} \frac{1}{Z_{\pi}^{app}} 2^{|A \cap \pi|} \\ &= \arg \max_{\pi \in \text{FEAS}_{\text{EX}}(I)} 2^{\sum_{A \in \mathbf{A}} |A \cap \pi|} \\ &= \arg \max_{\pi \in \text{FEAS}_{\text{EX}}(I)} \sum_{A \in \mathbf{A}} |A \cap \pi| \\ &= \text{MAXCARD}(I, \mathbf{A}). \end{aligned}$$

The last line follows from the fact that MAXCARD is exhaustive.

MAXCARD coincides thus with the MLE on I for the noise model \mathcal{M}_{app} . Moreover, since MAXCARD and MAXCOST coincide on unit-cost instances, the result also applies to MAXCOST. \square

This first result is only half satisfactory. Can we find an impossibility result similar to the one we had for $\text{NASH-SW}[sat^{card}]$ and $\text{NASH-SW}[sat^{cost}]$? It is actually easy to see that the proof we gave for Theorem 5.3.5 also works for both MAXCARD and MAXCOST.

Theorem 5.3.10. *There is no noise model \mathcal{M} such that either MAXCARD or MAXCOST is the MLE of \mathcal{M} , not even on unit-cost instances.*

PROOF. Consider the instance I used in the proof of Theorem 5.3.5. We claim that for all profiles that are relevant for the proof, MAXCARD and $\text{NASH-SW}[sat^{card}]$ coincide. We list them below.

$$\begin{aligned} \text{MAXCARD}(I, (\{p_1\})) &= \{\{p_1\}, \{p_1, p_2\}\} = \text{NASH-SW}[sat^{card}](I, (\{p_1\})). \\ \text{MAXCARD}(I, (\{p_2\})) &= \{\{p_2\}, \{p_1, p_2\}\} = \text{NASH-SW}[sat^{card}](I, (\{p_2\})). \\ \text{MAXCARD}(I, (\{p_1, p_2\})) &= \{\{p_1, p_2\}\} = \text{NASH-SW}[sat^{card}](I, (\{p_1, p_2\})). \\ \text{MAXCARD}(I, (\emptyset)) &= \text{FEAS}(I) = \text{NASH-SW}[sat^{card}](I, (\emptyset)). \end{aligned}$$

Given that on unit-cost instances MAXCARD and MAXCOST coincide, the result also applies to MAXCOST. \square

We conclude by discussing the rules that greedily approximate the outcome of MAXCARD and MAXCOST. These rules deserve a special place in our analysis since GREEDCOST is the most widely used rule in practice.

The first observation to make is that the three rules GREEDCARD, GREEDCOST and MAXCARD all coincide on unit-cost instances. Thus, both Proposition 5.3.9 and Theorem 5.3.10 apply to GREEDCARD and GREEDCOST as well. This also applies to any refinement of the rules, such as the *leximax rule* (see Chapter 6).

Corollary 5.3.11. *Under the assumption that the ground truth is exhaustive, both GREEDCARD and GREEDCOST are the MLE of \mathcal{M}_{app} for unit-cost instances.*

Moreover, for unconstrained ground truths, there is no noise model \mathcal{M} such that either GREEDCARD or GREEDCOST is the MLE of \mathcal{M} , not even on unit-cost instances.

We conclude our discussion of sat^{card} and sat^{cost} by showing for the sake of completeness that GREEDCARD and GREEDCOST fail weak reinforcement. Remember that this was not the case for MAXCARD and MAXCOST as they are monotonic argmax rules.

Proposition 5.3.12. *Both GREEDCARD and GREEDCOST fail weak reinforcement.*

PROOF. Consider an instance I with three projects denoted by p_1 , p_2 and p_3 , a budget limit of $b = 4$, and costs as shown in the table below. Moreover, consider two profiles \mathbf{A} and \mathbf{A}' in which the approval scores are as presented below.

	Cost	Approval score in \mathbf{A}	Approval score in \mathbf{A}'	Approval score in $\mathbf{A} + \mathbf{A}'$
p_1	2	10	1	11
p_2	2	1	10	11
p_3	3	9	9	18

One can easily check that both GREEDCARD and GREEDCOST would return $\{\{p_1, p_2\}\}$ on both \mathbf{A} and \mathbf{A}' . However, on the joint profile $\mathbf{A} + \mathbf{A}'$, the two rules would return $\{\{p_3\}\}$. They thus violates weak reinforcement. \square

Normalised Satisfaction

Let us conclude our formal analysis by briefly mentioning the utilitarian social welfare with the normalised satisfaction functions \overline{sat}^{card} and \overline{sat}^{cost} . Again abusing notation, we denote these two rules by $UTIL-SW[\overline{sat}^{card}]$ and $UTIL-SW[\overline{sat}^{cost}]$.

For the same reasons as for $NASH-SW[\overline{sat}^{card}]$ and $NASH-SW[\overline{sat}^{cost}]$, the two rules $UTIL-SW[\overline{sat}^{card}]$ and $UTIL-SW[\overline{sat}^{cost}]$ are not exhaustive. Analysing the epistemic status of these rules however turns out to be rather intricate, even on unit-cost instances. Indeed, it is less clear what a suitable noise model might look like, especially due to the complications related to the potential normalisation factor. Exploring these rules remains an interesting open problem.

5.4 Summary

In this chapter, we have presented an analysis of PB rules from an epistemic perspective. We reviewed a large set of rules, including the most celebrated in the literature (MES, SEQPHRAG, MAXCARD and MAXCOST), the most used in practice (GREEDCOST), and some others that received little attention so far. It proved rather difficult to find rules that can be interpreted as MLEs. Indeed, already a large number of rules failed weak reinforcement, a necessary condition for being an MLE. Then, even when focusing on a class of rules that all satisfy weak reinforcement, MLEs have been hard to find. In particular our two positive results concern rules that suffer other significant drawbacks. All the results presented in this chapter are displayed in Table 5.4.1.

	Unit-cost_{EX}	Unit-cost	General case
SEQPHRAG	X Proposition 5.2.1	X Proposition 5.2.1	X Proposition 5.2.1
MES[<i>sat</i>] for all <i>sat</i>	–	X Proposition 5.2.2	X Proposition 5.2.2
GREEDCARD and GREEDCOST	✓ Corollary 5.3.11	X Corollary 5.3.11	X Corollary 5.3.11
MAXCARD and MAXCOST	✓ Proposition 5.3.9	X Theorem 5.3.10	X Theorem 5.3.10
UTIL-SW[$\overline{sat^{card}}$]	–	?	?
UTIL-SW[$\overline{sat^{cost}}$]	–	?	?
NASH-SW[$\overline{sat^{card}}$] and NASH-SW[$\overline{sat^{cost}}$]	✓ Proposition 5.3.4	X Theorem 5.3.5	X Theorem 5.3.5
NASH-SW[$\overline{sat^{card}}$]	–	✓ Theorem 5.3.6	✓ Theorem 5.3.6
NASH-SW[$\overline{sat^{card}}$]	–	✓ Theorem 5.3.7	✓ Theorem 5.3.7

Table 5.4.1: Summary of the results for all the rules on specific classes of instances. A check-mark **✓** indicates that there exists a noise model for which the rule is an MLE and a cross-mark **X** the fact that it is impossible to find such a noise model. The _{EX} subscript signifies that we make the additional assumption that the ground truth is exhaustive. This assumption would not be meaningful for non-exhaustive rules (remember that SEQPHRAG is exhaustive on unit-cost instances).

Part Three

Variations on the Model

Chapter 6

A General Framework for Multi-Constraint Participatory Budgeting

This is the first chapter of Part **Three** where we investigate variations on the model. What this means is that in the coming chapters, we will study variants of the standard model of PB that we introduced in Chapter 2, each time introducing new aspects of PB processes into our formal analysis.

On our agenda for now is the study of PB scenarios with multiple constraints on the outcome. More specifically, we are interested in cases in which the definition of a *feasible* budget allocation does not only depend on its cost, but also on some of its other features. We will thus investigate scenarios for which there are additional constraints on top of the budget limit.

Let us delve into an example right away. Consider the case of Zoiville, a little town in the country of Friendtopia. Zoiville is, as its name suggests, run by Zoi, a powerful but benevolent mayor. Zoi was re-elected on the promise of developing an ambitious cultural programme for the city. The plan included the construction of a brand new international centre for the art of movement, whose design is to be decided in collaboration with the citizens. The architecture of the building has been decided by a committee of experts, but a PB process will be organised to decide on parts of the interior design.

After several public meetings, organised to collect proposals from the citizens, it appears that several of them got quite interested in the process. Two proposals caught Zoi's attention. They were submitted by Arianna and Sirin, two citizens known to have, among other things, very definite opinions when it comes to colour palettes. Both of them proposed to buy a set of sofas to allow visitors to sit in the main hall. However, while Arianna wants the sofas to be olive-green, Sirin wants them to be mustard-yellow. Not knowing, which one to pick, Zoi decides that both proposals

would make it to the next stage, and that voters will decide.³⁹ Of course, both projects cannot be selected together, so an additional constraint to rule out outcomes in which it happens has been added.

Going down the list of submitted proposals, Zoi noticed two other ones that necessitate special treatment. Indeed, after reading Emma’s proposal of having a bike pumping station in the building, Camille proposed to have it decorated by local artists. Once again, these two proposals are not independent as one can only be selected if the other also is. Zoi thus added yet another constraint to ensure that the decoration for the pumping station is only implemented if the latter also is implemented.

In the end, two additional constraints on the outcome have to be considered, each of a different kind. Zoi—who is inclined to follow the latest recommendations of the social choice literature—had already selected her favourite voting method based on the properties that it guarantees. However, she does not know how to adapt it to impose the constraints in a way that also preserves these desirable properties.

The problem highlighted in the above example is a typical one for any formal analysis: the mathematical models we study are very strictly defined and rarely provide flexibility in small variations of the model.

Moreover, while the above is completely fictive, it does exemplify scenarios that can occur in practice. For instance, until 2022, the Parisian PB process had some projects labelled “low income neighbourhood” and a specific number of them had to be selected in the final budget allocation. Another category of projects gathered the ones designed for the whole city (and not a specific neighbourhood), and a lower bound on the number of such projects to be selected was imposed.⁴⁰ Similar constraints have also been observed in PB processes in Lisbon (Allegretti and Antunes, 2014), in Amsterdam (City of Amsterdam, 2022), or in Lyon as we described in Section 1.4.2.

There is thus a need for a formal analysis that can easily account for such additional constraints. The good news is that this is exactly the research question for this chapter. More specifically, this chapter addresses the following question.

How can we design PB rules that are robust against variations of the feasibility constraints for the outcomes?

The approach we will develop consists in using an aggregation framework way more expressive than PB to reason about PB. Doing so will allow us to obtain the flexibility required by our research agenda. More specifically, we will use *Judgement*

³⁹Remember the typical structure of a PB process we outlined in Section 1.1.2. Usually projects brought to the vote are initially proposed by citizens.

⁴⁰It is unfortunately hard to find a trace of this now that the Paris municipality has renewed the platform for the PB process. But it is still mentioned at mairiepariscentre.paris.fr/pages/c-est-parti-pour-le-budget-participatif-2021-16701, for instance. The important part is that “after the vote, between 2 and 5 projects will be chosen per neighbourhood based on their demographics, with a bonus for low income neighbourhoods. Two projects for the entire city will also be selected”.

Aggregation (JA), a framework to aggregate opinions for scenarios where logical constraints on the outcome apply.

The idea of using a general framework does come with some potential downsides. The first one is computational in nature: Determining outcomes in JA is generally computationally demanding. One would then need to ensure that defining PB rules through JA does not render the rules completely unusable because of their computational cost. The second potential downside is axiomatic in nature: JA rules need to be analysed from a PB perspective to ensure that they provide similar guarantees than the PB ones. These two aspects will be covered in this chapter.

Let us highlight that this approach does provide an answer to Zoi’s problem. Indeed, it does provide great flexibility since when incorporating extra constraints for PB, one “only” has to investigate the specific encoding of the constraint on the JA side, thereby immediately making available all other results previously established for the basic framework without those constraints. In particular, once a rule has been proven to satisfy a certain axiom, it will continue to do so, regardless of the extra constraints.⁴¹

For the rest of this chapter, we will first present some additional related work. Then, we will formally define the JA framework and our approach for using JA to reason about PB (Section 6.1). We will then turn to proving the viability of the approach, first from a computational perspective (Section 6.2), and then from an axiomatic perspective (Section 6.3 and Section 6.4). A short summary will conclude this chapter (Section 6.5).

Additional Related Work. On top of the related work on PB presented in Chapter 3, we provide here some general references on Judgment Aggregation (JA). The setting was initially introduced by List and Pettit (2002). The study of how to embed preference aggregation problems into JA dates back to at least Dietrich and List (2007a). The systematic study of how to embed voting rules into JA was then initiated by Lang and Slavkovik (2013) and later refined by Endriss (2018). Our work on the algorithmic aspects of such embeddings is generalising results of De Haan (2018), whose paper is also the first example for work investigating the embedding of PB into JA. Lately, a similar approach has been followed by Chingoma, Endriss and De Haan (2022) to embed multi-winner voting into JA.

6.1 Frameworks

The main aim of this chapter is to present how to reason about PB by using Judgment Aggregation (JA). In this section we present a variation of the standard model

⁴¹To be precise, this is true for axioms that have a “universal flavour”, *i.e.*, for axioms that stipulate that certain conditions must be satisfied for all relevant situations. In the other case, when the axioms required the existence of a situation exhibiting a specific property of interest, it can be that additional constraints make it impossible for the property to occur, rendering the axiom vacuous for instance.

of PB where multiple resources are involved, and the basic definitions of these two frameworks. We also define the main concept of this paper, namely embeddings of PB instances into JA.

6.1.1 Participatory Budgeting with Multiple Resources

For the PB side, we mainly adopt the notation introduced in Chapter 2. However, for this chapter we will consider a model of PB in which the costs and the budget limit are expressed over several resources. Additional notation is thus needed. We will not redefine all the concepts as most of the “new” definitions for the multi-resource setting are simple rewriting of the single-resource case.

An instance of PB with multiple resources is a tuple $I = \langle \mathcal{P}, \mathcal{R}, c, \mathbf{b} \rangle$. As before \mathcal{P} denotes the set of projects. The *set of resources* is $\mathcal{R} = \{r_1, \dots, r_d\}$. The budget limit $\mathbf{b} = (b_1, \dots, b_d)$ is now a vector of length $|\mathcal{R}|$ in which $b_i \in \mathbb{R}_{>0}$ indicates the budget limit in terms of resource r_i . The cost function $c : \mathcal{P} \times \mathcal{R} \rightarrow \mathbb{R}_{>0}$ now maps any project $p \in \mathcal{P}$ and resource $r \in \mathcal{R}$ to p 's cost in terms of the resource r . Overloading notation, we use $\mathbf{c}(p) = (c(p, r_1), \dots, c(p, r_d))$ to denote the cost vector of project p . Finally, for any subset of projects $P \subseteq \mathcal{P}$, let $c(P, r) = \sum_{p \in P} c(p, r)$ and $\mathbf{c}(P) = \sum_{p \in P} \mathbf{c}(p)$. We denote by \mathcal{I} the set of all instances of PB with multiple resources.

A solution for an instance $I = \langle \mathcal{P}, \mathcal{R}, c, \mathbf{b} \rangle$ is a *budget allocation* $\pi \subseteq \mathcal{P}$ that is *feasible*, i.e., such that if $\mathbf{c}(\pi) \leq \mathbf{b}$, where for two same-sized vectors $\mathbf{v} = (v_1, \dots, v_k)$ and $\mathbf{v}' = (v'_1, \dots, v'_k)$, $\mathbf{v} \leq \mathbf{v}'$ indicates that $v_j \leq v'_j$ for all $j \in \{1, \dots, k\}$. We will consider irresolute rules in this chapter, so a *PB rule* maps each instance I and each profile \mathbf{A} to a non-empty set $F(I, \mathbf{A}) \subseteq \text{FEAS}(I)$ of feasible budget allocations.⁴²

6.1.2 Judgment Aggregation

The specific JA framework we use in this chapter is known as *binary aggregation with integrity constraints* (Grandi and Endriss, 2011). We introduce it in the following. While this framework is most convenient for our purposes, the original framework of List and Pettit (2002) could be used as well, given that it is known that the former can be efficiently embedded into the latter (Endriss, Grandi, De Haan and Lang, 2016).

Let $\mathcal{L}_{\mathfrak{X}}$ be the set of *propositional formulas* over a given set \mathfrak{X} of propositional atoms, using the usual connectives $\neg, \vee, \wedge, \rightarrow$, and logical constants \perp and \top . Propositional atoms and their negations are called *literals*. For any subset of atoms $X \subseteq \mathfrak{X}$, we write $\text{Lit}(X) = X \cup \{\neg x \mid x \in X\}$ for the set of literals corresponding to X . We often use x_i to denote atoms and ℓ_{x_i} to denote literals corresponding to x_i , i.e., $\ell_{x_i} \in \{x_i, \neg x_i\}$. We say that the literal ℓ_{x_i} is *positive* if $\ell_{x_i} = x_i$ and *negative* if $\ell_{x_i} = \neg x_i$. A *truth assignment* $\alpha : \mathfrak{X} \rightarrow \{0, 1\}$ is a mapping indicating for each atom

⁴²Observe that $\text{FEAS}(I)$ is never empty since the empty set of projects is always feasible. This, however, is not true for some of the extensions discussed in Section 6.2.

its truth value (1 being true and 0 false). For $\ell_{x_i} \in \text{Lit}(\mathfrak{X})$, let $\alpha(\ell_{x_i}) = \alpha(x_i)$ if ℓ_{x_i} is positive and let $\alpha(\ell_{x_i}) = 1 - \alpha(x_i)$ otherwise. We write $\alpha \models \varphi$ whenever α is a model of φ according to the usual semantics of propositional logic (Van Dalen, 2013).

In the context of JA, the atoms in \mathfrak{X} represent *propositions* an agent may either *accept* or *reject*. A *judgment* J is a set $J \subseteq \mathfrak{X}$, indicating which propositions are accepted. Let $\text{aug}(J) = J \cup \{\neg x \mid x \in \mathfrak{X} \setminus J\}$ be the judgment J augmented with the negative literals for the propositions for which no literal occurs in J . Observe that a judgment J can be equivalently described as the truth assignment α such that $\alpha(x) = 1$ if and only if $x \in J$. In our examples, when we do not explicitly specify the status of some of the propositions, it is assumed that we only consider judgments (and truth assignments) for which the unspecified propositions are mapped to 0.

An *integrity constraint* $\Gamma \in \mathcal{L}_{\mathfrak{X}}$ is a formula used to constrain the range of admissible judgments. A judgment J *satisfies* Γ —written $J \models \Gamma$ —if J , interpreted as a truth assignment, is a model of Γ . Such a judgment J is then said to be *admissible*. Let $\mathfrak{J}(\Gamma) = \{J \subseteq \mathfrak{X} \mid J \models \Gamma\}$ be the set of all admissible judgments for any $\Gamma \in \mathcal{L}_{\mathfrak{X}}$. A *JA instance* is simply an integrity constraint Γ .

We again use $\mathcal{N} = \{1, \dots, n\}$ to denote the set of *agents*. Each agent $i \in \mathcal{N}$ provides us with a judgment J_i , resulting in a *judgment profile* $\mathbf{J} = (J_1, \dots, J_n)$. For a profile \mathbf{J} and a literal $\ell \in \text{Lit}(\mathfrak{X})$, we write $n_{\ell}^{\mathbf{J}} = \sum_{i \in \mathcal{N}} \mathbb{1}_{\ell \in \text{aug}(J_i)}$ for the number of supporters of ℓ (note the analogy with *app*). The *majoritarian outcome* for a profile \mathbf{J} , denoted by $m(\mathbf{J})$, is the set of literals supported by a (strict) majority of agents:

$$m(\mathbf{J}) = \{\ell \in \text{Lit}(\mathfrak{X}) \mid n_{\ell}^{\mathbf{J}} > n/2\}.$$

A *JA rule* is a function F taking as input an integrity constraint Γ and a judgment profile \mathbf{J} and returning a non-empty set $F(\Gamma, \mathbf{J}) \subseteq \mathfrak{J}(\Gamma)$ of admissible judgments. Note how we use F to represent both a JA rule and a PB rule. Finally, observe that no assumption is made about the profile. In particular, we do not require $J_i \models \Gamma$ for any of the agents $i \in \mathcal{N}$.

Before reviewing a number of well-known concrete JA rules, let us first introduce a very general class of JA rules.

Definition 6.1.1 (Additive Rules). *A JA rule F is said to be additive rule if there exists a function $f : (2^{\mathfrak{X}})^n \times \text{Lit}(\mathfrak{X}) \rightarrow \mathbb{R}$ mapping judgment profiles and literals to real values, such that, for every integrity constraint $\Gamma \in \mathcal{L}_{\mathfrak{X}}$ and every profile $\mathbf{J} \in (2^{\mathfrak{X}})^n$, we have:*

$$F(\Gamma, \mathbf{J}) = \arg \max_{J \in \mathfrak{J}(\Gamma)} \sum_{\ell \in \text{aug}(J)} f(\mathbf{J}, \ell).$$

The class of additive rules generalises both the *scoring rules* introduced by Dietrich (2014) and the *additive majority rules* (AMRs) defined by Nehring and Pivato (2019).

More specifically, a scoring rule is associated with a scoring function $s : 2^{\mathcal{X}} \times \text{Lit}(\mathcal{X}) \rightarrow \mathbb{R}$ mapping judgments and literals to scores, and corresponds to the additive rule defined with respect to the function f such that:

$$f(\mathbf{J}, \ell) = \sum_{i \in \mathcal{N}} s(J_i, \ell).$$

An AMR is associated with a non-decreasing gain function $g : \{0, \dots, n\} \rightarrow \mathbb{R}$ with $g(k) < g(k')$ for any $k < \frac{n}{2} \leq k'$ that maps the number of supporters of a literal to a score, and is an additive rule defined with respect to the function f such that:

$$f(\mathbf{J}, \ell) = g(n_{\ell}^{\mathbf{J}}).$$

Three additive rules are of particular importance for our purposes:⁴³

- ▶ The *Slater rule* (Miller and Osherson, 2009; Lang, Pigozzi, Slavkovik and Van der Torre, 2011) selects the admissible outcome closest to the majoritarian outcome in terms of the number of propositions they agree on. It is the AMR associated with the following gain function g :

$$g(x) = \begin{cases} 0 & \text{if } 0 \leq x < \frac{n}{2}, \\ 1 & \text{if } \frac{n}{2} \leq x \leq n. \end{cases}$$

- ▶ The *Kemeny rule* (Pigozzi, 2006; Miller and Osherson, 2009) selects the feasible outcome that is the closest to the profile as a whole. It is both an AMR associated with the gain function $g(x) = x$, and a scoring rule associated with the scoring function $s(J, \ell) = \mathbb{1}_{\ell \in \text{aug}(J)}$.
- ▶ The *leximax rule* (Everaere, Konieczny and Marquis, 2014; Nehring and Pivato, 2019) favours the propositions supported by the largest majorities. It is the AMR defined by the gain function $g(x) = |\mathcal{X}|^x$.

Note that these three rules are all *majority-consistent*, meaning that whenever the majoritarian outcome is admissible, it is the unique judgement returned by the rules.

6.1.3 Embedding PB into Judgement Aggregation

The aim of this chapter is to provide an easy framework to discuss additional constraint for PB problems. To this end, we want to embed PB into JA and then use JA rules to compute budget allocations. A full schematic representation of the process is presented in Figure 6.1.1.

⁴³Most JA rules are known under a variety of different names and have been introduced independently from each other by several different sets of authors. We only provide a small number of representative citations and refer the reader to Endriss (2016) and Lang, Pigozzi, Slavkovik, Van der Torre and Vesic (2017) for further information.

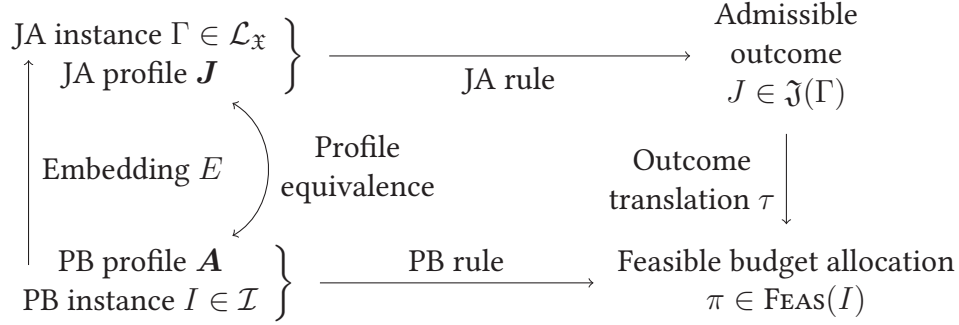


Figure 6.1.1: Full process to use JA rules for PB instances.

For a given PB instance $I = \langle \mathcal{P}, \mathcal{R}, c, \mathbf{b} \rangle$, we introduce one proposition for each project to form the set of propositional atoms \mathfrak{X} . There is thus a direct correspondence between budget allocations $\pi \subseteq \mathcal{P}$ and judgments $J \subseteq \mathfrak{X}$, and also between PB and JA profiles. Similarly, any JA outcome can be translated back into the PB setting.

Definition 6.1.2 (Outcome Translation). *Let $I = \langle \mathcal{P}, \mathcal{R}, c, \mathbf{b} \rangle$ be a PB instance and let $\Gamma \in \mathcal{L}_{\mathfrak{X}}$ be an integrity constraint expressed over the atoms $\mathfrak{X} = \{x_p \mid p \in \mathcal{P}\}$. The outcome translation $\tau : 2^{\mathfrak{X}} \rightarrow 2^{\mathcal{P}}$ maps any judgment $J \in 2^{\mathfrak{X}}$ to a budget allocation defined as:*

$$\tau(J) = \{p \in \mathcal{P} \mid x_p \in J\}.$$

We extend the outcome translation to sets $\mathcal{J} \subseteq 2^{\mathfrak{X}}$ of judgments by stipulating that $\tau(\mathcal{J}) = \{\tau(J) \mid J \in \mathcal{J}\}$.

We now define one of the fundamental elements of our approach: *embeddings*. An embedding is a function $E : \mathcal{I} \rightarrow \mathcal{L}_{\mathfrak{X}}$ that takes a PB instance as input and returns an integrity constraint (*i.e.*, a JA instance). Given an embedding, we can translate any input of a PB rule into an input for a JA rule, apply the JA rule, and finally translate the result obtained into a set of budget allocations (see Figure 6.1.1). Note that the whole process defines PB rules, through JA ones.

However, to be meaningful, the integrity constraint should express the budget constraint of the PB instance. This is captured by the notion of *correctness* that states that the outcome translation τ defines a bijection between the set of budget allocations on the PB side and the set of admissible judgments on the JA side.

Definition 6.1.3 (Correct Embedding). *An embedding $E : \mathcal{I} \rightarrow \mathcal{L}_{\mathfrak{X}}$ is said to be correct if, for every PB instance $I \in \mathcal{I}$, we have:*

$$\tau(\mathfrak{J}(E(I))) = \text{FEAS}(I).$$

In the following section we will study actual embeddings and show how they can be used to easily incorporate additional constraints in the study of PB.

We show that $\text{OUTCOMEDETERMINATION}(F)$ is solvable in polynomial time for any additive JA rule F for which the associated function f is *polynomial-time computable*.

We will need some notions of algebra in the proof. Let us provide some definitions here. A *semi-ring* $\langle A, \oplus, \otimes, e^\oplus, e^\otimes \rangle$ is an algebraic structure such that \oplus and \otimes are associative binary operations over A ; \oplus is commutative; e^\oplus is the identity element of \oplus and e^\otimes that of \otimes ; \otimes is left and right distributive over \oplus ; and finally $e^\oplus \otimes a = a \otimes e^\oplus = e^\oplus$ for any $a \in A$. A semi-ring is commutative if \otimes is commutative too.

Theorem 6.2.2. *Let F be an additive JA rule defined with respect to some polynomial-time computable function f . Then $\text{OUTCOMEDETERMINATION}(F)$ is polynomial-time solvable when the integrity constraint Γ in the input is represented as a DNNF circuit.*

PROOF. We show that when Γ is a DNNF circuit, we can use the *algebraic model counting* (AMC) problem to solve $\text{OUTCOMEDETERMINATION}(F)$. Given a propositional formula $\varphi \in \mathcal{L}_{\mathfrak{X}}$, a commutative semi-ring $\langle A, \oplus, \otimes, e^\oplus, e^\otimes \rangle$, and a labeling function $\lambda : \text{Lit}(\mathfrak{X}) \rightarrow A$, the AMC problem is to compute:

$$\text{AMC}(\varphi) = \bigoplus_{\substack{\alpha: \mathfrak{X} \rightarrow \{0,1\} \\ \alpha \models \varphi}} \bigotimes_{\substack{\ell \in \text{Lit}(\mathfrak{X}) \\ \alpha(\ell)=1}} \lambda(\ell).$$

The pair $\langle \oplus, \lambda \rangle$ is called *neutral* if and only for every propositional atom $x \in \mathfrak{X}$, $\lambda(x) \oplus \lambda(\neg x) = e^\oplus$. [Kimmig, Van den Broeck and De Raedt \(2017\)](#) proved that when φ is a DNNF circuit, \oplus is idempotent (for every $a \in A$, we have: $a \oplus a = a$), and $\langle \oplus, \lambda \rangle$ is neutral, then the AMC problem can be solved in polynomial time.

We now show that $\text{OUTCOMEDETERMINATION}(F)$ can be solved using the AMC problem when F is an additive rule. Let f be the function associated with F . We will consider AMC problem with the max-plus algebra—a commutative and idempotent semi-ring ([Akian, Bapat and Gaubert, 2006](#))—defined by $A = \mathbb{R} \cup \{-\infty, \infty\}$, $e^\oplus = -\infty$, and $e^\otimes = 0$, where \oplus and \otimes are the usual max and $+$ operators over $\mathbb{R} \cup \{-\infty, \infty\}$. Moreover, for a profile \mathbf{J} we introduce a labelling function $\lambda_{\mathbf{J}}(\cdot)$ defined for every literal $\ell_x \in \text{Lit}(\mathfrak{X})$ as:

$$\lambda_{\mathbf{J}}(\ell_x) = f(\mathbf{J}, \ell_x) - \max(f(\mathbf{J}, x), f(\mathbf{J}, \neg x)).$$

Since we have $\max(\lambda_{\mathbf{J}}(x), \lambda_{\mathbf{J}}(\neg x)) = 0$ for every $x \in \mathfrak{X}$, it is easy to see that the pair $\langle \lambda_{\mathbf{J}}, \oplus \rangle = \langle \lambda_{\mathbf{J}}, \max \rangle$ is neutral.

For every profile \mathbf{J} and labelling function $\lambda_{\mathbf{J}}$, we then have:

$$\begin{aligned} & \arg \max_{J \in \mathfrak{J}(\Gamma)} \sum_{\ell_x \in \text{aug}(J)} \lambda_{\mathbf{J}}(\ell_x) \\ &= \arg \max_{J \in \mathfrak{J}(\Gamma)} \left(\sum_{\ell_x \in \text{aug}(J)} \left(f(\mathbf{J}, \ell_x) - \max(f(\mathbf{J}, x), f(\mathbf{J}, \neg x)) \right) \right) \end{aligned} \quad (6.1)$$

$$= \arg \max_{J \in \mathfrak{J}(\Gamma)} \left(\sum_{\ell_x \in \text{aug}(J)} f(\mathbf{J}, \ell_x) - 2 \cdot \sum_{x \in \mathfrak{X}} \max(f(\mathbf{J}, x), f(\mathbf{J}, \neg x)) \right) \quad (6.2)$$

$$= \arg \max_{J \in \mathfrak{J}(\Gamma)} \sum_{\ell \in \text{aug}(J)} f(\mathbf{J}, \ell) \quad (6.3)$$

$$= F(\Gamma, \mathbf{J})$$

Let us briefly explain the above. The transition between lines (6.1) and (6.2) comes from the fact that $\text{aug}(J)$ include exactly one literal for each propositional atom in \mathfrak{X} . Observe then that the term $2 \cdot \sum_{x \in \mathfrak{X}} \max(f(\mathbf{J}, x), f(\mathbf{J}, \neg x))$ does not depend on the judgment J considered by the $\arg \max$ operator, and can thus be dropped, leading to line (6.3).

We can then solve $\text{OUTCOMEDETERMINATION}(F)$ by using the AMC problem. For Γ , \mathbf{J} and $L \subseteq \text{Lit}(\mathfrak{X})$ given as inputs of the $\text{OUTCOMEDETERMINATION}(F)$ problem, we will solve the AMC problem twice: first for $\varphi = \Gamma$ and then for $\varphi = \Gamma'$, where Γ' is obtained from Γ by fixing the value of the atoms as in L . If the solution of the AMC problem is the same in both cases, we answer the $\text{OUTCOMEDETERMINATION}(F)$ problem by the positive, and by the negative otherwise.

Overall, as the max-plus algebra is idempotent and $\langle \max, \lambda \rangle$ is neutral, the AMC problem can be solved in polynomial time when φ is a DNNF circuit (Kimig, Van den Broeck and De Raedt, 2017). Hence, our procedure to solve the $\text{OUTCOMEDETERMINATION}(F)$ problem also runs in polynomial time when Γ is a DNNF circuit. \square

Since all the rules we introduce in Section 6.1.2 are additive rules for which the corresponding function f is polynomial-time computable. Thus, Theorem 6.2.2 immediately implies tractability of the outcome determination problem for these rules.

Corollary 6.2.3. *When the integrity constraint is represented as a DNNF circuit, then the problem $\text{OUTCOMEDETERMINATION}(F)$ can be solved in polynomial time when F is either the Kemeny, the Slater, or the leximax rule.*

6.2.2 DNNF Circuit Embeddings

At this point, we know that we can efficiently compute the outcome of JA rules when the integrity constraint is represented as a DNNF circuit, but we still need to demonstrate that it is actually possible to encode PB problems as integrity constraints of this kind. So we move on to the description of embeddings of PB into JA returning integrity constraints represented as DNNF circuits. In doing so, we follow De Haan (2018) but use a slight generalisation of his approach, allowing us to deal with PB instances with multiple resources.

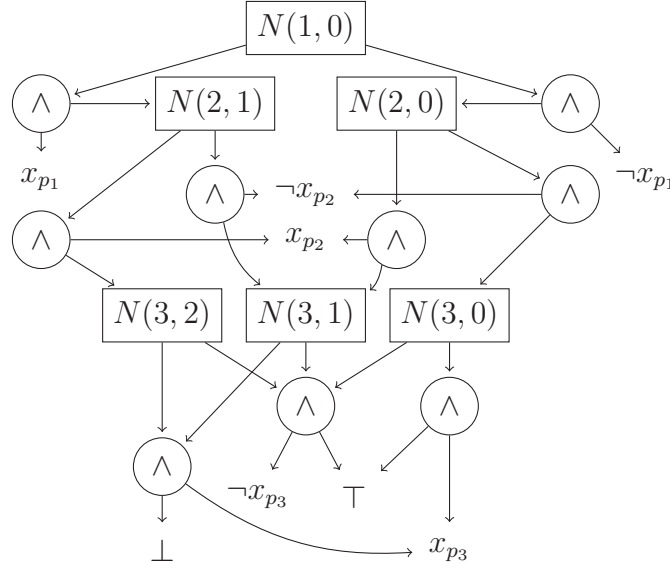


Figure 6.2.1: (Simplified) DNNF circuit produced by *TE* in Example 6.2.4.

Let us describe the basic construction. The idea is that every \vee -node in the DNNF circuit will represent the choice of selecting (or not) a given project. To know whether it is possible to select a given project or not, we keep track of the amount of resources that has been used so far. Selecting a project can thus only be done if it would not lead to a violation of the budget constraint.

For a project index j and a vector of used quantities per resources $\mathbf{v} \in \mathbb{R}_{\geq 0}^d$, we introduce the \vee -node $N(j, \mathbf{v})$, corresponding to the situation where we previously made a choice on projects with indices 1 to $j - 1$, and where for these choices we used resources according to \mathbf{v} . These nodes $N(j, \mathbf{v})$ are defined as follows:

$$N(j, \mathbf{v}) = \begin{cases} \top & \text{if } j = m + 1, \\ \vee \left(\begin{array}{l} (x_{p_j} \wedge N(j + 1, \mathbf{v} + c(p_j))) \\ (\neg x_{p_j} \wedge N(j + 1, \mathbf{v})) \end{array} \right) & \text{if } \mathbf{v} + c(p_j) \leq \mathbf{b}, \\ (\neg x_{p_j} \wedge N(j + 1, \mathbf{v})) \vee (x_{p_j} \wedge \perp) & \text{otherwise.} \end{cases}$$

For a PB instance $I = \langle \mathcal{P}, \mathcal{R}, c, \mathbf{b} \rangle$, we define the *tractable embedding* $TE(I)$ as the embedding that returns the integrity constraint defined by $N(1, \mathbf{0}_d)$, where $\mathbf{0}_d$ denotes the vector of length d whose components are all equal to 0.

Let us illustrate this embedding on a simple example.

Example 6.2.4. Consider an instance I with a single resource r and projects p_1 , p_2 , and p_3 . The costs of the projects in terms of r are $c(p_1) = c(p_2) = 1$ and $c(p_3) = 2$ and the budget limit is $b = 2$. Call x_{p_1} , x_{p_2} , and x_{p_3} the propositional atoms corresponding to p_1 , p_2 , and p_3 , respectively. *TE* on I would construct the DNNF circuit presented in Figure 6.2.1. Note that we simplified it a bit to improve its readability. \triangle

Providing an embedding is the first step, the next one is to show that it encodes PB instances correctly. We do so for TE in the following.

Proposition 6.2.5. *The tractable embedding TE is correct, and for any given PB instance $I = \langle \mathcal{P}, \mathcal{R}, c, \mathbf{b} \rangle$ returns an integrity constraint $TE(I)$ represented as a DNNF circuit of size in $\mathcal{O}(m \cdot |\{c(\pi) \mid \pi \subseteq \text{FEAS}(I)\}|)$.*

PROOF. Let $I = \langle \mathcal{P}, \mathcal{R}, c, \mathbf{b} \rangle$ be a PB instance, and Γ the integrity constraint returned by the tractable embedding, *i.e.*, $\Gamma = TE(I)$.

We first show that Γ is represented as DNNF circuit. First, observe that Γ is a Boolean circuit rooted in $N(1, \mathbf{0}_d)$. Next, it is clear that every \vee -node is of the form $(x \wedge \beta_1) \vee (\neg x \wedge \beta_2)$, where $x \in \mathfrak{X}$ is a propositional atom and β_1, β_2 are either \vee -nodes, or one of the logical constant \perp , or \top . This implies that Γ is represented as an NNF circuit. Because each project is only considered once, the propositional atom corresponding to the project cannot appear in two distinct conjuncts. Hence, Γ is a DNNF circuit.

Observe that there are at most $m \cdot |\{c(\pi) \mid \pi \subseteq \text{FEAS}(I)\}|$ many \vee -nodes in Γ —one for each $N(j, v)$ for which the budget is not exceeded—all of them having at most two child \wedge -nodes. There are moreover $2m + 2$ leaves, one per literal and two for \perp and \top , hence the size of the DNNF circuit.

We now show that the tractable embedding is correct. Observe that a branch leading to the \perp -leaf is chosen if and only if one would violate the budget limit by selecting a project p_j . Hence, finding an assignment that does not lead to a \perp leaf in Γ can only be done by selecting a feasible set of projects. The set of such assignments defines the set of outcomes satisfying Γ , so $\tau(E(I)) \subseteq \text{FEAS}(I)$. Now, consider $\pi \in \text{FEAS}(I)$. Since π is feasible, it is clear that there exists a branch in the DNNF circuit Γ along which the selected projects correspond exactly to those that are in π . We thus have $\tau(E(I)) = \text{FEAS}(I)$. \square

At this point, it should be noted that the exponential factor in the size of the produced DNNF circuit, namely $|\{c(\pi) \mid \pi \subseteq \text{FEAS}(I)\}|$, is bounded from above by the product of the budget limits for each resource *i.e.*, it is in $\mathcal{O}(m \cdot \prod_{r \in \mathcal{R}} b_r)$. This is thus pseudo-polynomial in the size of the PB instance when the number of resources is fixed.

The next natural question then is whether we can do better. For instance, is it possible to reduce the size to something pseudo-polynomial in the size of the PB instance regardless of the number of resources, *i.e.*, in $\mathcal{O}(m \cdot \sum_{r \in \mathcal{R}} b_r)$? The following result answers this question in the negative, or at least, shows that such a DNNF circuit cannot be *found* in time polynomial in $m + \sum_{r \in \mathcal{R}} b_r$ (but may exist).

Proposition 6.2.6. *There exists no embedding into a DNNF circuit that can be computed in time polynomial in $m + \sum_{r \in \mathcal{R}} b_r$ for any instance $I = \langle \mathcal{P}, \mathcal{R}, c, \mathbf{b} \rangle$, unless $\text{P} = \text{NP}$.*

PROOF. We first show that the following problem is strongly NP-complete.

MAXIMAL EXHAUSTIVE ALLOCATION

Input: A PB instance I and a natural number $k \in \mathbb{N}$.

Question: Is there a budget allocation $\pi \in \text{FEAS}_{\text{EX}}(I)$ such that $|\pi| \geq k$?

First, note that MAXIMAL EXHAUSTIVE ALLOCATION obviously is in NP, the certificate being a budget allocation of size at least k .

We now show that MAXIMAL EXHAUSTIVE ALLOCATION is strongly NP-hard. To do so we reduce from the 3-DIMENSIONAL MATCHING problem, which was shown to be NP-complete by Karp (1972). Note that, since its input does not involve numbers, the 3-DIMENSIONAL MATCHING problem is also strongly NP-complete.

3-DIMENSIONAL MATCHING

Input: A finite set T and a set $X \subseteq T \times T \times T$.

Question: Is there a set $M \subseteq X$ such that $|M| = |X|$, and for all (x_1, x_2, x_3) and (y_1, y_2, y_3) in M , we have $x_1 \neq y_1$ and $x_2 \neq y_2$ and $x_3 \neq y_3$?

Consider, without loss of generality, an instance of the 3-DIMENSIONAL MATCHING problem $\langle T, X \rangle$ such that $T = \{1, \dots, t\}$ and $X = \{x_1, \dots, x_{|X|}\}$. The corresponding PB instance is $I = \langle \mathcal{P}, \mathcal{R}, c, \mathbf{b} \rangle$ where the set of resources is $\mathcal{R} = \{r_i^j \mid i \in T, j \in \{1, 2, 3\}\}$ and for every resource $r \in \mathcal{R}$, we have $b_r = 1$. The set of projects is $\mathcal{P} = \{p_1, \dots, p_{|X|}\}$. Consider project $p_i \in \mathcal{P}$ corresponding to $x_i = (x_i^1, x_i^2, x_i^3) \in X$, its cost is 1 for the three resources $r_{x_i^1}^1, r_{x_i^2}^2$ and $r_{x_i^3}^3$ and 0 for any other resource. Moreover we set $k = |X|$. We claim that the answer for the 3-DIMENSIONAL MATCHING problem on $\langle T, X \rangle$ is yes if and only if the answer for the MAXIMAL EXHAUSTIVE ALLOCATION problem on $\langle I, k \rangle$ is yes too.

To a matching $M \subseteq X$, corresponds the budget allocation $\pi = \{p_i \mid x_i \in M\}$. A matching $M \subseteq X$ is a solution of the 3-DIMENSIONAL MATCHING problem if and only if no triplet in M share a coordinate. Because of the budget limit, this is possible if and only if the corresponding budget allocation π is feasible. Moreover $|M| = |X|$ if and only if $|A| = |X| = k$. Note that in this case π would be exhaustive, which proves the claim. The reduction is clearly done in polynomial time which shows that the MAXIMAL EXHAUSTIVE ALLOCATION problem is strongly NP-complete.

To conclude the proof, we now show that if there exists an embedding of PB into a DNNF circuit that runs in time polynomial in $m + \sum_{r \in \mathcal{R}} b_r$, then we would be able to solve the MAXIMAL EXHAUSTIVE ALLOCATION problem in pseudo-polynomial time. This would imply that $P = NP$ as MAXIMAL EXHAUSTIVE ALLOCATION is strongly NP-complete.

Let Γ be the integrity constraint returned by a suitable exhaustive embedding on an arbitrary instance I . Note that the answer of the MAXIMAL EXHAUSTIVE ALLOCATION problem is yes if and only if the outcome of the ALGEBRAIC MODEL COUNTING (AMC) problem is at least $k - m$ when run on Γ with the max-plus algebra (see the proof of Theorem 6.2.2 for the relevant definitions) and the labeling function λ such that $\lambda(x) = 0$ and $\lambda(\neg x) = -1$ for all $x \in \mathfrak{X}$. Since Γ is a DNNF circuit and the pair $\langle \lambda, \oplus \rangle$ is neutral, we can compute the outcome of the AMC problem in time polynomial in the size of Γ . This proves the claim. \square

Even though there is no hope to find pseudo-polynomial embeddings when the number of resources is unbounded, we still argue that the embedding is efficient for realistic scenarios. First, in most typical PB processes the number of resources is expected to be small. Indeed the difficulty of assessing the different costs and of the deliberation and voting processes increases significantly with the number of dimensions. It thus seems particularly unlikely that the cost will be expressed in more than, say, five dimensions. Moreover, if the number of resources is fixed, or at least bounded, the size of the DNNF circuit will not really present a serious limitation with state-of-the-art solvers.

In the remainder of this section we investigate to what extent this approach allows us to introduce additional distributional constraints for PB.

6.2.3 Participatory Budgeting with Project Dependencies

We now consider the situation where the completion of some projects is directly dependent on the completion of some others. The idea here is to incorporate in the formal model statements such as “constructing a bike shed only makes sense if the project of the bicycle lane also is implemented”, or “one cannot build a fountain in the middle of the park at the same time as some children amusement facilities”.

Let us formally introduce what we mean here. Take a PB instance $I = \langle \mathcal{P}, \mathcal{R}, c, \mathbf{b} \rangle$. We introduce a set of implications, $Imp \subseteq \mathcal{L}_{\mathfrak{X}}$, linking projects together. A set of implications is a set of propositional formulas of the form $\ell_{x_p} \rightarrow \ell_{x_{p'}}$, where p and p' are two projects in \mathcal{P} , and ℓ_{x_p} and $\ell_{x_{p'}}$ the corresponding literals. Note that this corresponds to 2-CNF formulas.⁴⁵ In the case that ℓ_{x_p} is positive (respectively negative), such an implication indicates that p can be selected (respectively not selected) only if p' is selected, when $\ell_{x_{p'}}$ is positive, or not selected, when $\ell_{x_{p'}}$ is negative. A budget allocation π satisfies the set of implications Imp if and only if the previously described semantics is satisfied. For an instance I , we denote by $FEAS(I, Imp)$ the set of feasible

⁴⁵A propositional logic formula $\varphi \in \mathcal{L}_{\mathfrak{X}}$ is in Conjunctive Normal Form (CNF) if it is expressed as a conjunction of clauses, where a clause is a disjunction of literals. It is moreover a 2-CNF formula if all clauses are of size 2, i.e., if φ is a conjunction of disjunctions over two literals.

budget allocations satisfying Imp . Moreover, we will write $\ell_{x_p} \rightarrow^* \ell_{x_{p'}}$ if there is a chain of implication in Imp linking ℓ_{x_p} to $\ell_{x_{p'}}$.

In terms of applications, this approach to model dependencies is quite flexible. It allows us to model the fact that project p_1 can only be implemented if projects p_2 and p_3 would also be implemented for instance. This would be encoded as:

$$Imp = \{x_{p_1} \rightarrow x_{p_2}, x_{p_1} \rightarrow x_{p_3}\}.$$

At the same time, we can also model “negative” dependencies, or “incompatibilities”, where p_1 can only be implemented if p_2 is not: $Imp = \{x_{p_1} \rightarrow \neg x_{p_2}\}$. Note that we are not aware of any previous work that would allow for such constraint to be expressed in the PB framework.

The first step in our study of PB with project dependencies is to show that finding a feasible budget allocation when there are implications between project is an NP-complete problem.

Proposition 6.2.7. *Let $I = \langle \mathcal{P}, \mathcal{R}, c, \mathbf{b} \rangle$ be a PB instance and Imp a set of implications over I . Deciding whether $FEAS(I, Imp)$ is empty or not is an NP-complete problem, even for the case of a single resource and with unit costs.*

PROOF. The problem of finding a feasible budget allocation satisfying Imp is clearly in NP. Indeed, checking that the budget limit is not exceeded can be done by summing the costs of the selected projects. Moreover, verifying that the set of implications is satisfied simply amounts at checking the truth value of each implication in Imp . Both of these problems can be solved in polynomial time.

To show that the problem is NP-hard, we reduce from the NP-complete problem 2-CNF MINIMAL MODEL (Ben-Eliyahu and Dechter, 1996).

2-CNF MINIMAL MODEL

Input: A formula $\varphi \in \mathcal{L}_{\mathfrak{X}}$ in conjunctive normal form with exactly two literals per clauses, and $k \in \mathbb{N}$.
Question: Is there a model α such that $\alpha \models \varphi$ and $|\{p \in \mathfrak{X} \mid \alpha(p) = 1\}| \leq k$?

Take an instance $\langle \varphi, k \rangle$ of the 2-CNF MINIMAL MODEL problem. We construct the following participatory budgeting instance I . The set of resources is $\mathcal{R} = \{r\}$ with budget limit $b_r = k$. There is one project per propositional atom in φ , i.e., $\mathcal{P} = \{p_x \mid x \in P\}$. For every project $p \in \mathcal{P}$, we have $c(p) = 1$. Finally, the set of implications Imp is the set of clauses in φ (remember that an implication between two variable can be equivalently expressed as a clause with two literals, and reciprocally).

We claim that there exists a model of φ setting no more than k variables to true if and only if there exists a feasible budget allocation for I that satisfies the

set of implications Imp . Indeed, to a given truth assignment α , corresponds the budget allocation $\pi = \{p_x \mid \alpha(x) = 1\}$. Observe first that there exists $\alpha \models \varphi$ if and only if π satisfies Imp . Moreover, since every project is of cost 1, the cost of π is exactly the number of project that are selected. As the set of projects is the set of variables in φ , the cost of π is also equal to the number of variables set to true in α . Because the budget limit for r is k , the budget allocation π satisfies the budget limit if and only no more than k propositional atoms are set to true in α .

Observing that this reduction clearly can be done in polynomial time concludes the proof. \square

Because of this result we cannot hope to find an efficient embedding into a DNNF circuit, even with a small number of resources. However, we can still define an interesting embedding whose size is parameterized by some parameter on the structure of the implication set (in the spirit of parameterized complexity, [Downey and Fellows, 2013](#)). The parameter in question is the pathwidth ([Robertson and Seymour, 1983](#); [Bodlaender, 1998](#)) of the interconnection graph. Let us define these two terms.

First, we introduce the *interconnection graph* G of a set of implications Imp . It is the graph $G = \langle \mathcal{P}, E \rangle$ where there is an edge $\{p, p'\} \in E$ between projects p and p' if and only if there exists an implication in Imp linking the two projects, *i.e.*, Imp includes at least one implication of the form $\ell_{x_p} \rightarrow \ell_{x_{p'}}$ for some $\ell_{x_p} \in \{x_p, \neg x_p\}$, $\ell_{x_{p'}} \in \{x_{p'}, \neg x_{p'}\}$.

Second, let us discuss the *pathwidth* of a graph $G = \langle V, E \rangle$. A path-decomposition of a graph G is a vector of subset of vertices (V_1, \dots, V_q) , called *bags*, such that:

- (i) for every edge $(v_1, v_2) \in E$, there is a bag V_i such that v_1 and v_2 are in V_i ;
- (ii) for every $i \leq j \leq k$, we have $V_i \cap V_k \subseteq V_j$.

The second property should be understood as saying that the set of bags in which a given vertex appears is contiguous. The *width* of a tree decomposition is the size of the largest bag minus one. The *pathwidth* of a graph G is the minimum width of any of its path-decomposition. Interestingly, given a graph $G = \langle V, E \rangle$, we can compute an optimal path-decomposition of G in FPT-time, where the parameter is the pathwidth itself ([Bodlaender and Kloks, 1996](#)).

We are now equipped with all the definitions we need to formulate our embedding for dependencies, denoted by TE_{dep} .

Theorem 6.2.8. *Let $I = \langle \mathcal{P}, \mathcal{R}, c, \mathbf{b} \rangle$ be a PB instance and Imp a set of implications over I . Then, there exists a correct embedding from I and Imp to an integrity constraint expressed as a DNNF circuit Γ with size in $\mathcal{O}(m \cdot |\{c(\pi) \mid \pi \subseteq \text{FEAS}(I, Imp)\}| \cdot 2^k)$, where k is the pathwidth of the interconnection graph of Imp .*

PROOF. The proof will be presented in several steps. We will first present our embedding TE_{dep} , then investigate the size of the integrity constraint returned and finally show that the embedding is correct. In the following we assume that \mathfrak{X} is exactly the set $\{x_p \mid p \in \mathcal{P}\}$.

Let $G = \langle \mathcal{P}, E \rangle$ be the interconnection graph of Imp . We order the projects in the same order in which they are introduced in an optimal path-decomposition of G , with p_1 being the first project, p_2 the second and so forth. Then, following the idea developed for TE , we introduce \vee -nodes $N(j, \mathbf{v}, L)$ where j is a project index, $\mathbf{v} \in \mathbb{R}_{\geq 0}^d$ a vector of used quantities per resource and $L \subseteq Lit(\mathfrak{X})$ a subset of literals. Intuitively, the set L specifies the literals that we selected and that we should remember because they might trigger implications later on. The nodes $N(j, \mathbf{v}, L)$ are then defined according to the following cases (where each \blacktriangleright should be understood as an “otherwise”):

- \blacktriangleright If $j = m + 1$, then $N(j, \mathbf{v}, L) = \top$;
- \blacktriangleright If both the positive literal x_{p_j} and the negative literal $\neg x_{p_j}$ are implied by some literal in L according to Imp , then $N(j, \mathbf{v}, L) = \perp$;
- \blacktriangleright If the positive literal x_{p_j} is implied by some literal in L according to Imp , and $\mathbf{v} + c(p_j) \leq \mathbf{b}$, then $N(j, \mathbf{v}, L) = N(j + 1, \mathbf{v} + c(p_j), L \cup \{x_{p_j}\})$;
- \blacktriangleright If the positive literal x_{p_j} is implied by some literal in L according to Imp , but there exists a resource $r_q \in \mathcal{R}$ such that $v_q + c(p_j, r_q) > b_q$, then $N(j, \mathbf{v}, L) = \perp$;
- \blacktriangleright If the negative literal $\neg x_{p_j}$ is implied by some literal in L according to Imp , then $N(j, \mathbf{v}, L) = N(j + 1, \mathbf{v}, L \cup \{\neg x_{p_j}\})$;
- \blacktriangleright If $\mathbf{v} + c(p_j) \leq \mathbf{b}$, then:

$$N(j, \mathbf{v}, L) = \vee \begin{array}{l} (x_{p_j} \wedge N(j + 1, \mathbf{v} + c(p_j), L \cup \{x_{p_j}\})) \\ (\neg x_{p_j} \wedge N(j + 1, \mathbf{v}, L \cup \{\neg x_{p_j}\})) \end{array} ;$$

- \blacktriangleright In all other cases, $N(j, \mathbf{v}, L) = (x_{p_j} \wedge \perp) \vee (\neg x_{p_j} \wedge N(j + 1, \mathbf{v}, L \cup \{\neg x_{p_j}\}))$.

The tractable embedding with dependencies, written TE_{dep} , refers to the integrity constraint defined by $N(1, \mathbf{0}_m, \emptyset)$.

It is clear that this construction ensures that the integrity constraint returned by TE_{dep} is represented as a DNNF circuit. The proof is very similar to the one for the tractable embedding (Proposition 6.2.5).

We now investigate the maximum size of the DNNF circuit. We need to count the maximum number of \vee -nodes, that is the number of $N(j, \mathbf{v}, L)$. At a first

glance, the number of possible $L \subseteq \text{Lit}(\mathfrak{X})$ is upper-bounded by $2^{|\mathcal{P}|}$. However, we can have a more fine-grained analysis of this last term. The set L is used to keep track of the projects for which a truth value has already been assigned and that could imply the truth value of some other project appearing later in the ordering. However, since the projects are considered following an optimal path-decomposition of the interconnection graph, we know that we never need to remember the truth value of more than $k+1$ projects, where k is the pathwidth of the interconnection graph. Indeed, by definition of a path-decomposition, for any project p_j , whenever we consider another project $p_{j'}$ such that $p_{j'}$ never appears in a bag together with p_j , we no longer need to keep track of the truth value associated with x_{p_j} as p_j will never be involved in implications with any subsequent projects.

Overall, we can update the definition of the nodes $N(j, \mathbf{v}, L)$ to implement this “forgetting” operation. We chose not to over-complicate a chapter that is already rather technical, but it should be clear from these explanations how to do so. The size of the DNNF circuit would then be in $\mathcal{O}(m \cdot |\{c(P) \mid P \subseteq \mathcal{P}\}| \cdot 2^k)$.

Finally we show that TE_{dep} is correct. Remember that this is the case if for every instance I :

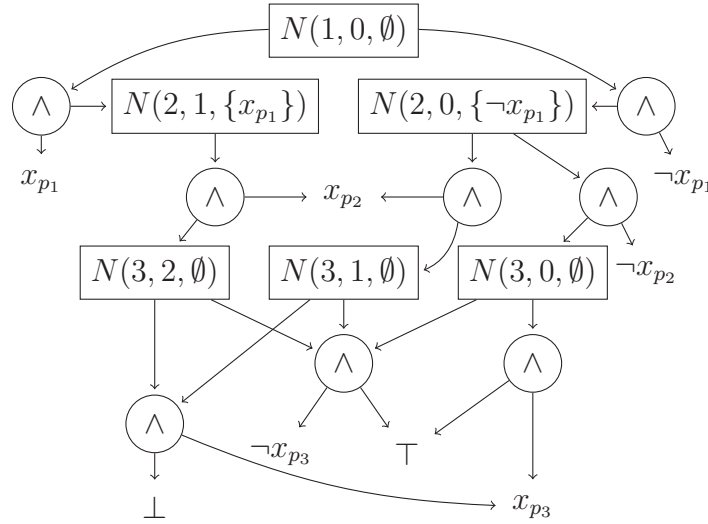
$$\tau(\mathfrak{J}(TE_{dep}(I, Imp))) = \text{FEAS}(I, Imp).$$

Consider an outcome $J \in \mathfrak{J}(TE_{dep}(I, Imp))$ on the JA side. It is clear that $\tau(J)$ does not exceed the budget limit as every time where selecting a project could lead to a too high cost, the branch in the DNNF circuit ends up in the \perp leaf. We then need to prove that $\tau(J)$ satisfies Imp . Observe that every time a literal is implied by a literal that has been previously given a truth value, we follow the implication. Hence it can never be the case that the premise of an implication is set to true but not the conclusion. Moreover, every time triggering implications would lead to an inconsistent outcome (a project being both selected and not selected), the branch in the DNNF circuit also leads to the \perp leaf. Overall $\tau(J)$ satisfies Imp . We have thus proved that $\{\tau(J) \mid J \in \mathfrak{J}(TE_{dep}(I, Imp))\} \subseteq \text{FEAS}(I, Imp)$. The proof for the reversed inclusion is exactly as that presented in Proposition 6.2.5. \square

The size of the DNNF circuit produced by the embedding includes a factor 2^k where k is the pathwidth of the interconnection graph of Imp . The value of k is in general upper bounded by the number of projects, for instance in the case where the interconnection graph is complete (when there are dependencies between any two projects). Once again it seems rather fair to assume projects not to be very interconnected, leading to small values of k in practice.

Let us illustrate this embedding on an example.

Example 6.2.9. Consider the instance described in Example 6.2.4. Assume that, if project p_1 is selected, then also project p_2 should be selected. In our model, this means that $Imp = \{x_{p_1} \rightarrow x_{p_2}\}$. An optimal path-decomposition of the interconnection graph is thus $(\{p_1, p_2\}, \{p_3\})$. So we will consider the projects in the following order: p_1, p_2, p_3 . The embedding TE_{dep} would return the DNNF circuit presented in below.



It is important to understand how we can “forget” about the truth values of x_{p_1} and x_{p_2} once we consider project p_3 . \triangle

We conclude this section with a small detour to the field of Knowledge Compilation (see, e.g., [Marquis, 2015](#)). Embedding PB problems with dependencies into a DNNF circuit is equivalent to asking whether the conjunction of a 2-CNF formula with a DNNF circuit can be efficiently encoded as a DNNF circuit of polynomial size. It turns out to be impossible to do so, unless the polynomial hierarchy collapses.

Proposition 6.2.10. *It is not possible to compile any 2-CNF formula $\varphi \in \mathcal{L}_x$ into a DNNF circuit of size polynomial in the size of φ , unless the polynomial hierarchy collapses to the second level.*

PROOF. We will prove that if 2-CNF formulas can be compiled into polynomial-size DNNF circuits, then the NP-complete problem CLIQUE would be in P/poly, using similar techniques as [Cadoli, Donini, Liberatore and Schaerf \(2002\)](#). Because of the Karp-Lipton theorem ([Karp and Lipton, 1980](#)), this would immediately entail the polynomial hierarchy to collapse at the second level.

Let us first introduce the problem CLIQUE, shown to be NP-complete by [Karp \(1972\)](#). Given an undirected graph $G = \langle V, E \rangle$, we say that a subset of vertices $V' \subseteq V$ forms a clique in G if for every $v_1, v_2 \in V'$, we have $\{v_1, v_2\} \in E$. The decision problem CLIQUE is then defined as follows.

 CLIQUE

Input: An undirected graph $G = \langle V, E \rangle$ and an integer $k \in \mathbb{N}$.

Question: Is there a clique $V' \subseteq V$ in G such that $|V'| = k$?

For any two integers $n, k \in \mathbb{N}$, we introduce a 2-CNF formula $\varphi(n, k)$ such that the answer of CLIQUE on a graph $G = \langle V, E \rangle$ with $|V| = n$ and k can be deduced by means of queries of the following problem the we call MAX SAT EXTENSION: Given a partial truth assignment α of φ , what is the maximum number of variables that can be set to true in any extension of α in such a way that α satisfies $\varphi(n, k)$. Importantly, the formula $\varphi(n, k)$ is the same for all graphs with n vertices and inquired clique size k . In the following we will assume that the vertices in V are named v_1, \dots, v_n .

Let us describe how to construct $\varphi(n, k)$. For every $i, i' \in \{1, \dots, n\}$ and each $1 \leq j < j' \leq k$, we introduce the variable $x_{j,j',i,i'}$. These variables indicate which vertex was selected at a given position in an arbitrary ordering of the clique. So having $x_{j,j',i,i'}$ set to true means that in the ordering of the clique in which vertex v_i is at position j , then vertex $v_{i'}$ is at position j' . Now for this to be correct, we need to enforce that no two vertices can be selected at the same position. For every $j_1, j_2, j_3, j_4 \in \{1, \dots, k\}$ such that $j_1 < j_2$, $j_1 < j_3$ and $j_4 < j_2$, and for every $i_1, i_2, i_3, i_4 \in \{1, \dots, n\}$, we thus add the following two clauses to the formula $\varphi(n, k)$:

$$\begin{aligned} & (\neg x_{j_1, j_2, i_1, i_2} \vee \neg x_{j_1, j_3, i_3, i_4}) \\ & (\neg x_{j_1, j_2, i_1, i_2} \vee \neg x_{j_4, j_2, i_3, i_4}) \end{aligned}$$

Now, we want to check whether any given graph $G = \langle V, E \rangle$ with $|V| = n$ has a clique of size k . Consider the formula $\varphi(n, k)$. We first create a partial truth assignment α that sets $x_{j,j',i,i'}$ and $x_{j,j',i',i}$ to false for all $1 \leq j < j' \leq k$ and all $i, i' \in \{1, \dots, n\}$ such that $\{v_i, v_{i'}\} \notin E$. Now it is easy to see that G has a clique of size k if and only if we can extend this α to a (non-partial) truth assignment that satisfies $\varphi(n, k)$ and that sets at least $k(k-1)/2$ variables to true.

At this point, it is important to observe that the problem MAX SAT EXTENSION can be solved in polynomial time if the formula provided is a DNNF circuit. It is a simple variant of the MAXIMUM MODEL problem that we can solve via dynamic programming on DNNF circuits (Darwiche and Marquis, 2002).

Now, suppose we can compile 2-CNF formulas into DNNF circuits in polynomial-space, then we compile $\varphi(n, k)$ into a DNNF for each n and each k . Call this DNNF circuit $D(n, k)$. From the above, it should be clear now that with $D(n, k)$ we can solve the CLIQUE problem on a graph G with n nodes and an inquired clique size of k in polynomial time.

To conclude the proof we show that this would then all mean that CLIQUE is in P/poly. We first informally define P/poly, we refer the reader to Chapter 6

in the book by [Arora and Barak \(2009\)](#) for a formal definition. P/poly is a complexity class that contains problems that can be solved in polynomial time by a deterministic Turing machine that has access to one advice string per size of the input. In our case, the advice string for an input of size n will be the sequence $(D(n, 1), \dots, D(n, n))$. Note that the size of the advice string is polynomial in n as each DNNF circuit in it is. For any input of CLIQUE, we can then find the corresponding DNNF circuit $D(n, k)$ to solve the problem in polynomial time. This proves that CLIQUE would be in P/poly, and so that $\text{NP} \subseteq \text{P/poly}$ (since CLIQUE is NP-complete). The Karp-Lipton theorem ([Karp and Lipton, 1980](#)) then implies that the polynomial hierarchy would collapse to the second level. \square

We already knew that there was little hope for a DNNF circuit encoding of PB instances with dependencies in polynomial-time (unless $\text{P} = \text{NP}$) when the number of resources is small. Under a stronger computational complexity assumption, we now have gone further by showing that it is also unlikely that we will not be able to do this in polynomial space.

6.2.4 Participatory Budgeting with Quotas on Types of Projects

The second extension of PB we consider is when projects are grouped into categories, or types, that are constrained by quotas.⁴⁶ The idea is that the projects belong to various types (health, education, environment to name a few) and that certain quotas over these types are to be respected by the final budget allocation (at least two health-related projects must be funded, for instance).

We start by presenting a model of PB with quotas over types of projects. We first introduce the notion of a type system. For a given PB instance $I = \langle \mathcal{P}, \mathcal{R}, c, \mathbf{b} \rangle$, a *type system* is a tuple $\mathfrak{T} = \langle \mathcal{T}, \mathcal{Q}, q, f \rangle$, where:

- ▶ $\mathcal{T} \in 2^{2^{\mathcal{P}}}$ is a set of types, each type being a subset of projects⁴⁷;
- ▶ $\mathcal{Q} = \langle Q, \oplus, e^{\oplus}, \leq_Q \rangle$ is an ordered group⁴⁸ over which the quotas are expressed;

⁴⁶After our initial work on this topic ([Rey, Endriss and de Haan, 2020](#)), [Jain, Sornat, Talmon and Zehavi \(2021\)](#) studied a specific subcase of our model, namely PB instances in which projects are grouped into categories and quotas regarding the total cost of projects from within each categories need to be satisfied (see Section 3.6.2). There is a small overlap between this chapter and the work of [Jain, Sornat, Talmon and Zehavi \(2021\)](#): the result of Proposition 6.2.11 is directly implied by Theorem 10 of [Jain, Sornat, Talmon and Zehavi \(2021\)](#). Other results are incomparable and they complement each other nicely.

⁴⁷Note that in Chapter 7 we will use the term “type” to refer to groups of agents, instead of projects.

⁴⁸A group $\langle Q, \oplus, e^{\oplus} \rangle$ is an algebraic structure equipped with a binary operation \oplus over Q that is associative, that has an identity element e^{\oplus} , and such that for every $a \in Q$, there exists a unique $b \in Q$ such that $a \oplus b = e^{\oplus}$ and $b \oplus a = e^{\oplus}$. An ordered group $\langle Q, \oplus, e^{\oplus}, \leq_Q \rangle$ is a group $\langle Q, \oplus, e^{\oplus} \rangle$ equipped with a total order \leq_Q over Q .

- ▶ $q : \mathcal{T} \rightarrow Q^2$ is a quota function such that for any type $T \in \mathcal{T}$, we have $q(T) = (a, b)$ with $a, b \in Q$ satisfying $a \leq_Q b$;
- ▶ $f : \mathcal{T} \times \text{FEAS}(I) \rightarrow Q$ is a type aggregator.

For any $T \in \mathcal{T}$ for which $q(T) = (a, b)$, we write $q(T)^- = a$ and $q(T)^+ = b$. They indicate the lower and upper quota for type T respectively. A budget allocation π satisfies the type system $\mathfrak{T} = \langle \mathcal{T}, Q, q, f \rangle$ if the quotas are respected, *i.e.*, if for every type $T \in \mathcal{T}$, we have:

$$q(T)^- \leq_Q f(T, \pi) \leq_Q q(T)^+.$$

We denote by $\text{FEAS}(I, \mathfrak{T})$ the set of feasible budget allocation that satisfy the type system \mathfrak{T} for instance I .

The type aggregator $f(\cdot)$ can be defined in several different ways. We provide two type aggregators that are very natural.

- ▶ **Cardinality-type aggregator.** The quotas express lower and upper bounds on the number of projects selected for each type. We have $Q = \mathbb{N}$, \oplus is the usual addition operator on numbers with identity element 0, and \leq_Q is the usual linear order on numbers. Now for every type $T \in \mathcal{T}$ and budget allocation π , the cardinality-type aggregator is defined as follows:

$$f^{\text{card}}(T, \pi) = |T \cap \pi|.$$

- ▶ **Cost-type aggregator.** The quotas define lower and upper bounds on the total cost of the projects selected for each type. Here $Q = \mathbb{R}_{\geq 0}^d$, \oplus is the component-wise addition over vectors with identity element $\mathbf{0}_d$, and \leq_Q is the component-wise order over vectors. Now for every type $T \in \mathcal{T}$ and budget allocation π , the cost-type aggregator is defined as follows:

$$f^{\text{cost}}(T, \pi) = \sum_{p \in \pi \cap T} c(p).$$

Turning to the formal analysis of PB with quotas over types of projects, we first show that deciding whether there is a feasible budget allocation satisfying given type system is NP-complete, for both the cardinality- and the cost-type aggregator.

Proposition 6.2.11. *Let $I = \langle \mathcal{P}, \mathcal{R}, c, \mathbf{b} \rangle$ be a PB instance and $\mathfrak{T} = \langle \mathcal{T}, Q, q, f \rangle$ a type system over I . Deciding whether $\text{FEAS}_{\text{EX}}(I, \mathfrak{T})$ is empty or not is an NP-complete problem when f is either the cardinality or the cost-type aggregator, even for the case of a single resource and with unit costs.*

PROOF. The problem is clearly in NP. Indeed, verifying that the budget allocation π does not exceed the budget limit can be done easily by summing the costs of the selected projects. Note that both the cardinality- and the cost-type aggregators can be computed in polynomial time. Then, verifying that every quota is respected is just a matter of scanning π for each quota.

Let us show now that the problem is NP-hard. To do so we reduce from the NP-hard problem SET SPLITTING (Garey and Johnson, 1979) described below.

SET SPLITTING

Input: A collection C of subsets of a given set S .
Question: Are there two sets S_1 and S_2 partitioning S such that for any $c \in C$, we have $c \not\subseteq S_1$ and $c \not\subseteq S_2$?

Let $\langle C, S \rangle$ be an instance of SET SPLITTING. We construct a participatory budgeting instance $I = \langle \mathcal{P}, \mathcal{R}, c, \mathbf{b} \rangle$ such that $\mathcal{R} = \{r\}$, and $b_r = |S|$. There is one project per element in S , i.e., $\mathcal{P} = \{p_s \mid s \in S\}$, and $c(p_s) = 1$ for every $s \in S$. Thus, the budget limit can never be exceeded. The corresponding set of types is $\mathcal{T} = \{\{p_s \mid s \in c\} \mid c \in C\}$ so that there is one type for subset in C . For a given type $T \in \mathcal{T}$, the quota is $q(T) = (1, |T| - 1)$. With one resource and projects whose costs are in $\{0, 1\}$, the cardinality-type aggregator and the cost-type aggregator coincide.

We claim that $\langle C, S \rangle$ is a yes-instance of SET SPLITTING if and only if there exists a feasible budget allocation in the instance I with the previous type system. For a given partition of S , (S_1, S_2) , a suitable corresponding budget allocation is $\pi = S_1$ (or equivalently $\pi = S_2$).

A partition (S_1, S_2) is a solution of the SET SPLITTING problem if and only if for every $c \in C$, at least one element of c is in S_1 and at least one element of c is not in S_1 (and is then in S_2). Based on the type system we defined, this is equivalent to π satisfying the quota associated to c . Moreover, observe that every budget allocation respects the budget limit. Hence (S_1, S_2) is a solution of the SET SPLITTING problem if and only if π is a feasible budget allocation.

The reduction is clearly done in polynomial time, hence the problem of finding whether a feasible budget allocation exists is NP-complete when using both the cost and the cardinality type aggregator. \square

Once again, this implies that no efficient embedding can be defined for this extension, even in the case of a small number of resources.

Given this computational impossibility, we present a parameterized embedding for PB with types and quotas in the following. The embedding works for any *additive* type aggregator $f : \mathcal{T} \times \text{FEAS}(I) \rightarrow Q$, that is, any type aggregator f for which there exists a *score type function* s that takes as input a project $p \in \mathcal{P}$ and returns an

element in Q such that for every type $T \in \mathcal{T}$ and every allocation $\pi \in \text{FEAS}(I)$:

$$f(T, \pi) = \bigoplus_{p \in A} s(p).$$

Note that both the cardinality- and the cost-type aggregators are additive, with the following score type function: $s^{\text{card}}(p) = 1$ and $s^{\text{cost}}(p) = c(p)$ respectively.

As for the dependencies case, the size of our embedding will be parameterized by the pathwidth of a graph. This time, it will be the *overlap graph* of a type system. Let I be an instance and $\langle \mathcal{T}, \mathcal{Q}, q, f \rangle$ a type system over I , the overlap graph of the type system is the graph $G = \langle \mathcal{T}, E \rangle$, where there is an edge $\{T, T'\} \in E$ between types T and T' if and only if $T \cap T' \neq \emptyset$, i.e., T and T' overlap.

Theorem 6.2.12. *Let $I = \langle \mathcal{P}, \mathcal{R}, c, \mathbf{b} \rangle$ be a PB instance and $\mathfrak{T} = \langle \mathcal{T}, \mathcal{Q}, q, f \rangle$ a type system where f is an additive type aggregator defined with respect to the function s . Then, there exists a correct embedding for I and $\langle \mathcal{T}, \mathcal{Q}, q, f \rangle$ that returns an integrity constraint represented as a DNNF circuit whose size is in:*

$$\mathcal{O} \left(m \cdot |\{c(\pi) \mid \pi \subseteq \text{FEAS}(I, \mathfrak{T})\}| \cdot \max_{T \in \mathcal{T}} (|\{f(T, A) \mid \pi \in \text{FEAS}(I, \mathfrak{T})\}|)^{k+1} \right),$$

where k is the pathwidth of the overlap graph $G = \langle \mathcal{T}, E \rangle$ of \mathfrak{T} .

PROOF. We use a similar strategy as for Theorem 6.2.8. The general idea is that because the type aggregator is additive, we can keep track of the current value of the quotas, and then, when deciding whether a project can be selected or not, we can check the current quota value before making our choice.

Let $G = \langle \mathcal{T}, E \rangle$ be the overlap graph of $\langle \mathcal{T}, \mathcal{Q}, q, f \rangle$. We order the projects in the same order in which they are introduced in an optimal path-decomposition of G , with p_1 being the first project, p_2 the second and so forth. As before, we will then define the \vee -nodes $N(j, \mathbf{v}, \mathbf{q})$, where j is a project index, $\mathbf{v} \in \mathbb{R}_{\geq 0}^d$ a vector of used resources and $\mathbf{q} \in Q^{|\mathcal{T}|}$ is a vector of current quota value.

We first introduce notation. Let $\mathcal{T}_{p_j} = \{T \in \mathcal{T} \mid p_j \in T\}$ be the set of types containing project p_j . For $\mathbf{q} \in Q^{|\mathcal{T}|}$, define \mathbf{q}^{p_j} as $\mathbf{q}_T^{p_j} = \mathbf{q}_T \oplus s(p_j)$ for every $T \in \mathcal{T}_{p_j}$ and $\mathbf{q}_T^{p_j} = \mathbf{q}_T$ for every $T \notin \mathcal{T}_{p_j}$. So \mathbf{q}^{p_j} is the updated values for the quotas if p_j is selected given a vector of current quota value $\mathbf{q} \in Q^{|\mathcal{T}|}$.

The \vee -nodes $N(j, \mathbf{v}, \mathbf{q})$ are then defined as follows (where each \blacktriangleright should be understood as an “otherwise”):

- \blacktriangleright If $j = m + 1$, we have $N(j, \mathbf{v}, \mathbf{q}) = \top$;
- \blacktriangleright If there are two types $T_1, T_2 \in \mathcal{T}_{p_j}$ such that p_j is the last project in T_2 to be considered, and selecting project p_j would lead to a violation for T_1 , namely $\mathbf{q}_{T_1}^{p_j} >_Q q(T_1)^+$, but not selecting project p_j would lead to a violation for T_2 , namely $\mathbf{q}_{T_2}^{p_j} <_Q q(T_2)^-$, then $N(j, \mathbf{v}, \mathbf{q}) = \perp$;

- If there is a type $T \in \mathcal{T}_{p_j}$ such that $\mathbf{q}_T^{p_j} >_Q q(T)^+$, then:

$$N(j, \mathbf{v}, \mathbf{q}) = (x_{p_j} \wedge \perp) \vee (\neg x_{p_j} \wedge N(j+1, \mathbf{v}, \mathbf{q}));$$

- If there is a type $T \in \mathcal{T}_{p_j}$ such that p_j is the last project from T to be considered and the quota over T satisfies $\mathbf{q}_T <_Q q(T)^-$ and $\mathbf{q}_T^{p_j} \geq_Q q(T)^-$, then:

$$N(j, \mathbf{v}, \mathbf{q}) = \begin{cases} (x_{p_j} \wedge N(j+1, \mathbf{v} + c(p_j), \mathbf{q}^{p_j})) \vee (\neg x_{p_j} \wedge \perp) & \text{if } \mathbf{v} + c(p_j) \leq \mathbf{b}, \\ \perp & \text{otherwise;} \end{cases}$$

- If $\mathbf{v} + c(p_j) \leq \mathbf{b}$, then:

$$N(j, \mathbf{v}, \mathbf{q}) = (x_{p_j} \wedge N(j+1, \mathbf{v} + c(p_j), \mathbf{q}^{p_j})) \vee (\neg x_{p_j} \wedge N(j+1, \mathbf{v}, \mathbf{q}));$$

- In all other cases, $N(j, \mathbf{v}, \mathbf{q}) = (x_{p_j} \wedge \perp) \vee (\neg x_{p_j} \wedge N(j+1, \mathbf{v}, \mathbf{q}))$.

The tractable embedding for types and quotas, written TE_{quo} , refers to the integrity constraint defined by $N(1, \mathbf{0}_m, \mathbf{0}_{|\mathcal{T}|})$.

It is clear that TE_{quo} returns an integrity constraint represented as a DNNF circuit. We now show that the size of DNNF circuit is in:

$$\mathcal{O} \left(m \cdot |\{c(\pi) \mid \pi \subseteq \text{FEAS}(I, \mathfrak{T})\}| \cdot \max_{T \in \mathcal{T}} (|\{f(T, A) \mid \pi \in \text{FEAS}(I, \mathfrak{T})\}|)^{k+1} \right),$$

where k is the path-width of the overlapping graph.

Let us begin by observing that the first two terms of the product above are the same as for the other tractable embeddings presented before, we do not expend on them. Then, observe that a quota can take at most $\max_{T \in \mathcal{T}} (|\{f(T, A) \mid \pi \in \text{FEAS}(I, \mathfrak{T})\}|)$ different values. Hence, each component of \mathbf{q} can only have $\max_{T \in \mathcal{T}} (|\{f(T, A) \mid \pi \in \text{FEAS}(I, \mathfrak{T})\}|)$ distinct values. Now, how do we get it to the power k and not $|\mathcal{T}|$ in the size of the DNNF circuit? The idea is similar to that of the proof of Theorem 6.2.8. Projects are ordered according to the ordering of types in an optimal path-decomposition of the overlap graph. By doing so, in each node $N(j, \mathbf{v}, \mathbf{q})$, we can “forget” all types in \mathbf{q} for which we already considered all projects as the value of the corresponding quota will no longer change. Hence, the maximum number of types we need to keep track of their quota is upper bounded by $k + 1$. That proves the claim about the size of the DNNF circuit.

We now show that the embedding is correct. Let us consider an arbitrary JA outcome $J \in \mathfrak{J}(TE_{quo}(I, \mathfrak{T}))$. It should be clear from the correctness proofs of both TE and TE_{dep} that $\tau(J)$ satisfies the budget constraint. Moreover, in TE_{quo} , before considering any project, if the current value of a quota violate the quota

constraint associated with it, the branch in the DNNF circuit ends up on the \perp leaf. Because the type aggregator is additive, we know that the quota values in \mathbf{v} are the quota values of the corresponding budget allocation. Hence we have $\tau(j) \in \text{FEAS}(I, \mathfrak{T})$. We have thus showed that $\{\tau(J) \mid J \in \mathfrak{J}(TE_{dep}(I, \mathfrak{T}))\} \subseteq \text{FEAS}(I, \mathfrak{T})$ holds. To prove correctness, we need to additionally prove that the reversed inclusion holds. We omit this part of the proof as the details are exactly as in the proof of Proposition 6.2.5. \square

It should be noted that the factor $\max_{T \in \mathcal{T}} (|\{f(T, A) \mid \pi \in \text{FEAS}(I, \mathfrak{T})\}|)^{k+1}$ in the size of the DNNF circuit produced can be very high. However, for the cardinality and the cost-type aggregators, we can derive the following bounds:

$$\begin{aligned} \max_{T \in \mathcal{T}} |\{f^{\text{card}}(T, A) \mid \pi \in \text{FEAS}(I, \mathfrak{T})\}| &= \max_{T \in \mathcal{T}} q(T)^+ \leq |\mathcal{P}|, \\ \max_{T \in \mathcal{T}} |\{f^{\text{cost}}(T, A) \mid \pi \in \text{FEAS}(I, \mathfrak{T})\}| &= \max_{T \in \mathcal{T}} q(T)^+ \leq \prod_{r \in \mathcal{R}} b_r. \end{aligned}$$

This implies that the integrity constraint for these quota aggregators would be of reasonable size, as long as k —the pathwidth of the overlap graph—is small.

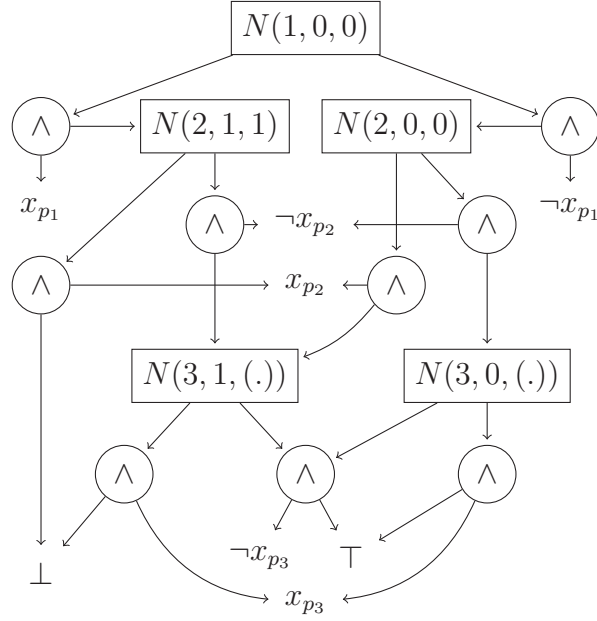
The question is then how small or large can k be? Of course, in principle, it can be as large as the number of projects (plus one). That is the case if all projects appear in all types. On the other hand, the pathwidth of the overlapping graph would be 1 in the case that the types are not overlapping (no project appears in more than one type). This special case is actually very natural. For instance, if one uses types to represent areas of a city that are to be developed, then no project concerning one area will also concern another area. We get the following statement for non-overlapping types.

Corollary 6.2.13. *Let $I = \langle \mathcal{P}, \mathcal{R}, c, \mathbf{b} \rangle$ be a PB instance and $\mathfrak{T} = \langle \mathcal{T}, \mathcal{Q}, q, f \rangle$ a type system where f is an additive type aggregator defined over the score type function s . If the types are not overlapping, that is, if the overlapping graph of $\langle \mathcal{T}, \mathcal{Q}, q, f \rangle$ is the empty graph, then the size of the integrity constraint returned by TE_{quo} is in*

$$\mathcal{O} \left(m \cdot |\{c(P) \mid P \subseteq \mathcal{P}\}| \cdot \max_{T \in \mathcal{T}} (|\{f(T, A) \mid A \in \text{FEAS}(I, \mathfrak{T})\}|) \right).$$

Let us conclude this section by presenting TE_{quo} on our running example.

Example 6.2.14. Consider the instance described in Example 6.2.4. Assume that, projects p_1 and p_2 are health-related projects and that no more than one should be selected. In terms of our model, this means that we are considering the type system $\langle \{T\}, \langle \mathbb{N}, +, 0, \leq \rangle, q, f^{\text{card}} \rangle$, where $T = \{p_1, p_2\}$ and $q(T) = (0, 1)$. An optimal path-decomposition of the overlap graph is thus $(\{p_1, p_2\}, \{p_3\})$. So we will consider the projects in the following order: p_1, p_2, p_3 . The embedding TE_{quo} would return the DNNF circuit presented next.



It is important to see how we can “forget” about the truth values of x_{p_1} and x_{p_2} once we consider project p_3 as the types corresponding to these projects will not play a role any more. \triangle

6.3 Enforcing Exhaustiveness

As we mentioned already at several points in this thesis (see, e.g., Section 3.4.1), *exhaustiveness* is usually a very basic requirement for PB rules, especially under the assumption that no project yields negative satisfaction to anyone. Because the scenarios typically modelled using JA are rather different from PB, the exhaustiveness axiom is not satisfied by common JA rules. Indeed, JA rules are usually designed to be majority-consistent—they would always return the majoritarian outcome if it is admissible—which is incompatible with exhaustiveness. This has to do with the semantics of rejection (of a proposition) in the context of JA: submitting a judgment in which an atom is mapped to false is usually considered as a clear rejection of the proposition, while in PB, not approving a project is usually not considered as a rejection (or is at least ambiguous, see our discussion in Section 3.1.2). In this section, we discuss how to enforce exhaustiveness in our setting.

Let us first extend the definition of exhaustiveness to JA rules and embeddings. An embedding $E : \mathcal{I} \rightarrow \mathcal{L}_{\mathfrak{X}}$ is said to be *exhaustive* if, for every instance $I \in \mathcal{I}$, we have $\tau(\mathfrak{J}(E(I))) \subseteq \text{FEAS}_{\text{EX}}(I)$. On the other hand, an exhaustive embedding E is *correct* if $\text{FEAS}_{\text{EX}}(I) = \tau(\mathfrak{J}(E(I)))$ for every instance I . Finally, a JA rule F is said to be *exhaustive* if for every correct embedding E , every instance $I \in \mathcal{I}$ and every profile \mathbf{A} , it is the case that $\tau(F(E(I), \mathbf{A})) \subseteq \text{FEAS}_{\text{EX}}(I)$.

We now show the incompatibility between majority-consistency and exhaustiveness that we mentioned above.

Proposition 6.3.1. *No majority-consistent JA rule is exhaustive.*

PROOF. Consider a correct but not exhaustive embedding E (for instance TE as defined above). As E is not exhaustive, there exists a PB instance I such that there is at least one admissible JA outcome $J \in \mathfrak{J}(E(I))$ with $\tau(J) \notin \text{FEAS}_{\text{EX}}(I)$. Now consider a profile \mathbf{A} with n agents in which $\lceil n/2 \rceil + 1$ agents only approve of the projects in $\tau(J)$; the other agents being unconstrained. On the JA side, the majoritarian outcome will be J . Since the majoritarian outcome is admissible, any majority-consistent rule F must return $\{J\}$ on $E(I)$ and \mathbf{A} , which does not correspond to an exhaustive budget allocation. \square

This result is far-reaching because exhaustiveness really clashes with basic properties of JA.⁴⁹ To circumvent this problem and to enforce exhaustiveness, we will investigate two approaches: either encoding exhaustiveness in the integrity constraint or designing new JA rules.

6.3.1 Exhaustive Embeddings for Single-Resource Instances

We introduce the *exhaustive tractable embedding*, which is an adaptation of TE designed to maintain exhaustiveness when there is exactly one resource.

Consider a PB instance $I = \langle \mathcal{P}, \mathcal{R}, c, \mathbf{b} \rangle$ with $\mathcal{R} = \{r\}$. Similarly to the previous embeddings, we introduce the \vee -nodes of the integrity constraint as $N(j, v, c^*)$, where j is a project index, v is the budget used in terms of the unique resource r , and c^* is the cost of the cheapest non-selected project. They are defined as follows:

► If $j = m + 1$, then we have:

$$N(j, v, c^*) = \begin{cases} \top & \text{if } c^* > b - v \\ \perp & \text{otherwise.} \end{cases}$$

► If $v + c(p_j) \leq b$, then we have:

$$N(j, v, c^*) = (x_{p_j} \wedge N(j + 1, v + c(p_j), c^*)) \vee (\neg x_{p_j} \wedge N(j + 1, v, \min(c^*, c(p_j)))) .$$

► Otherwise, $N(j, v, c^*) = (\neg x_{p_j} \wedge N(j + 1, v, c^*)) \vee (x_{p_j} \wedge \perp)$.

The exhaustive tractable embedding E_{TE} returns the integrity constraint defined by $N(1, 0, \max_{p \in \mathcal{P}} c(p))$. We prove that the embedding E_{TE} behaves as it is expected to.

⁴⁹Were it not for our assumption that every project must have at least one supporter (which rules out certain profiles), Proposition 6.3.1 could be strengthened to say that no unanimous JA rule is exhaustive (F is *unanimous* if $F(J, \dots, J) = \{J\}$ for all judgments J satisfying the integrity constraint Γ).

Proposition 6.3.2. *The exhaustive tractable embedding is correct and exhaustive, and returns an integrity constraint Γ represented as a DNNF circuit of size $\mathcal{O}(m^2 \cdot |\{c(\pi) \mid \pi \subseteq \text{FEAS}(I)\}|)$, for any $I = \langle \mathcal{P}, \mathcal{R}, c, \mathbf{b} \rangle$ where $|\mathcal{R}| = 1$.*

PROOF. The structure of the embedding is very similar to that of the tractable embedding TE presented above. We only prove that the embedding is exhaustive. Let $I = \langle \mathcal{P}, \mathcal{R}, c, \mathbf{b} \rangle$ be a PB instance. Consider an outcome $J \in \mathfrak{J}(ETE(I))$ and the budget allocation π such that $\pi = \tau(J)$. Note that the budget allocation π is exhaustive if and only if the cheapest not selected project does not fit in it. Observe that in the exhaustive tractable embedding we are keeping track of the cheapest project that has not been selected so that whenever all the projects have been considered two cases are left. If the cheapest project fits in the budget allocation then the latter is not exhaustive and we link the branch to the \perp leaf. Otherwise the budget allocation is exhaustive and the \top leaf is linked to it. This proves that the embedding is exhaustive. The fact that the constructed integrity constraint is a DNNF circuit of the suitable size is almost immediate given all the proofs we have already seen on that topic. \square

The embedding ETE is only defined for instances with a single resource and, unfortunately, the idea does not generalise if computational efficiency is required. The reason is that, when there are several resources, then there could be exponentially many “cheapest projects”. In the following we turn to another way of enforcing exhaustiveness, based on the use of different rules, the *asymmetric* ones.

6.3.2 Asymmetric Judgment Aggregation Rules

As we saw above, typical JA rules fail exhaustiveness because rejection is interpreted differently in JA than in PB. Therefore, to implement PB via JA we need to adapt them so that not selecting a project (*i.e.*, not accepting a proposition) is not interpreted as a rejection. To this end we introduce a new family of *asymmetric* JA rules. They avoid the symmetric treatment of acceptance and rejection common in most, if not all, established JA rules.

Definition 6.3.3 (Asymmetric Additive Rules). *Let F be an additive JA rule associated with $f : (2^{\mathfrak{X}})^n \times \text{Lit}(\mathfrak{X}) \rightarrow \mathbb{R}_{\geq 0}$. Then its asymmetric counterpart F_{asy} is the rule where for every integrity constraint Γ and every profile \mathbf{J} , we have:*

$$F_{asy}(\Gamma, \mathbf{J}) = \arg \max_{J \in \mathfrak{J}(\Gamma)} \sum_{\substack{\ell \in \text{aug}(J) \\ \ell \text{ is positive}}} f(\mathbf{J}, \ell) + \epsilon,$$

where ϵ is a small positive constant such that:

$$0 < \epsilon < \frac{1}{|\mathfrak{X}|} \cdot \min \{f(\mathbf{J}, \ell) \neq 0 \mid \mathbf{J} \in (2^{\mathfrak{X}})^n, \ell \in \text{aug}(J), \ell \text{ is positive}\}.$$

Importantly, this definition applies only if f is $\mathbb{R}_{\geq 0}$ -valued. The use of ϵ guarantees that accepting positive literals will always be more appealing than accepting negative ones, while being small enough so as to not impact the relative values of positive literals. Note that $\epsilon = \frac{1}{|\mathbb{X}|+1}$ is a suitable choice for the Slater, the Kemeny and the leximax rules.

The class of asymmetric additive rules is particularly interesting for us, as we can show that every rule in this class satisfies exhaustiveness.

Proposition 6.3.4. *Let F be an additive JA rule associated with an $\mathbb{R}_{\geq 0}$ -valued function f . Then, the asymmetric counterpart of F satisfies exhaustiveness.*

PROOF. Executing F_{asy} involves computing a score for every admissible outcome J . By definition, no negative literal in J can contribute to its score, while every positive literal makes a strictly positive contribution of at least ϵ . Thus, flipping a negative literal always results in an increased score. So F_{asy} only returns admissible judgments for which flipping any negative literal would violate the integrity constraint. This exactly corresponds to exhaustiveness. \square

Observe that the asymmetric counterpart of any additive rule is itself additive. The asymmetric counterpart of a scoring rule is also a scoring rule. Indeed, F_{asy} is the scoring rule defined with respect to $s_{asy} : (J, \ell) \mapsto \mathbb{1}_{\ell \text{ is positive}} \cdot s(J, \ell)$, where s is the scoring function corresponding to F . Note that this does not work for AMR since the gain function cannot filter out negative literals (since it is a function from \mathbb{N} to \mathbb{R}).

Finally, it is interesting to note that the asymmetric variant of the leximax rule is very similar to the GREEDCOST PB rule.⁵⁰

6.4 Axiomatic Analysis of Judgment Aggregation Rules

In this section we investigate to what extent important axioms proposed in the literature on PB are satisfied by JA rules, when used for the purpose of PB.

We focus on the monotonicity axioms introduced in Section 3.4.2, generalising the definitions to allow for multiple resources and irresolute rules.

Formally, for a given axiom \mathcal{A} about PB rules, we say that the JA rule F satisfies \mathcal{A} with respect to embedding E if the PB rule mapping any instance I and profile \mathbf{A} to $\tau(F(E(I), \mathbf{A}))$ satisfies \mathcal{A} .

Moreover, for a resolute rule F , the monotonicity axioms of Section 3.4.2 are usually stated as “when one moves from an instance/profile pair (I, \mathbf{A}) to another pair

⁵⁰GREEDCOST actually corresponds to the asymmetric variant of a refinement of the leximax rule known as the *ranked-agenda rule* (Lang, Pigozzi, Slavkovik and Van der Torre, 2011).

(I', \mathbf{A}') , then if $F(I, \mathbf{A})$ satisfies a certain property, $F(I', \mathbf{A}')$ should satisfy a corresponding property.” We generalise these axioms to the irresolute case by requiring that, if *every* budget allocation returned by our rule for (I, \mathbf{A}) satisfies the property in question, then *every* budget allocation for (I', \mathbf{A}') should satisfy the corresponding property. This is a universal extension of the axioms, in contrast to the existential extension we discussed at the end of Section 3.4.2.

The first axiom we study is called *limit monotonicity* (Definition 3.4.3). Recall that it states that after any increase in the budget limit that is not so substantial as to make some previously unaffordable project affordable, any funded project should continue to get funded.

Definition 6.4.1 (Limit Monotonicity). *An irresolute PB rule F is said to satisfy limit-monotonicity if, for any two PB instances $I = \langle \mathcal{P}, \mathcal{R}, c, \mathbf{b} \rangle$ and $I' = \langle \mathcal{P}, \mathcal{R}, c, \mathbf{b}' \rangle$ with $\mathbf{b} \leq \mathbf{b}'$ and $c(p) \leq \mathbf{b}$ for all projects $p \in \mathcal{P}$, it is the case that $\bigcap F(I, \mathbf{A}) \subseteq \bigcap F(I', \mathbf{A})$ for all profiles \mathbf{A} .*

In the following, we show that this axiom is not satisfied by any of the JA rules of interest, even when there is only one resource.

Proposition 6.4.2. *None of the Kemeny, the Slater, or the leximax rules, or their asymmetric counterparts satisfy limit monotonicity with respect to any correct embedding.*

PROOF. Assume the embedding is correct. Consider two instances $I = \langle \mathcal{P}, \mathcal{R}, c, \mathbf{b} \rangle$ and $I' = \langle \mathcal{P}, \mathcal{R}, c, \mathbf{b}' \rangle$ such that $|\mathcal{R}| = 1$. There are three projects p_1, p_2 , and p_3 , and the budget limits are $\mathbf{b} = (4)$ and $\mathbf{b}' = (5)$. The profile of interest \mathbf{A} together with the instances I and I' are presented below, I on the left and I' on the right.

	p_1	p_2	p_3		p_1	p_2	p_3
Cost	3	2	1	Cost	3	2	1
A_1	✓	✓	✓	A_1	✓	✓	✓
A_2	✓	✓	✓	A_2	✓	✓	✓
A_3	✓	✓	✓	A_3	✓	✓	✓
A_4	✓	✓	✗	A_4	✓	✓	✗
A_5	✓	✗	✗	A_5	✓	✗	✗
	$\mathbf{b} = (4)$				$\mathbf{b}' = (5)$		

We claim that on I , the Kemeny, the leximax rules, and their asymmetric counterpart would all return $\{\{p_1, p_3\}\}$. However, they would all return $\{\{p_1, p_2\}\}$ on I' . Project p_3 is thus a witness of the violation of limit monotonicity.

For the Slater rule and its asymmetric counterpart, consider the situation depicted below with two instances $I = \langle \mathcal{P}, \mathcal{R}, c, \mathbf{b} \rangle$ and $I' = \langle \mathcal{P}, \mathcal{R}, c, \mathbf{b}' \rangle$ involving

three projects and a single resource. The budget limits are $\mathbf{b} = (4)$ and $\mathbf{b}' = (6)$. We denote by \mathbf{A} (on the left) and \mathbf{A}' (on the right) the corresponding profiles.

	p_1	p_2	p_3		p_1	p_2	p_3
Cost	1	2	4	Cost	1	2	4
A_1	✓	✓	✓	A_1	✓	✓	✓
$\mathbf{b} = (4)$				$\mathbf{b}' = (6)$			

On I and \mathbf{A} , both the Slater rule and its asymmetric counterpart would return $\{\{p_1, p_2\}\}$. On the other hand, on I' and \mathbf{A} , both rules would return $\{\{p_1, p_2\}, \{p_1, p_3\}, \{p_2, p_3\}\}$. Since $\bigcap \{\{p_1, p_2\}, \{p_1, p_3\}, \{p_2, p_3\}\} = \emptyset$, both projects p_1 and p_2 are witnesses of the violation of limit monotonicity. \square

We move on to *discount monotonicity* (Definition 3.4.2), an axiom stating that, if the cost of a selected project p decreases, then p should continue to be selected.

Definition 6.4.3 (Discount Monotonicity). *An irresolute PB rule F is said to satisfy discount-monotonicity if, for any two PB instances $I = \langle \mathcal{P}, \mathcal{R}, c, \mathbf{b} \rangle$, $I' = \langle \mathcal{P}, \mathcal{R}, c', \mathbf{b} \rangle$ such that for some distinguished project $p^* \in \mathcal{P}$, we have $c(p^*) \geq c'(p^*)$, and $c(p) = c'(p)$ for all $p \in \mathcal{P} \setminus \{p^*\}$, it is the case that $p^* \in \bigcap F(I, \mathbf{A})$ implies $p^* \in \bigcap F(I', \mathbf{A})$ for all profiles \mathbf{A} .*

To study how JA rules deal with discount monotonicity, we introduce a new axiom for JA. As we will prove, this axiom is a sufficient condition for discount monotonicity.

Definition 6.4.4 (Constraint monotonicity). *A JA rule F is said to satisfy constraint-monotonicity if, for any two integrity constraints $\Gamma, \Gamma' \in \mathcal{L}_x$ with $\mathfrak{J}(\Gamma) \subseteq \mathfrak{J}(\Gamma')$ and any profile \mathbf{J} , it is the case that $F(\Gamma', \mathbf{J}) \setminus F(\Gamma, \mathbf{J}) \subseteq \mathfrak{J}(\Gamma') \setminus \mathfrak{J}(\Gamma)$.*

This axiom states that if the integrity constraint Γ is weakened into Γ' in the sense that more judgements are now admissible, then any new judgments returned by F must be taken from the set of newly admissible judgments.

We obtain the following formal connection between the two axioms.

Proposition 6.4.5. *Every constraint-monotonic JA rule is discount-monotonic with respect to any correct embedding.*

PROOF. Let F be a JA rule that is constraint-monotonic. Let E be a correct embedding. Consider the instances $I = \langle \mathcal{P}, \mathcal{R}, c, \mathbf{b} \rangle$ and $I' = \langle \mathcal{P}, \mathcal{R}, c', \mathbf{b} \rangle$, where a project $p^* \in \mathcal{P}$ became cheaper from I to I' as in Definition 6.4.3.

Let \mathbf{A} be an arbitrary profile such that p^* is always selected, i.e., we have $p^* \in \bigcap \tau(F(E(I), \mathbf{A}))$. We need to show that $p^* \in \bigcap \tau(F(E(I'), \mathbf{A}))$ holds. Observe that $\text{FEAS}(I) \subseteq \text{FEAS}(I')$. Because E is correct, we also have $\mathfrak{J}(E(I)) \subseteq \mathfrak{J}(E(I'))$. Moreover, for every new feasible budget allocation $\pi \in \text{FEAS}(I') \setminus \text{FEAS}(I)$, it should be the case that $p^* \in \pi$ as only $c'(p^*)$ changed in I' . Hence, for every newly admissible outcome $J \in \mathfrak{J}(E(I')) \setminus \mathfrak{J}(E(I))$, we also have $p^* \in \tau(J)$. Since the JA rule F is constraint-monotonic, for every profile profile \mathbf{A} , we have:

$$F(E(I'), \mathbf{A}) \subseteq F(E(I), \mathbf{A}) \cup (\mathfrak{J}(E(I')) \setminus \mathfrak{J}(E(I)))'$$

Thus, for every $J \in F(E(I'), \mathbf{A})$, it should be the case that $p^* \in \tau(J)$. This proves that F satisfies discount-monotonicity. \square

Interestingly, a large set of JA rules satisfy constraint monotonicity.

Proposition 6.4.6. *Every additive rule is constraint-monotonic.*

PROOF. Consider any additive rule F . Suppose, that F is *not* constraint-monotonic. Then there exist two integrity constraints Γ and Γ' with $\mathfrak{J}(\Gamma) \subseteq \mathfrak{J}(\Gamma')$ and a profile \mathbf{J} for which there exists a $J \in F(\Gamma', \mathbf{J}) \setminus F(\Gamma, \mathbf{J})$ with $J \notin \mathfrak{J}(\Gamma') \setminus \mathfrak{J}(\Gamma)$, that is, such that $J \in \mathfrak{J}(\Gamma') \cap \mathfrak{J}(\Gamma)$. Since by assumption on Γ and Γ' , we have $\mathfrak{J}(\Gamma) \subseteq \mathfrak{J}(\Gamma')$, it is the case that $\mathfrak{J}(\Gamma') \cap \mathfrak{J}(\Gamma) = \mathfrak{J}(\Gamma)$, and thus that $J \in \mathfrak{J}(\Gamma)$. As $J \notin F(\Gamma, \mathbf{J})$, there exists some other outcome $J' \in \mathfrak{J}(\Gamma)$ with a higher total score than that of J . Moreover, since $\mathfrak{J}(\Gamma) \subseteq \mathfrak{J}(\Gamma')$, this same J' would outperform J also under Γ' . This implies that $J \notin F(\Gamma', \mathbf{J})$, which is a contradiction. \square

Recall that several well-known JA rules are additive and thus subject to this result.

Corollary 6.4.7. *The Kemeny, Slater, and leximax rules as well as their asymmetric counterparts are all discount-monotonic with respect to any correct embedding.*

We last consider *splitting monotonicity* (Definition 3.4.4) and *merging monotonicity* (Definition 3.4.5), two axioms that deal with situations where projects are either split into subprojects, or merged into super-projects. Remember the action of splitting a project $p \in \mathcal{P}$ into P , and the dual action of merging a subset of projects P into p as defined in Definition 3.4.4. These actions are naturally extended to the multi-resource setting. We present below our adaptations of the axioms to irresolute rules.

Definition 6.4.8 (Splitting Monotonicity). *An irresolute PB rule F is said to satisfy splitting-monotonicity if, for any two PB instances $I = \langle \mathcal{P}, \mathcal{R}, c, \mathbf{b} \rangle$, $I' = \langle \mathcal{P}', \mathcal{R}, c', \mathbf{b} \rangle$ with corresponding profiles \mathbf{A} and \mathbf{A}' such that I' and \mathbf{A}' are the result of splitting project p into P given I and \mathbf{A} , it is the case that if $p \in \bigcap F(I', \mathbf{A})$ then $\pi' \cap P \neq \emptyset$ for all $\pi' \in F(I', \mathbf{A})$.*

Definition 6.4.9 (Merging Monotonicity). *An irresolute PB rule F is said to satisfy merging-monotonicity if, for any two PB instances $I = \langle \mathcal{P}, \mathcal{R}, c, \mathbf{b} \rangle$, $I' = \langle \mathcal{P}', \mathcal{R}, c', \mathbf{b} \rangle$ with corresponding profiles \mathbf{A} and \mathbf{A}' such that I' and \mathbf{A}' are the result of merging project set P into project p given I and \mathbf{A} , it is the case that $P \subseteq \bigcap F(I, \mathbf{A})$ implies $p \in \bigcap F(I', \mathbf{A}')$.*

We first show that splitting monotonicity is satisfied by AMRs and their asymmetric counterparts.

Proposition 6.4.10. *Every AMR as well as the asymmetric counterpart of every AMR are splitting-monotonic with respect to any correct embedding.*

PROOF. Let F be either an AMR or the asymmetric counterpart of an AMR, and let E be a correct and exhaustive embedding. Consider a PB instance $I = \langle \mathcal{P}, \mathcal{R}, c, \mathbf{b} \rangle$ and a profile \mathbf{A} . Let \mathbf{J} be the JA profile corresponding to \mathbf{A} . Let $I' = \langle \mathcal{P}', \mathcal{R}, c', \mathbf{b} \rangle$ and \mathbf{A}' be the instance and profile resulting from splitting a given project $p^* \in \bigcap \tau(F(E(I), \mathbf{J}))$ into the set of projects P^* . Let \mathbf{J}' be the JA profiles corresponding to \mathbf{A}' .

Consider an outcome $J_1 \in F(E(I), \mathbf{J})$. Note that any outcome J that is admissible before and after the splitting, i.e., any $J \in \mathfrak{J}(E(I)) \cap \mathfrak{J}(E(I'))$, cannot include either p^* or any project from P^* . Since $p^* \in \bigcap \tau(F(E(I), \mathbf{J}))$, this implies that J_1 has a higher total score than any $J \in \mathfrak{J}(E(I)) \cap \mathfrak{J}(E(I'))$.

Consider now an outcome $J'_1 = (J_1 \setminus \{p^*\}) \cup \{p\}$ for some newly created project $p \in P^*$. By definition of the new cost function, and since E is a correct embedding, we have $J'_1 \models E(I')$. J'_1 thus determines an admissible outcome for the constraint corresponding to I' . Based on the definition of \mathbf{J}' , it is clear that $n_\ell^{\mathbf{J}} = n_\ell^{\mathbf{J}'}$ for every $\ell \in \text{aug}(J_1 \setminus \{x_{p^*}\})$ and that $n_{x_{p^*}}^{\mathbf{J}} = n_{x_p}^{\mathbf{J}'}$. Hence, because the internal score used by F only depends on the number of supporters, we know that J_1 and J'_1 have the same total score. This implies that J'_1 has a higher total score than any $J \in \mathfrak{J}(E(I)) \cap \mathfrak{J}(E(I'))$. Thus, $\mathfrak{J}(E(I)) \cap F(E(I'), \mathbf{J}') = \emptyset$. As for every newly admissible judgement $J' \in \mathfrak{J}(E(I')) \setminus \mathfrak{J}(E(I))$ we have $P^* \cap \tau(J') \neq \emptyset$, every outcome returned by F would have a non-empty intersection with P^* . \square

Interestingly, this result provides a sufficient condition for PB rules to satisfy splitting monotonicity: any rule behaving as an AMR (with a suitable definition of AMR for PB) will satisfy it.

For the specific set of rules we study, we obtain the following corollary.

Corollary 6.4.11. *The Kemeny, Slater, and leximax rules as well as their asymmetric counterparts are all splitting-monotonic with respect to any correct embedding.*

Interestingly, when enforcing exhaustiveness through the embedding, this last result is no longer valid for symmetric rules.

Proposition 6.4.12. *None of the Kemeny, the Slater, or the leximax rules satisfy splitting monotonicity with respect to any correct and exhaustive embedding.*

PROOF. Consider the following pairs of instances and three-agent profiles: I and A on the left and I' and A' on the right. Both involve just one resource.

$p_1 \quad p_2 \quad p_3 \quad p_4$	$p_1 \quad p_2^1 \quad p_2^2 \quad p_2^3 \quad p_2^4 \quad p_3 \quad p_4$																																																				
<table style="width: 100%; border-collapse: collapse;"> <tr> <td style="text-align: left; padding-right: 10px;">Cost</td> <td style="text-align: center;">2</td> <td style="text-align: center;">2</td> <td style="text-align: center;">1</td> <td style="text-align: center;">1</td> </tr> <tr> <td style="text-align: left; padding-right: 10px;">A_1</td> <td style="text-align: center;">✓</td> <td style="text-align: center;">✗</td> <td style="text-align: center;">✗</td> <td style="text-align: center;">✗</td> </tr> <tr> <td style="text-align: left; padding-right: 10px;">A_2</td> <td style="text-align: center;">✓</td> <td style="text-align: center;">✗</td> <td style="text-align: center;">✗</td> <td style="text-align: center;">✗</td> </tr> <tr> <td style="text-align: left; padding-right: 10px;">A_3</td> <td style="text-align: center;">✗</td> <td style="text-align: center;">✓</td> <td style="text-align: center;">✓</td> <td style="text-align: center;">✓</td> </tr> </table>	Cost	2	2	1	1	A_1	✓	✗	✗	✗	A_2	✓	✗	✗	✗	A_3	✗	✓	✓	✓	<table style="width: 100%; border-collapse: collapse;"> <tr> <td style="text-align: left; padding-right: 10px;">Cost</td> <td style="text-align: center;">2</td> <td style="text-align: center;">1/2</td> <td style="text-align: center;">1/2</td> <td style="text-align: center;">1/2</td> <td style="text-align: center;">1/2</td> <td style="text-align: center;">1</td> <td style="text-align: center;">1</td> </tr> <tr> <td style="text-align: left; padding-right: 10px;">A_1</td> <td style="text-align: center;">✓</td> <td style="text-align: center;">✗</td> <td style="text-align: center;">✗</td> <td style="text-align: center;">✗</td> <td style="text-align: center;">✗</td> <td style="text-align: center;">✗</td> <td style="text-align: center;">✗</td> </tr> <tr> <td style="text-align: left; padding-right: 10px;">A_2</td> <td style="text-align: center;">✓</td> <td style="text-align: center;">✗</td> <td style="text-align: center;">✗</td> <td style="text-align: center;">✗</td> <td style="text-align: center;">✗</td> <td style="text-align: center;">✗</td> <td style="text-align: center;">✗</td> </tr> <tr> <td style="text-align: left; padding-right: 10px;">A_3</td> <td style="text-align: center;">✗</td> <td style="text-align: center;">✓</td> <td style="text-align: center;">✓</td> <td style="text-align: center;">✓</td> <td style="text-align: center;">✓</td> <td style="text-align: center;">✓</td> <td style="text-align: center;">✓</td> </tr> </table>	Cost	2	1/2	1/2	1/2	1/2	1	1	A_1	✓	✗	✗	✗	✗	✗	✗	A_2	✓	✗	✗	✗	✗	✗	✗	A_3	✗	✓	✓	✓	✓	✓	✓
Cost	2	2	1	1																																																	
A_1	✓	✗	✗	✗																																																	
A_2	✓	✗	✗	✗																																																	
A_3	✗	✓	✓	✓																																																	
Cost	2	1/2	1/2	1/2	1/2	1	1																																														
A_1	✓	✗	✗	✗	✗	✗	✗																																														
A_2	✓	✗	✗	✗	✗	✗	✗																																														
A_3	✗	✓	✓	✓	✓	✓	✓																																														
$\mathbf{b} = (4)$	$\mathbf{b} = (4)$																																																				

Observe that I' and A' are the result of splitting p_2 into $\{p_2^1, p_2^2, p_2^3, p_2^4\}$, given I and A . We leave the relevant computations to the reader, but the Kemeny, the Slater, and the leximax rules would all return $\{\{p_1, p_2\}\}$ on (I, A) when used with a correct and exhaustive embedding. However, they would return $\{\{p_1, p_3, p_4\}\}$ on (I', A') . Hence, p_2 is a witness of a violation of splitting monotonicity. \square

We finally investigate merging monotonicity. It turns out that none of the rules we are considering in this chapter satisfy it.

Proposition 6.4.13. *None of the Kemeny, the Slater, or the leximax rules, or their asymmetric counterparts satisfy merging monotonicity with respect to any correct embedding.*

PROOF. Consider the following pairs of instances and one-agent profiles: I and A on the left and I' and A' on the right. Both involve just one resource.

$p_1 \quad p_2 \quad p_3 \quad p_4 \quad p_5 \quad p_6$	$p_1 \quad p_2 \quad p'_3$																						
<table style="width: 100%; border-collapse: collapse;"> <tr> <td style="text-align: left; padding-right: 10px;">Cost</td> <td style="text-align: center;">2</td> <td style="text-align: center;">2</td> <td style="text-align: center;">1</td> <td style="text-align: center;">1</td> <td style="text-align: center;">1</td> <td style="text-align: center;">1</td> </tr> <tr> <td style="text-align: left; padding-right: 10px;">A_1</td> <td style="text-align: center;">✓</td> <td style="text-align: center;">✓</td> <td style="text-align: center;">✓</td> <td style="text-align: center;">✓</td> <td style="text-align: center;">✓</td> <td style="text-align: center;">✓</td> </tr> </table>	Cost	2	2	1	1	1	1	A_1	✓	✓	✓	✓	✓	✓	<table style="width: 100%; border-collapse: collapse;"> <tr> <td style="text-align: left; padding-right: 10px;">Cost</td> <td style="text-align: center;">2</td> <td style="text-align: center;">2</td> <td style="text-align: center;">4</td> </tr> <tr> <td style="text-align: left; padding-right: 10px;">A_1</td> <td style="text-align: center;">✓</td> <td style="text-align: center;">✓</td> <td style="text-align: center;">✓</td> </tr> </table>	Cost	2	2	4	A_1	✓	✓	✓
Cost	2	2	1	1	1	1																	
A_1	✓	✓	✓	✓	✓	✓																	
Cost	2	2	4																				
A_1	✓	✓	✓																				
$\mathbf{b} = (4)$	$\mathbf{b} = (4)$																						

Observe that I' and A' are the result of merging $\{p_3, p_4, p_5, p_6\}$ into project p'_3 , given I and A . It is easy to see that the Kemeny, the Slater, and the leximax rules would all return $\{\{p_3, p_4, p_5, p_6\}\}$ on (I, A) when used with a correct embedding. However, they would return $\{\{p_1, p_2\}\}$ on (I', A') . Hence, p'_3 is a witness of a violation of merging monotonicity. Moreover, since the only agent approves of every projects, the same hold for the asymmetric counterpart of the rules. \square

	Kemeny, Slater and Leximax		
	Symmetric	Symmetric ex. embedding	Asymmetric
Exhaustiveness	✗ Proposition 6.3.1	✓ Proposition 6.3.2	✓ Proposition 6.3.4
Limit Monotonicity	✗ Proposition 6.4.2	✗ Proposition 6.4.2	✗ Proposition 6.4.2
Discount Monotonicity	✓ Corollary 6.4.7	✓ Corollary 6.4.7	✓ Corollary 6.4.7
Splitting Monotonicity	✓ Proposition 6.4.10	✗ Proposition 6.4.12	✓ Proposition 6.4.10
Merging Monotonicity	✗ Proposition 6.4.13	✗ Proposition 6.4.13	✗ Proposition 6.4.13

Table 6.5.1: Axiomatic results for JA rules when used with a correct embedding. “Symmetric ex. embedding” indicates that we use the symmetric rule with a correct and exhaustive embedding.

6.5 Summary

In this chapter, we have presented how to use the expressive power of JA to reason about PB instances with additional constraints.

Our focus has mainly been computational. We spent some time discussing the computational complexity of embedding PB instances into well-structured instances of JA, namely into integrity constraint represented as DNNF circuits. Enforcing the integrity constraints to be DNNF circuits directly entailed that the outcome of common JA rules could be computed efficiently, thus proving the viability of our approach. For two very generic types of constraints—namely dependencies between projects and quotas over types of projects—we also developed specific embeddings, exemplifying the flexibility of our approach when it comes to incorporating additional constraints on the PB side.

The computational analysis only proves that we *can* use JA rules to determine budget allocations. It does not inform us about the quality of the budget allocations that are reached that way. To provide some elements of answer to this second point, we also presented an axiomatic analysis of JA rules, in terms of PB axioms. We discussed exhaustiveness and monotonicity axioms. Our findings are summarised in Table 6.5.1. Interestingly, relative to the range of axioms we have considered here, we may summarise the situation by saying that JA rules perform similarly to other PB rules in normative terms (see Table 3.4.2 for standard PB rules).

Chapter 7

A Long-Term Approach to Participatory Budgeting

We already mentioned how difficult it is to put forward an exact definition of what a Participatory Budgeting (PB) process is. Because of the wide variety of real-life implementations, authors tend to define PB processes in terms criteria they satisfy, as we saw in Section 1.1.1. When discussing some of the requisites we already presented in the introduction, [Sintomer, Herzberg, Röcke and Allegretti \(2012\)](#) note that:

« *It has to be a repeated process over years. Consequently, if a participatory process is already planned as a unique event, we would not consider it as PB: one meeting, one referendum on financial issues are not examples of participatory budgeting. In English, the expression of “participatory budgeting” has been used from the late 1990s in order to stress this notion of an ongoing process (“budgeting”) rather than an outcome (“budget”).* »

Interestingly, they emphasise that a PB process *has* to span several years. The standard model we have studied so far fails to capture this temporal aspect. Fortunately, it can be extended to discuss long-term perspectives of PB. The goal of this chapter is to present one such extension of the standard model, the so-called *Perpetual Participatory Budgeting* (PPB) model.

It is clear that the long-term perspective *should* be studied when investigating PB processes. However, it is not necessarily clear *what* should be studied in an investigation of long-term PB. The naïve approach would be to consider a long-term PB model simply as a repetition of the standard model. Following this approach, one could pick the voting method one prefers and simply repeatedly apply it at each round. This would be a perfectly valid way of handling the long-term approach, and, it is worth

noting, is how PB processes are implemented in real-life. Why did I refer to this approach as a *naïve* one then? My claim, and that of this chapter, is that by doing so, we would miss an opportunity to actually use past information to take better decisions, round after round.

As we like to do for each variation on the model we introduce, let us illustrate the motivation of this chapter through an example. Consider Dean, the mayor of our fictional town. He is absolutely convinced that the only way to obtain an efficient and well-functioning democratic process is by enabling citizens to weight in on the decisions of the public institutions. This is why he has implemented a PB process in his town. There is one budgeting decision per year, over an indefinite time horizon. However, after a couple of successful years, the PB process is now less and less popular, and the turnout has significantly dropped during the last few rounds. In an attempt to identify the causes of this decline, he hired Julian, a decision-making expert. After having reviewed the data and conducted some interviews, Julian has an explanation to offer to Dean: The citizens of smaller neighbourhoods gave up on the PB process because none of their favourite projects ever got selected. The assessment is simple: the way the process is organised does not implement any kind of temporal fairness and the decisions made year after year always favour larger groups of the population.⁵¹

This example, specifically tailored to our needs, illustrates one impact of using the naïve approach discussed above: ignoring the temporal aspects of a PB process can lead to always favouring the same voters, leading to a decision that is unfair, not necessarily for each individual round, but regarding the process in its entirety. The motivation for this chapter is thus to investigate how to deal with long-term fairness. We will try to answer the following question:

*How should fairness properties be defined, and how can we enforce them,
in a long-term model for PB?*

Having now reached Chapter 7, one can claim that we have acquired quite some experience with fairness in PB. Still, defining fairness properties in the perpetual PB (PPB) framework does require some additional pondering. Indeed, the temporal aspect of the model immediately brings up two further challenges.

The first challenge of the PPB model is that the identity of the voters cannot be traced back from one round to the other. This is because ballots are submitted anonymously. This fundamental requirement of any democratic process does entail that it is not possible to know the entire ballot history of a voter. In particular, this implies that we cannot discuss fairness regarding a specific voter that is based on the past voting behaviours. To circumvent this, we will only consider fairness properties that apply to groups of agents, instead of individuals. We assume that the groups are

⁵¹Interestingly, a similar situation—where some citizens complained of having never seen any project they voted for being selected—was described by employees of the municipality in charge of PB when we discussed the PB processes implemented in Amsterdam.

defined exogenously from the process, based on some publicly observable variables such as where the ballot was cast, the gender of the voter, or their age group. These are details that are often available together with the ballots (see pabulib.org for some examples coming from real-life PB processes). Defining fairness criteria for groups of agents—we will call them types of agents—is the first challenge we will be facing.

The second challenge of the PPB model is that we are now discussing an environment that is both dynamic and online (meaning that decisions are taken without knowledge of what is to come). The consequences are twofold. On the one hand, because the environment is dynamic, we need to take into account the fact that new projects will be considered, that the ballots of the agents can change over time, etc...⁵² On the other hand, because the decisions are made in an online fashion, nothing about the future is known when deciding on which projects to select in a given round. We thus need to ensure that decisions are robust against any instance that can occur in subsequent rounds. This is our second challenge.

To answer the research question stated above, while facing the aforementioned challenges, we will follow a similar approach as the one we developed in Chapter 4. We will first define what we consider to represent a case of perfect fairness. Because perfect fairness can rarely be achieved, we will introduce several relaxations that will be shown to be satisfiable. The specifics are to be discovered in the coming sections.

As is customary by now, we will start by mentioning additional related work. The formal model of perpetual PB will then be introduced (Section 7.1). All the components will then be in place for us to present the fairness theory we devised for this model (Section 7.2). We will then begin the formal study by focusing on perfect fairness (Section 7.3). As it cannot always be guaranteed, we will discuss two relaxations of perfect fairness. The first one consists in optimising for fairness (Section 7.4) by selecting projects that lead to an optimal value of some fairness measure. The second one consists in guaranteeing fairness, but only for an infinite horizon (Section 7.5). We will finally draw a conclusion (Section 7.6).

Additional Related Work. The present chapter is the only study of the repetitive aspect of a PB process that we are aware of within the social choice literature. However, several other social choice frameworks have been studied through the long-term lens. Our inspiration for this chapter stems from the *perpetual voting* framework of Lackner (2020), and the study of proportionality in this setting (Lackner and Maly, 2023). A similar long-term perspective has also been studied for score-based aggregation (Freeman, Zahedi and Conitzer, 2017; Freeman, Zahedi, Conitzer and Lee, 2018).

7.1 Perpetual Participatory Budgeting

In the following we describe our long-term model for PB, called *perpetual participatory budgeting* (PPB). In essence, it consists of a sequence instances of the standard PB

⁵²To simplify the model, we assume that the set of agents remains the same throughout the process.

model, each representing a round.

Because there are several rounds, we need one extra layer of notation to refer to projects. We will use \mathbb{P} to represent the (potentially infinite) set of all the projects that can occur throughout the process. We naturally extend the *cost function* $c : \mathbb{P} \rightarrow \mathbb{R}_{>0}$ so that it applies to \mathbb{P} as well, mapping any project $p \in \mathbb{P}$ to its cost $c(p) \in \mathbb{R}_{>0}$. Note that we are only considering uni-dimensional costs here. We will still write $c(P)$ instead of $\sum_{p \in P} c(p)$ for any $P \subseteq \mathbb{P}$. As before, an instance is a tuple $I = \langle \mathcal{P}, c, b \rangle$ where $\mathcal{P} \subseteq \mathbb{P}$ is finite.

Moving on to the temporal side of the model. A *perpetual participatory budgeting instance* of length $k \in \mathbb{N}_{>0} \cup \{\infty\}$ (or k -PPB instance) is a sequence of k instances $\mathbf{I} = (I_1, \dots, I_k)$. Using this notation an ∞ -PPB instance is simply an infinite sequence of PB instances. The instance occurring at round $j \in \mathbb{N}_{>0}$ is denoted by $I_j = \langle \mathcal{P}_j, c, b_j \rangle$, where $\mathcal{P}_j \subseteq \mathbb{P}$ is finite. Notice how the cost function repeats itself at each round (since we extended it to \mathbb{P}). Given a k -PPB instance $\mathbf{I} = (I_1, \dots, I_k)$, a vector $\boldsymbol{\pi} = (\pi_1, \pi_2, \dots)$ of budget allocations such $|\boldsymbol{\pi}| \leq k$ and $\pi_j \subseteq \mathcal{P}_j$ for every round $j \in \{1, \dots, |\boldsymbol{\pi}|\}$ is called a *solution* for \mathbf{I} . A solution $\boldsymbol{\pi}$ for a k -PPB instance \mathbf{I} is called *partial* if $|\boldsymbol{\pi}| < k$. For a given solution $\boldsymbol{\pi} = (\pi_1, \pi_2, \dots)$ and any round $j \in \{1, \dots, |\boldsymbol{\pi}|\}$, we denote by $\boldsymbol{\pi}_{[j]} = (\pi_1, \dots, \pi_j)$ the solution containing the first j elements of $\boldsymbol{\pi}$. We use the convention that $\boldsymbol{\pi}_{[0]}$ is the empty solution.

A solution $\boldsymbol{\pi}$ for \mathbf{I} is said to be *feasible* if every budget allocation $\pi_j \in \boldsymbol{\pi}$ is feasible for I_j . We similarly say that $\boldsymbol{\pi}$ is *exhaustive* if all budget allocations in $\boldsymbol{\pi}$ are exhaustive for their respective instances. Finally a solution $\boldsymbol{\pi}$ is *non-empty* if none of the budget allocations it contains are empty.

Coming to the agents now, we assume that the set of agents does not change from one round to the next. We thus do not need additional notation and use \mathcal{N} to denote the set of agents. A novelty of this chapter is that agents belongs to different *types*. A type represents any set of characteristics that can be used to group agents together (district of residence, age group...). We denote by \mathcal{T} the set of all the types.⁵³ The *type function* $T : \mathcal{N} \rightarrow \mathcal{T}$ associates each agent $i \in \mathcal{N}$ with the type $T(i)$ they belong to, with the assumption that every agent belongs to a type. For simplicity, we will mainly consider a type $t \in \mathcal{T}$ as the set $\{i \in \mathcal{N} \mid T(i) = t\}$ of the agents belonging to type t , *i.e.*, t 's preimage under T . In that view, $|t|$ denotes the number of agents having type $t \in \mathcal{T}$. Types are constant over time.

Agents submit their opinions about the projects through approval ballots. For a given k -PPB instance $\mathbf{I} = (I_1, \dots, I_k)$, at round $j \in \{1, \dots, k\}$ every agent $i \in \mathcal{N}$ is asked to submit an approval ballot denoted by $A_i^j \subseteq \mathcal{P}_j$. A profile for round $j \in \{1, \dots, k\}$ is a vector $\mathbf{A}^j = (A_1^j, \dots, A_n^j)$ of ballots. A *PPB profile* is a vector $\mathbf{A} = (\mathbf{A}^1, \dots, \mathbf{A}^k)$ of profiles, one per round.⁵⁴













⁵³Note that we are using here the same notation for type of *agents* as the one used in Chapter 6, Section 6.2.4, for types of *projects*. This dual usage of the terminology and notation should not be problematic given that the context is rather clear.

⁵⁴Note that we have to superscript the rounds here since the subscript position is used for the agents.

We conclude this section by an example presenting all the components of the PPB model. This will also be our running example.

Example 7.1.1. Consider a town that is running a PB process. For simplicity, assume that only five inhabitants are active citizens and are participating in the PB process. The latter is planned to run for as long as possible, with one instance being organised each year. A total budget of 30 000 € is dedicated to the PB process each year (we will divide all budget limits and costs by 1 000 to simplify). The rules also stipulate that exactly four projects should advance to the voting stage at the end of shortlisting stage, when the projects that will be voted on are selected (see either Section 1.1.2 or Chapter 8 for more details). We are now at the voting stage of the third year, the ballots have been cast, and a budget allocation has to be selected.

The specifics of this instance are presented below. Twelve projects are considered and the ones selected in years 1 and 2 are boxed in dashed lines. Five voters, including the mayor himself are voting in the process.

		Year 1 $b_1 = 30$				Year 2 $b_2 = 30$				Year 3 $b_3 = 30$			
													
	Cost	18	6	6	12	15	15	9	6	21	21	4	29
West type	Alina	✓			✓		✓	✓	✓	✓		✓	
	Daira	✓		✓			✓		✓	✓		✓	
	Julian	✓			✓		✓	✓	✓	✓		✓	
East type	Dean		✓		✓	✓		✓	✓			✓	✓
	Lwenn			✓	✓	✓		✓				✓	✓

The town is organised in two districts, the east and the west one, each clearly separated from the other. The mayor, Dean, has decided to ensure long-term fairness, grouping the voters by their location. There are thus two types, the east type and the west type. Note that the information presented above allows us to track down, for every voter, the ballot they submitted in each year. This would not be possible in practice because of the anonymity of the ballots but simplifies the presentation.

Let us describe how to capture this information with our notation. We are facing a 3-PPB instance $I = (I_1, I_2, I_3)$ where the set of all projects is $\mathbb{P} = \{p_1, \dots, p_{12}\}$. The cost function c is as described in the above table. The instance corresponding to a given round $j \in \{1, 2, 3\}$ is $I_j = \langle \mathcal{P}_j, c, b_j \rangle$ with $\mathcal{P}_j = \{p_{4(j-1)+1}, \dots, p_{4(j-1)+4}\}$ and

This slight inconsistency of the notation for the rounds in this chapter ensures consistency for the notation of the agents throughout the thesis.

$b_j = 30$. In the first round, the budget allocation $\pi_1 = \{p_1, p_4\}$ was selected; while $\pi_2 = \{p_6, p_7, p_8\}$ was selected for the second round. There are five agents, grouped into two types, namely $\mathcal{T} = \{t_{west}, t_{east}\}$, so that $t_{west} = \{Alina, Daira, Julian\}$ and $t_{east} = \{Dean, Lwenn\}$. The PPB profile $\mathbf{A} = (\mathbf{A}^1, \mathbf{A}^2, \mathbf{A}^3)$ is composed of three profiles. For the first round, we have $\mathbf{A}^1 = (A_{Alina}^1, A_{Daira}^1, A_{Julian}^1, A_{Dean}^1, A_{Lwenn}^1)$, where, for instance, Dean's ballot is $A_{Dean}^1 = \{p_2, p_4\}$. The rest of the ballots and profiles should be clear from the table. \triangle

7.2 A Fairness Theory for Perpetual Participatory Budgeting

As we have seen already on multiple occasions (especially in Chapters 3 and 4), fairness is one of the central theme of the research on PB. The present chapter will make no exception, and fairness will be our main focus in the analysis of the perpetual PB model. In the following, we introduce our fairness theory.

7.2.1 Evaluation Functions

We start by introducing evaluation functions. These are functions that assess the quality of a solution for a given type of agents.

Definition 7.2.1 (Evaluation Function). *An evaluation function Φ is a function taking as inputs a k -PPB instance \mathbf{I} , a PPB profile \mathbf{A} , a (partial) solution π , and a type $t \in \mathcal{T}$, and returning $\Phi(\mathbf{I}, \mathbf{A}, \pi, t) \in \mathbb{R}_{\geq 0}$, an evaluation of the quality of the solution π for type t . Moreover, for a round $j \in \{1, \dots, |\pi|\}$ the marginal evaluation function is defined as:*

$$\Phi_{marg}(\mathbf{I}, \mathbf{A}, \pi, t, j) = \Phi(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t) - \Phi(\mathbf{I}, \mathbf{A}, \pi_{[j-1]}, t).$$

Given everything we have seen so far, the reader probably expects evaluation functions to mimic satisfaction functions, but for types of agents (satisfaction functions only apply to individuals). This intuition is correct, but only partially. An evaluation function is more general and does not need to be about satisfaction. It can also be about the influence of a type (as in power indices), for instance, or any other measure that is deemed relevant.

In the following, we will introduce three evaluation functions, based either on the cost satisfaction function sat^{cost} , on its relative version $relsat_{sat^{cost}}$, or on the share $share$ (that we interpret as an influence measure and not a satisfaction function). The relative satisfaction $relsat_{sat^{cost}}$ will be defined in the following. Both sat^{cost} and $share$ have already been introduced at an earlier point of this thesis, though for individuals and not for types. They will then be reframed as evaluation functions.

We start with the cost satisfaction function sat^{cost} , which we reinterpret as a evaluation function for the PPB model.

Definition 7.2.2 (Cost Evaluation Function). *Let $\mathbf{I} = (I_1, \dots, I_k)$ be a k -PPB instance, $\mathbf{A} = (A^1, \dots, A^k)$ a PPB profile, and $\boldsymbol{\pi} = (\pi_1, \pi_2, \dots)$ a (partial) solution for \mathbf{I} , with $|\boldsymbol{\pi}| \leq k$. The cost evaluation function Φ^{cost} of a given type $t \in \mathcal{T}$ is defined as:*

$$\Phi^{cost}(\mathbf{I}, \mathbf{A}, \boldsymbol{\pi}, t) = \sum_{j=1}^{|\boldsymbol{\pi}|} \frac{1}{|t|} \sum_{i \in t} sat^{cost}(\pi_j \cap A_i^j) = \sum_{j=1}^{|\boldsymbol{\pi}|} \frac{1}{|t|} \sum_{i \in t} c(\pi_j \cap A_i^j).$$

For a given round $j \in \{1, \dots, |\boldsymbol{\pi}|\}$, the marginal cost evaluation function is:

$$\Phi_{marg}^{cost}(\mathbf{I}, \mathbf{A}, \boldsymbol{\pi}, t, j) = \frac{1}{|t|} \sum_{i \in t} sat^{cost}(\pi_j \cap A_i^j).$$

So, according to the cost evaluation function, the quality of a solution for a given type $t \in \mathcal{T}$ is measured as the average cost satisfaction of the members of t , summed up over all rounds.

Let us illustrate this evaluation function on our running example.

Example 7.2.3. Consider the PPB instance \mathbf{I} and PPB profile \mathbf{A} that have been described in Example 7.1.1. Remember that we had $\pi_1 = \{p_1, p_4\}$ and $\pi_2 = \{p_6, p_7, p_8\}$. At the end of the second round, the cost evaluation function would produce the following assessments for the two types:

$$\begin{aligned} \Phi^{cost}(\mathbf{I}, \mathbf{A}, (\pi_1, \pi_2), t_{west}) &= \sum_{j=1}^2 \frac{1}{3} (sat_{Alina}^{cost}(\pi_j) + sat_{Daira}^{cost}(\pi_j) + sat_{Julian}^{cost}(\pi_j)) \\ &= \frac{30 + 18 + 30}{3} + \frac{30 + 21 + 30}{3} = 53, \\ \Phi^{cost}(\mathbf{I}, \mathbf{A}, (\pi_1, \pi_2), t_{east}) &= \sum_{j=1}^2 \frac{1}{2} (sat_{Dean}^{cost}(\pi_j) + sat_{Lwenn}^{cost}(\pi_j)) \\ &= \frac{12 + 12}{2} + \frac{15 + 9}{2} = 24. \end{aligned}$$

Anticipating on what is to come, the solution (π_1, π_2) can thus be deemed unfair given the huge discrepancy between the evaluations of the two types. \triangle

One potential drawback of the cost satisfaction function is its strong dependence on the size of voters' approval sets. For example, if agent 1 approves a proper subset of agent 2's approved projects (if we have $A_1^j \subset A_2^j$) and all their approved projects are funded (if $A_1^j \subset A_2^j \subseteq \pi_j$), then agent 2 is considered more satisfied than agent 1. However, it can be argued that the welfare of both agents should be equal as all projects they wanted to be funded have actually been funded; neither agent 1 or 2 can be made happier (subject to the available information). This motivates us to define the relative cost satisfaction of a voter.

The *relative satisfaction* of any satisfaction function sat normalises the satisfaction of a voter by the maximum satisfaction achievable. It is defined as:⁵⁵

$$relsat_{sat}(P) = \frac{sat(P)}{\max\{sat(P') \mid P' \in \text{FEAS}(I) \text{ and } P' \subseteq A_i\}}.$$

It is interesting to note that this does not define a satisfaction function as defined in Definition 2.2.1 since it depends on the full approval ballot of the voter, and not only the approved and selected projects. From this definition, we derive the relative cost evaluation function, defined below.

Definition 7.2.4 (Relative Cost Evaluation Function). *Let $\mathbf{I} = (I_1, \dots, I_k)$ be a k -PPB instance, $\mathbf{A} = (\mathbf{A}^1, \dots, \mathbf{A}^k)$ a PPB profile, and $\boldsymbol{\pi} = (\pi_1, \pi_2, \dots)$ a (partial) solution for \mathbf{I} , with $|\boldsymbol{\pi}| \leq k$. The relative cost evaluation function $\Phi^{relcost}$ of a given type $t \in \mathcal{T}$ is thus defined as:*

$$\begin{aligned} \Phi^{relcost}(\mathbf{I}, \mathbf{A}, \boldsymbol{\pi}, t) &= \sum_{j=1}^{|\boldsymbol{\pi}|} \frac{1}{|t|} \sum_{i \in t} relsat_{sat^{cost}}(\pi_j \cap A_i^j) \\ &= \sum_{j=1}^{|\boldsymbol{\pi}|} \frac{1}{|t|} \sum_{i \in t} \frac{c(\pi_j \cap A_i^j)}{\max\{c(P) \mid P \in \text{FEAS}(I_j) \text{ and } P \subseteq A_i^j\}}. \end{aligned}$$

For a given round $j \in \{1, \dots, |\boldsymbol{\pi}|\}$, the marginal relative cost evaluation function is:

$$\Phi_{marg}^{relcost}(\mathbf{I}, \mathbf{A}, \boldsymbol{\pi}, t, j) = \frac{1}{|t|} \sum_{i \in t} relsat_{sat^{cost}}(\pi_j \cap A_i^j).$$

Thus, when using the relative cost evaluation function, a type assesses a solution as the average relative cost satisfaction of its members, where the latter is the proportion of the best case (highest cost satisfaction achievable) that is actually achieved.

Once again, let us exemplify this evaluation function on our running example.

Example 7.2.5. Let us consider \mathbf{I} , \mathbf{A} , π_1 , and π_2 as defined in Example 7.1.1. At the end of the second round, we have:

$$\begin{aligned} \Phi^{relcost}(\mathbf{I}, \mathbf{A}, (\pi_1, \pi_2), t_{west}) &= \frac{1 + 18/24 + 1}{3} + \frac{1 + 1 + 1}{3} = \frac{23}{12} \approx 1.917, \\ \Phi^{relcost}(\mathbf{I}, \mathbf{A}, (\pi_1, \pi_2), t_{east}) &= \frac{12/18 + 12/18}{2} + \frac{15/30 + 9/24}{2} = \frac{53}{48} \approx 1.104. \end{aligned}$$

Interestingly, the evaluation changed quite a lot when using $\Phi^{relcost}$ compared to Φ^{cost} . Type t_{west} now evaluates solution (π_1, π_2) only 74% higher than type t_{east} (it was 120% higher with Φ^{cost}). The solution is then still somewhat unfair, but not as much as it may have seemed initially. \triangle

⁵⁵Note here that the relative cost evaluation as introduced in this chapter is similar in spirit to the *normalised satisfaction* of Chapter 5 (see Section 5.3.1), though radically different in terms of behaviour.

We finally consider an evaluation function based on the share. A whole chapter has been dedicated to the study of fairness criteria based on the share (Chapter 4). We use it here as a measure of the influence of the types on a solution.

Definition 7.2.6 (Share Evaluation Function). *Let $\mathbf{I} = (I_1, \dots, I_k)$ be a k -PPB instance, $\mathbf{A} = (\mathbf{A}^1, \dots, \mathbf{A}^k)$ a PPB profile, and $\boldsymbol{\pi} = (\pi_1, \pi_2, \dots)$ a (partial) solution for \mathbf{I} , with $|\boldsymbol{\pi}| \leq k$. The share evaluation function Φ^{share} of a given type $t \in \mathcal{T}$ is:*

$$\begin{aligned} \Phi^{share}(\mathbf{I}, \mathbf{A}, \boldsymbol{\pi}, t) &= \sum_{j=1}^{|\boldsymbol{\pi}|} \frac{1}{|t|} \sum_{i \in t} share(I_j, \mathbf{A}^j, \pi_j, A_i^j) \\ &= \sum_{j=1}^{|\boldsymbol{\pi}|} \frac{1}{|t|} \sum_{i \in t} \sum_{p \in \pi^j \cap A_i^j} \frac{c(p)}{|\{i' \in \mathcal{N} \mid p \in A_{i'}^j\}|}. \end{aligned}$$

For a given round $j \in \{1, \dots, |\boldsymbol{\pi}|\}$, the marginal share evaluation function is:

$$\Phi_{marg}^{share}(\mathbf{I}, \mathbf{A}, \boldsymbol{\pi}, t, j) = \frac{1}{|t|} \sum_{i \in t} share(I_j, \mathbf{A}^j, \pi_j, A_i^j).$$

Analogously to Φ^{cost} and $\Phi^{relcost}$, the share evaluation function sums up the average share of the members of a type, over all rounds.

We now analyse Example 7.1.1 through the lens of the share evaluation function.

Example 7.2.7. We consider once again \mathbf{I} , \mathbf{A} , π_1 , and π_2 as defined in Example 7.1.1. For the solution (π_1, π_2) , we have:

$$\begin{aligned} \Phi^{share}(\mathbf{I}, \mathbf{A}, (\pi_1, \pi_2), t_{west}) &= \frac{9 + 6 + 9}{3} + \frac{35/4 + 13/2 + 35/4}{3} = 16, \\ \Phi^{share}(\mathbf{I}, \mathbf{A}, (\pi_1, \pi_2), t_{east}) &= \frac{3 + 3}{2} + \frac{15/4 + 9/4}{2} = 6. \end{aligned}$$

Interestingly, according to Φ^{share} , the two types have significantly different evaluations, with type t_{west} having more influence in the solution (π_1, π_2) than t_{east} . This is the largest relative difference between the evaluation of the two types across the three evaluation functions we consider. \triangle

We have briefly mentioned fairness in the examples above. In the following section we will define everything more formally.

7.2.2 Fairness Criteria

The foundation of our fairness theory is that the fairest solution is one in which all types evaluate the solution the same way. This is our first fairness criterion.

Definition 7.2.8 (EQUAL- Φ). Given an evaluation function Φ , a solution π for the k -PPB instance \mathbf{I} such that $|\pi| \leq k$ and the PPB profile \mathbf{A} satisfies EQUAL- Φ at round $j \in \{1, \dots, |\pi|\}$ if for every two types $t, t' \in \mathcal{T}$, we have:

$$\Phi(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t) = \Phi(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t').$$

We will also say that π satisfies EQUAL- Φ if it satisfies EQUAL- Φ at round j for all rounds $j \in \{1, \dots, |\pi|\}$.

In the case of Φ^{share} , it should be noted that enforcing EQUAL- Φ^{share} is very similar to requiring all agents to reach their fair share, as defined in Chapter 4. It is not—strictly speaking—equivalent. In the definition of the fair share we corrected for the fact that agents may not be able to reach a share of b/n , which we do not do in the definition of EQUAL- Φ^{share} . However, this connection still entails that in any 1-PPB instance in which all agents belong to their own type (this corresponds to a regular instance of the standard PB model), a solution satisfying EQUAL- Φ^{share} corresponds to a budget allocation satisfying fair share. This implies, in particular, that EQUAL- Φ^{share} cannot always be guaranteed as fair share cannot always be (see Proposition 4.2.2). We also expect that the computational problems that are hard for the fair share will also be hard for EQUAL- Φ^{share} , though this cannot be turned into a formal statement.

In the illustrating examples from the previous section, we already hinted at the definition of EQUAL- Φ . Let us get back to them now that it has been properly defined.

Example 7.2.9. We are still considering \mathbf{I} , \mathbf{A} , π_1 and π_2 as defined in Example 7.1.1. The question we are facing now is: Given the ballots of the third round, can we find π_3 such that (π_1, π_2, π_3) satisfies EQUAL- Φ ?

Remember from Example 7.2.3 that at the end of the second round, the cost evaluation of (π_1, π_2) was of 53 for t_{west} and of 24 for t_{east} . Now, one can check that selecting $\pi_3 = \{p_{12}\}$ would make (π_1, π_2, π_3) satisfy EQUAL- Φ^{cost} . Both types would evaluate the solution at 53. Similarly, selecting $\pi'_3 = \{p_{10}, p_{11}\}$ in the third round would lead to a solution (π_1, π_2, π'_3) that satisfies EQUAL- Φ^{share} . Both types would then have a share evaluation of 17.

Interestingly, there is no $\pi''_3 \subseteq \{p_9, p_{10}, p_{11}, p_{12}\}$ for which the solution (π_1, π_2, π''_3) would satisfy EQUAL- $\Phi^{relcost}$. △

EQUAL- Φ implements a strict equality approach to fairness. It is well known, at least from the literature on social choice, that perfect equality can rarely be guaranteed in practice. Partially disclosing some of the results that will come, this will also be observed for EQUAL- Φ . Because of that, we introduce two relaxations.

When perfect fairness cannot be achieved, one can try to get as close to it as possible. This optimisation-based approach is the first relaxation of EQUAL- Φ we introduce. This idea is particularly relevant in a long-term perspective as subsequent rounds can be used to compensate for unfairness in previous rounds. We pursue this approach by

considering the Gini coefficient (Gini, 1912) of a solution—a well-known measure of inequality given a multi-set of values—that can be used as a minimisation objective. In the following, we use the standard formulation (Blackorby and Donaldson, 1978).⁵⁶

Definition 7.2.10 (Φ -GINI). *Let $\mathbf{v} = (v_1, \dots, v_q) \in \mathbb{R}_{\geq 0}^q$ be a vector of non-negative values ordered in non-increasing order, i.e., such that $v_i \geq v_j$ for all $1 \leq i \leq j \leq q$. The Gini coefficient of \mathbf{v} is given by:*

$$\text{GINI}(\mathbf{v}) = 1 - \frac{\sum_{i=1}^q (2i-1)v_i}{q \sum_{i=1}^q v_i}.$$

For an evaluation function Φ , the Φ -GINI coefficient of a solution π for the k -PPB instance \mathbf{I} and of the PPB profile \mathbf{A} at round $j \in \{1, \dots, |\pi|\}$ is then:

$$\text{GINI}_{\Phi}(\mathbf{I}, \mathbf{A}, \pi_{[j]}) = \text{GINI}(\Phi^{\downarrow}(\mathbf{I}, \mathbf{A}, \pi_{[j]})),$$

where $\Phi^{\downarrow}(\mathbf{I}, \mathbf{A}, \pi_{[j]})$ is a vector containing $\Phi(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t)$ for all types $t \in \mathcal{T}$, ordered in non-increasing order.

A solution π satisfies Φ -GINI at round j with respect to a set \mathfrak{S} of solutions for \mathbf{I} , if there is no solution $\pi' \in \mathfrak{S}$ with $\text{GINI}_{\Phi}(\mathbf{I}, \mathbf{A}, \pi'_{[j]}) < \text{GINI}_{\Phi}(\mathbf{I}, \mathbf{A}, \pi_{[j]})$.

Informally, the Gini coefficient measures the distance between a set of values, and an optimally fair distribution of the values. It has a very visual intuition. For a given solution π , let $\bar{\Phi} = \sum_{t \in \mathcal{T}} \Phi(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t)$, and define the Lorenz curve as the curve of the cumulative distribution, i.e., the one that plots for any $x \in [0, 1]$, the proportion of $\bar{\Phi}$ that is achieved by the $x\%$ worst off types. Then, the Φ -GINI coefficient measures the distance between the diagonal (the case of perfect equality), and the Lorenz curve. The above definition is a discrete version of this idea.

We state couple of interesting facts about the Gini coefficient that will prove useful in the coming proofs. Proving these simple facts could be a good exercise for a reader interested in understanding the Gini coefficient in more depth.

Fact 7.2.11. *For any vector $\mathbf{v} = (v_1, \dots, v_q) \in \mathbb{R}_{\geq 0}^q$ ordered non-increasingly, we have:*

$$0 \leq \text{GINI}(\mathbf{v}) \leq 1.$$

Moreover, we have $\text{GINI}(\mathbf{v}) = 0$ if and only if $v_1 = \dots = v_q$.

Fact 7.2.12. *Let $d \in \mathbb{R}_{\geq 0}$ and $x \in \mathbb{R}_{\geq 0}$. For the two-dimensional vector $\mathbf{v}_{x,d} = (x + d, x)$, we have:*

$$\frac{\delta \text{GINI}(\mathbf{v}_{x,d})}{\delta x} \leq 0 \quad \text{and} \quad \frac{\delta \text{GINI}(\mathbf{v}_{x,d})}{\delta d} \geq 0,$$

that is, $\text{GINI}((x + d, x))$ is a decreasing function in x , and an increasing function in d .

⁵⁶Blackorby and Donaldson (1978) can be difficult to read. For a definition of the Gini coefficient that closely resembles ours, see the definition given, in French, in Sylvain Bouveret's PhD thesis (Bouveret, 2007, page 39). Keeping in mind that in the notation used there, \bar{u} denotes the average value, and u^{\uparrow} is ordered non-decreasingly, one can see that the definitions are the same (up to reversing the ordering of the vector of values so that it is ordered non-increasingly).

In particular, Fact 7.2.11 implies that Φ -GINI indeed is a relaxation of EQUAL- Φ : for every evaluation function Φ , a solution π satisfies EQUAL- Φ if and only if its Φ -GINI coefficient reaches 0, its minimum.

Let us explore our running example in the light of Φ -GINI.

Example 7.2.13. Let us go back to the instance and profile described in Example 7.1.1. We already know from Example 7.2.9 that we can select projects in the third round such that either EQUAL- Φ^{cost} is satisfied. The corresponding solution would thus have a Φ^{cost} -GINI coefficient of 0, the optimum value. The same is true for Φ^{share} .

Let us thus focus on $\Phi^{relcost}$. Remember that we could not find a set of projects for the third round that would lead to a solution satisfying EQUAL- $\Phi^{relcost}$. Still, we can find a budget allocation π_3 such that (π_1, π_2, π_3) satisfies $\Phi^{relcost}$ -GINI, i.e., has an optimal $\Phi^{relcost}$ -GINI coefficient. Consider $\pi_3 = \{p_{12}\}$. Then, we have:

$$\begin{aligned}\Phi^{relcost}(\mathbf{I}, \mathbf{A}, (\pi_1, \pi_2, \pi_3), t_{west}) &= \frac{1 + 18/24 + 1}{3} + \frac{1 + 1 + 1}{3} + 0 = \frac{23}{12}, \\ \Phi^{relcost}(\mathbf{I}, \mathbf{A}, (\pi_1, \pi_2, \pi_3), t_{east}) &= \frac{12/18 + 12/18}{2} + \frac{15/30 + 9/24}{2} + \frac{1 + 1}{2} = \frac{101}{48}.\end{aligned}$$

We thus have:

$$\text{GINI}_{\Phi^{relcost}}(\mathbf{I}, \mathbf{A}, (\pi_1, \pi_2, \pi_3)) = 1 - \frac{3 \cdot 23/12 + 101/48}{2(23/12 + 101/48)} = \frac{9}{386} \approx 0.023.$$

Now, for $\pi'_3 = \{p_{11}\}$, we get:

$$\text{GINI}_{\Phi^{relcost}}(\mathbf{I}, \mathbf{A}, (\pi_1, \pi_2, \pi'_3)) = 1 - \frac{3 \cdot 2545/1392 + 23/12}{2(23/12 + 2545/1392)} = \frac{123}{10426} \approx 0.011.$$

The other budget allocations would not lead to a better $\Phi^{relcost}$ -GINI coefficient as the difference between the evaluation of t_{west} and t_{east} would only grow larger. In this example, $\pi = (\pi_1, \pi_2, \pi'_3)$ satisfies $\Phi^{relcost}$ -GINI with respect to the set of solutions in which π_1 and π_2 were selected in the first two rounds. \triangle

Another approach to circumvent the problem that perfect fairness cannot always be achieved is to require perfect fairness, but only in the long run. This approach is perfectly suitable with a long-term perspective, and captures some of the motivation we detailed in the introduction. We introduce below EQUAL- Φ -CONV which formalizes the idea of asymptotically equalising the evaluations of the solution for the different types.

Definition 7.2.14 (EQUAL- Φ -CONV). Given an evaluation function Φ , an infinite solution π for the ∞ -PPB instance \mathbf{I} and the infinite PPB profile \mathbf{A} satisfies EQUAL- Φ -CONV if for every two types t and $t' \in \mathcal{T}$, the following holds:

$$\frac{\Phi(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t)}{\Phi(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t')} \xrightarrow{j \rightarrow +\infty} 1.$$

It is important to note that the above definition only applies to infinite instances, profiles, and solutions as it relies on the notion of convergence. The impact of this technical requirement is however limited. Indeed, we can for instance create an infinite PPB-instance from a finite one by simply repeating the original instance infinitely many times.

Once again, this indeed relaxes EQUAL- Φ in the sense that if a solution satisfies EQUAL- Φ in all rounds, it would trivially satisfy EQUAL- Φ -CONV as well.

Now that all the elements of our fairness theory for the PPB model are in place, we turn to its formal investigation.

7.3 Achieving Perfect Fairness: EQUAL- Φ

We first explore the criterion that represents a situation of perfect fairness: EQUAL- Φ . We will first focus on existence guarantees, and then shift to computational problems.

In the light of what we observed when trying to provide every agent with their fair share, we already know that EQUAL- Φ^{share} cannot always be guaranteed. It is thus not surprising that this is also the case for EQUAL- Φ^{cost} and EQUAL- $\Phi^{relcost}$, independently of the number of rounds we consider.

Proposition 7.3.1. *For any $k \in \mathbb{N}_{>0}$, there exists a k -PPB instance \mathbf{I} and a PPB profile \mathbf{A} such that no non-empty solution π for \mathbf{I} satisfies EQUAL- Φ^{cost} , EQUAL- $\Phi^{relcost}$, or EQUAL- Φ^{share} . This holds even in the unit-cost setting and if all types are of size 1.*

PROOF. Let $\mathbf{I} = (I_1, \dots, I_k)$ be a k -PPB instance, for a given $k \in \mathbb{N}_{>0}$. The instance $I_j = \langle \mathcal{P}_j, c, b_j \rangle$ at round $j \in \{1, \dots, k\}$ is such that $P_j = \{p_1^j, p_2^j\}$ and $b_j = 1$ for any round $j \in \{1, \dots, k\}$. We thus have $\mathbb{P} = P_1 \cup \dots \cup P_k = \{p_1^1, p_2^1, \dots, p_1^k, p_2^k\}$. The cost function is such that $c(p) = 1$ for all projects $p \in \mathbb{P}$.

Consider a PPB profile with two agents 1 and 2, of types t_1 and t_2 respectively. In the first round, agent 1 approves only of project p_1^1 and agent 2 only of p_2^1 . In any subsequent round $j \in \{2, \dots, k\}$, both agents only approve of p_1^j .

Assume without loss of generality that p_1^1 is selected in the first round, *i.e.* $\pi_1 = \{p_1^1\}$. Then, for any round $j \in \{1, \dots, k\}$ we have:

$$\Phi(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t_1) = 1 + \Phi(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t_2),$$

for any evaluation function $\Phi \in \{\Phi^{cost}, \Phi^{relcost}, \Phi^{share}\}$. This is a clear violation of EQUAL- Φ . \square

The fact that EQUAL- Φ cannot be satisfied in general for our evaluation functions does not imply that it is *never* satisfiable. In the following, we investigate the computational complexity of checking whether EQUAL- Φ can be satisfied for a given in-

stance. Uninspired, we call this computational problem EQUAL- Φ SATISFIABILITY. It is defined as follows.

EQUAL- Φ SATISFIABILITY	
Input:	A k -PPB instance $\mathbf{I} = (I_1, \dots, I_k)$, a PPB profile \mathbf{A} and a solution $\boldsymbol{\pi} = (\pi_1, \dots, \pi_{k-1})$ for (I_1, \dots, I_{k-1}) .
Question:	Is there a non-empty and feasible $\pi_k \in \text{FEAS}(I_k) \setminus \{\emptyset\}$ such that the solution $\boldsymbol{\pi}' = (\pi_1, \dots, \pi_{k-1}, \pi_k)$ satisfies EQUAL- Φ ?

If $k = 1$ in the input, we assume $\boldsymbol{\pi}$ to be the empty solution.

As we will prove shortly, it turns out that for all three welfare measures, checking the existence of an EQUAL- Φ solution is an NP-complete problem. Importantly, this hardness result implies that for our evaluation functions, there cannot be an efficient algorithm that always returns a solution satisfying EQUAL- Φ when one exists, unless $P = NP$. We first prove our claim for Φ^{cost} and Φ^{relcost} .

Proposition 7.3.2. *The EQUAL- Φ^{cost} SATISFIABILITY and EQUAL- Φ^{relcost} SATISFIABILITY problems are strongly NP-complete, even in the unit-cost setting and if there is only one round.*

PROOF. Membership in NP is clear, the certificate being the solution itself. To prove NP-hardness, we will reduce from the problem ONE-IN-THREE 3-SAT known to be strongly NP-hard (Garey and Johnson, 1979; Schaefer, 1978).

ONE-IN-THREE 3-SAT	
Input:	A propositional formula φ in conjunctive normal form with exactly three literals per clause (3-CNF).
Question:	Is there a truth assignment α for φ so that each clause in φ has exactly one literal set to true?

Consider a 3-CNF formula φ . Denote by $\text{Var}(\varphi)$ the set of propositional variables appearing in φ , and by $\text{Clause}(\varphi)$ the set of clauses of φ . We construct a 1-PPB instance $\mathbf{I} = (I_1)$ where $I_1 = \langle \mathcal{P}, c, b \rangle$ as follows. The set of projects is

$$\mathcal{P} = \bigcup_{x \in \text{Var}(\varphi)} \{p_x, p_{\neg x}\}.$$

All projects have cost 1 and the budget limit is $b = |\text{Var}(\varphi)|$.

Let us turn to the PPB profile $\mathbf{A} = (\mathbf{A}^1)$ now. For each propositional variable $x \in \text{Var}(\varphi)$, there is an agent i_x approving of both p_x and $p_{\neg x}$. Moreover, for each clause $c \in \text{Clause}(\varphi)$, there is an agent i_c approving of the three projects

corresponding to the literals in c . Every agent belongs to a unique type and is the only one belonging to that type.

We claim that there exists a truth assignment for φ that sets exactly one literal to true in every clause of φ if and only if there exists a non-empty and feasible solution π for I that provides EQUAL- Φ^{cost} . Indeed, since π has to be non-empty, at least one “variable” agent i_x , for $x \in Var(\varphi)$, will have a cost-satisfaction of 1. Thus, to satisfy EQUAL- Φ^{cost} , this implies that every agent should have a cost-satisfaction of at least 1. We show that it also cannot be more than 1. Note that the approval ballots of the “variable” agents are all disjoint. Since the budget limit is $b = |Var(\varphi)|$ and there are $|Var(\varphi)|$ “variable” agents, EQUAL- Φ^{cost} cannot thus be satisfied if an agent has a cost satisfaction above 1. Overall, EQUAL- Φ^{cost} in I is equivalent to every agent having cost-satisfaction 1. Satisfying the latter is equivalent to selecting exactly one project among p_x and $p_{\neg x}$ for every $x \in Var(\varphi)$, and exactly one project among the ones corresponding to the literals in c for all $c \in Clause(\varphi)$. Call π such a budget allocation. Consider then the truth assignment α that sets a propositional variable $x \in Var(\varphi)$ to true (respectively false) if and only if p_x (respectively $p_{\neg x}$) has been selected in π . Since π has to be feasible, it is clear that α is a suitable truth assignment for the ONE-IN-THREE 3-SAT problem if and only if the solution $\pi = (\pi)$ for I exists.

Since this reduction can be done in polynomial time, the proof is complete for EQUAL- Φ^{cost} .

We now describe how to adapt the reduction for EQUAL- $\Phi^{relcost}$. The idea here is to add one project p^* that is approved by every “variable” agent i_x for $x \in Var(\varphi)$. That way, every agent $i \in \mathcal{N}$ submits an approval ballot of length 3. From there, it should be clear that there exist a suitable truth assignment for φ if and only if there is a solution π in which all agents have a relative cost-satisfaction of $1/3$. Indeed, because of the budget limit, to satisfy EQUAL- $\Phi^{relcost}$, a solution π needs to provide every agent with a relative cost-satisfaction of $1/3$. This cannot be done by selecting p^* because doing so would provide every “variable” agent a relative cost-satisfaction of $1/3$, and there are no ways to match this $1/3$ for the “clause” agents without increasing the relative cost-satisfaction of some “variable” agents (since $Var(\varphi)$ contains all the variables appearing in a clause). \square

When it comes to EQUAL- Φ^{share} , one could be tempted to simply refer to the hardness result we showed related to fair share in Chapter 4 (Proposition 4.2.3 to be precise). This is however not formally correct because of the minimum operator used in the definition of the fair share. Still, the same statement holds. We prove weak NP-completeness below, strong NP-hardness has been shown by Klein Goldewijk (2022).

Proposition 7.3.3. *The EQUAL- Φ^{share} SATISFIABILITY problem is weakly NP-complete, even if there is only one round and two agents.*

PROOF. The problem is in NP since checking that all agents have the same share simply requires to compute all the shares, which can be done in polynomial time.

We show the NP-hardness via a reduction from the SUBSET-SUM problem, known to be weakly NP-hard (Karp, 1972; Garey and Johnson, 1979).

SUBSET-SUM

Input: A finite set $Z \subseteq \mathbb{Z} \setminus \{0\}$.

Question: Is there a non-empty $Z' \subseteq Z$ such that $\sum_{z \in Z'} z = 0$?

Given a suitable set $Z \subset \mathbb{Z}$, we construct the following 1-PPB instance $I = (I_1)$ where $I_1 = \langle \mathcal{P}, c, b \rangle$. The budget limit is $b = \sum_{z \in Z} |z|$. For every $z \in Z$, there is a corresponding project $p_z \in \mathcal{P}$ with cost $|z|$. There are also $b - \min_{z \in Z} |z|$ additional projects, denoted by p_j^* for $j \in \{1, \dots, b - \min_{z \in Z} |z|\}$, all of cost 1. We also construct a PPB profile $\mathbf{A} = (\mathbf{A}^1)$ with two agents 1 and 2, both belonging to a different type: $1 \in t_+$ and $2 \in t_-$. In the first, and only, round, agent 1 approves of project all projects p_z such that $z > 0$, and agent 2 approves of all p_z for which $z < 0$. Furthermore, both agents approve of all projects p_j^* for $j \in \{1, \dots, b - \min_{z \in Z} |z|\}$. We denote the ballot of agent 1 by A_1 , and that of agent 2 by A_2 .

We claim that for any $Z' \subseteq Z$, we have $\sum_{z \in Z'} z = 0$ if and only if the solution defined by $\pi = (\pi)$ with $\pi = \{p_z \mid z \in Z'\} \cup \{p_1^*, \dots, p_{b - \sum_{z \in Z'} |z|}\}$ satisfies $\text{EQUAL-}\Phi^{\text{share}}$. Note that for π we have:

$$c(A_1 \cap \pi) = \sum_{\substack{z \in Z' \\ z > 0}} |z| + b - \sum_{z \in Z'} |z|,$$

$$c(A_2 \cap \pi) = \sum_{\substack{z \in Z' \\ z < 0}} |z| + b - \sum_{z \in Z'} |z|.$$

So $c(A_1 \cap \pi) = c(A_2 \cap \pi)$ if and only if $\sum_{z \in Z', z > 0} |z| = \sum_{z \in Z', z < 0} |z|$. Since the ballots are disjoint, we have $c(A_i \cap \pi) = \text{share}_i(\pi)$ for $i \in \{1, 2\}$. Therefore, π satisfies $\text{EQUAL-}\Phi^{\text{share}}$ if and only if $\sum_{z \in Z'} z = 0$. \square

From a parametrised complexity perspective, it is also interesting to observe that the above reduction also applies to $\text{EQUAL-}\Phi^{\text{cost}}$, showing hardness when there are only two agents and one round. This holds because each project is only approved by a single agent, meaning that Φ^{cost} and Φ^{share} coincide.

Interestingly, for the proofs of both Proposition 7.3.2 and Proposition 7.3.3 we made sure to only use exhaustive budget allocations. They thus also show that checking whether there is an *exhaustive* solution that satisfies $\text{EQUAL-}\Phi^{\text{cost}}$, $\text{EQUAL-}\Phi^{\text{relcost}}$ or $\text{EQUAL-}\Phi^{\text{share}}$ is NP-complete. The computational problem is defined as follows.

EXHAUSTIVE EQUAL- Φ SATISFIABILITY

Input:	A k -PPB instance $\mathbf{I} = (I_1, \dots, I_k)$, a PPB profile \mathbf{A} and a solution $\boldsymbol{\pi} = (\pi_1, \dots, \pi_{k-1})$ for (I_1, \dots, I_{k-1}) that is exhaustive.
Question:	Is there a feasible and exhaustive budget allocation $\pi_k \in \text{FEAS}_{\text{EX}}(I_k)$ such that the solution $\boldsymbol{\pi} = (\pi_1, \dots, \pi_{k-1}, \pi_k)$ satisfies EQUAL- Φ ?

Corollary 7.3.4. *Both the problem EXHAUSTIVE EQUAL- Φ^{cost} SATISFIABILITY and the problem EXHAUSTIVE EQUAL- Φ^{relcost} SATISFIABILITY are strongly NP-complete. Moreover, the problem EXHAUSTIVE EQUAL- Φ^{share} SATISFIABILITY is weakly NP-complete.*

PROOF. The proof directly follows from the observation that the proofs of proposition 7.3.2 and 7.3.3 only make use of exhaustive budget allocations. \square

Our analysis of EQUAL- Φ is now complete. We will turn to Φ -GINI.

7.4 Optimising for Fairness: Φ -GINI

Let us now turn our attention to Φ -GINI, the criterion according to which we optimise for fairness.

By definition, for every instance, there always is a solution that satisfies Φ -GINI, for any evaluation function Φ . Indeed, since a solution satisfies Φ -GINI if it achieves an optimal (*i.e.*, minimal) Φ -GINI coefficient, and there are finitely many solutions, there will always be one with a lowest Φ -GINI coefficient. Therefore, the main questions here concern computational problems, and not existence guarantees. We will distinguish two cases, depending on whether exhaustiveness is required or not.

Since this section is rather long and technical, we already disclose the main conclusion of the results proven below: checking whether a solution satisfies Φ -GINI is a hard problem—coNP-hard to be precise—for all three evaluation functions we defined.

7.4.1 Among all Solutions

At first, we do not impose any specific requirements on the solution (except for non-emptiness). The computational problem we are interested in is the following.

 Φ -GINI SATISFACTION

Input:	A k -PPB instance $\mathbf{I} = (I_1, \dots, I_k)$, a PPB profile \mathbf{A} , and a solution $\boldsymbol{\pi} = (\pi_1, \dots, \pi_k)$ for \mathbf{I} .
Question:	Does $\boldsymbol{\pi}$ satisfy Φ -GINI with respect to \mathfrak{S} , the set of all the solutions $\boldsymbol{\pi}'$ such that every $\pi_j \in \boldsymbol{\pi}'$ is non-empty and feasible for I_j ?

Note that proving that this problem is hard for a given Φ implies that there is no hope for efficiently computing a solution that satisfies Φ -GINI, if the Φ -GINI-coefficient of a solution can be efficiently computed (which is the case for all three evaluation functions of interest here).

We first show that Φ -GINI is coNP-complete for both cost-satisfaction and share.

Proposition 7.4.1. *The problems Φ^{cost} -GINI SATISFACTION and Φ^{share} -GINI SATISFACTION are weakly coNP-complete, even if there are only two agents and one round.*

PROOF. Let us start with Φ^{cost} -GINI SATISFACTION. We show that it is coNP-complete by showing that its co-problem—checking whether a solution does not satisfy Φ^{cost} -GINI, *i.e.*, does not have an optimum Φ^{cost} -GINI coefficient—is NP-complete. We call this problem Φ^{cost} -GINI VIOLATION.

It is clear that Φ^{cost} -GINI VIOLATION is in NP, as we can just guess a non-empty and feasible solution and check if it has a lower Φ^{cost} -GINI coefficient than the input solution π in polynomial time. Let us now prove that it is NP-hard. We do so by a reduction from the SUBSET-SUM problem (see the proof of Proposition 7.3.3 for the definition).

Let $Z = \{z_1, z_2, \dots\}$ be a SUBSET-SUM instance. We construct a Φ^{cost} -GINI VIOLATION instance as follows. We consider a 1-PPB instance $I = (I)$ where $I = \langle \mathcal{P}, c, b \rangle$. The set of projects is $\mathcal{P} = \{p_1, \dots, p_{|Z|}, p_+, p_-\}$. It is such that:

- For every element $z_i \in Z$, there is a project p_i with $c(p_i) = 4|z_i|$;
- There are two additional projects, project p_+ with $c(p_+) = 8 \sum_{z \in Z} |z|$, and project p_- with $c(p_-) = 8 \sum_{z \in Z} |z| + 1$.

The budget limit is $b = c(\mathcal{P}) = 1 + 20 \sum_{z \in Z} |z|$, which means that in principle all the projects can be funded. There are two agents 1 and 2, with the type t_+ and t_- respectively. The PPB profile $\mathbf{A} = (\mathbf{A}^1)$ is formed with the following two approval ballots:

$$A_1 = \{p_i \mid i \in \{1, \dots, |Z|\} \text{ and } z_i \geq 0\} \cup \{p_+\},$$

$$A_2 = \{p_i \mid i \in \{1, \dots, |Z|\} \text{ and } z_i < 0\} \cup \{p_-\}.$$

The solution for I we consider is $\pi = (\pi_1)$ with $\pi_1 = \{p_+, p_-\}$. The question we want to answer is thus whether π fails Φ^{cost} -GINI or not.

We claim that there is a non-empty and feasible solution that has a lower Φ^{cost} -GINI coefficient than π if and only if Z is a positive SUBSET-SUM instance.

First, assume that there exists $Z' \subseteq Z$ such that $\sum_{z \in Z'} z = 0$. We show that the solution $\pi' = (\pi'_1)$ where $\pi'_1 = \{p_i \mid i \in \{1, \dots, |Z|\} \text{ and } z_i \in Z'\}$ has a lower Φ^{cost} -GINI coefficient than π .

On the one hand, with π agent 1 enjoys a cost satisfaction of $c(p_+)$ and agent 2 of $c(p_-)$. We thus have $\Phi^{cost\downarrow}(\mathbf{I}, \mathbf{A}, \pi) = (c(p_-), c(p_+))$. Overall, we have:

$$\begin{aligned} \text{GINI}_{\Phi^{cost}}(\mathbf{I}, \mathbf{A}, \pi) &= 1 - \frac{3 \cdot c(p_+) + c(p_-)}{2(c(p_+) + c(p_-))} \\ &= 1 - \frac{24 \cdot \sum_{z \in Z} |z| + 1}{24 \cdot \sum_{z \in Z} |z| + 2}, \end{aligned}$$

which is clearly larger than 0.

On the other hand, when considering π' , agents 1 and 2 enjoy the same cost-satisfaction of $\frac{1}{2} \cdot \sum_{z \in Z'} 4|z|$ each. This is because $\sum_{z \in Z'} z = 0$, so the sum of the positive numbers in Z' is equal to that of the negative numbers in Z' . From Fact 7.2.11, we thus immediately have:

$$\text{GINI}_{\Phi^{cost}}(\mathbf{I}, \mathbf{A}, \pi') = 0 < \text{GINI}_{\Phi^{cost}}(\mathbf{I}, \mathbf{A}, \pi).$$

This concludes the first direction of the proof.

Now assume there is a non-empty and feasible solution $\pi' = (\pi'_1)$ that has a lower EQUAL- Φ^{cost} coefficient than π . We claim that in this case, we have:

$$\Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi', t_+) = \sum_{p \in A_1^1 \cap \pi'_1} c(p) = \sum_{p \in A_2^1 \cap \pi'_1} c(p) = \Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi', t_-).$$

For the sake of contraction, assume that $\Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi', t_+) \neq \Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi', t_-)$.

Given Fact 7.2.12, we know that $\text{GINI}_{\Phi^{cost}}(\mathbf{I}, \mathbf{A}, \pi')$ is decreasing in the difference of the cost-satisfaction of agents 1 and 2. By construction, if agents 1 and 2 do not enjoy the same cost-satisfaction, the difference between the cost-satisfaction of the two agents must be at least 3. Indeed, their cost-satisfaction is always a multiple of 4, unless p_- is selected, in which case the cost-satisfaction of agent 1 is a multiple of 4, and that of agent 2 is 1 plus a multiple of 4. Fact 7.2.12 also implies that $\text{GINI}_{\Phi^{cost}}(\mathbf{I}, \mathbf{A}, \pi')$ is decreasing in the total cost-satisfaction of agents 1 and 2. Since we know that $\Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi, t_+) + \Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi, t_-) = c(\pi')$, we can thus lower bound $\text{GINI}_{\Phi^{cost}}(\mathbf{I}, \mathbf{A}, \pi')$ by the value achieved with a solution that funds all the projects. Indeed, this corresponds to both the highest total cost-satisfaction of agents 1 and 2, and the smallest difference in their cost-satisfaction (the difference being 3). Since π' cannot be empty, this constitutes a best case for $\text{GINI}_{\Phi^{cost}}(\mathbf{I}, \mathbf{A}, \pi')$, *i.e.*, lower bounds the Φ^{cost} -GINI coefficient of any suitable solution π' .

Let us then assume the best case, that is, $\pi' = \mathcal{P}$. In this case, we have:

$$\begin{aligned} \sum_{p \in A_1^1 \cap \pi'} c(p) &= 2 \sum_{z \in Z} |z| + 2 + c(p_+), \\ \sum_{p_i \in A_2^1 \cap \pi'} c(p) &= 2 \sum_{z \in Z} |z| - 2 + c(p_-). \end{aligned}$$

Then, the Φ^{cost} -GINI coefficient for π' is:

$$\begin{aligned} \text{GINI}_{\Phi^{cost}}(\mathbf{I}, \mathbf{A}, \pi') &= 1 - \frac{3(2 \sum_{z \in Z} |z| - 2 + c(p_-)) + 2 \sum_{z \in Z} |z| + 2 + c(p_+)}{2(2 \sum_{z \in Z} |z| - 2 + c(p_-) + 2 \sum_{z \in Z} |z| + 2 + c(p_+))} \\ &= 1 - \frac{30 \sum_{z \in Z} |z| - 3 + 10 \sum_{z \in Z} |z| + 2}{20 \sum_{z \in Z} |z| - 2 + 20 \sum_{z \in Z} |z| + 4} \\ &= 1 - \frac{40 \sum_{z \in Z} |z| - 1}{40 \sum_{i \leq k} |z_i| + 2}. \end{aligned}$$

Now, we can compare the Φ^{cost} -GINI coefficient of π and π' . To do so, we compute their difference:

$$\frac{24 \sum_{z \in Z} |z| + 1}{24 \sum_{z \in Z} |z| + 2} - \frac{40 \sum_{z \in Z} |z| - 1}{40 \sum_{z \in Z} |z| + 2} = \frac{32 \sum_{z \in Z} |z| + 4}{(24 \sum_{z \in Z} |z| + 2)(40 \sum_{z \in Z} |z| - 2)} > 0.$$

Therefore, $\text{GINI}_{\Phi^{cost}}(\mathbf{I}, \mathbf{A}, \pi') > \text{GINI}_{\Phi^{cost}}(\mathbf{I}, \mathbf{A}, \pi)$, a contradiction. This means that if there is a solution π' with a lower Φ^{cost} -GINI coefficient than π , then it must be the case that:

$$\sum_{p \in A_1^1 \cap \pi'_1} c(p) = \sum_{p \in A_2^1 \cap \pi'_1} c(p).$$

Since $c(p)$ is even for all $p \in \mathcal{P} \setminus \{p_-\}$ and $c(p_-)$ is odd, this implies that $p_- \notin \pi'$. Because $c(p_+) > \sum_{p \in \mathcal{P} \setminus \{p_-, p_+\}} c(p)$, the above equality can thus only hold if $p_+ \notin \pi'$. This yields that $Z' = \{z_i \mid p_i \in \pi'\}$ is a solution of the SUBSET-SUM instance Z .

So far, the proof has only focused on Φ^{cost} -GINI. However, in our construction the ballots of both voters are disjoint, meaning that Φ^{cost} and Φ^{share} coincide in this case. Hence, the same reduction also shows that Φ^{share} -GINI SATISFACTION is coNP-complete, even if we have only two voters. \square

We now prove the same result for the relative cost-satisfaction.

Proposition 7.4.2. *The problem $\Phi^{relcost}$ -GINI SATISFACTION is weakly coNP-complete, even if there are only one round and two agents.*

PROOF. The proof for $\Phi^{relcost}$ -GINI SATISFACTION is essentially the same as the one we provided for Φ^{cost} -GINI SATISFACTION in Proposition 7.4.1. In the following, we explain how to modify the reduction so that it works for $\Phi^{relcost}$ as well.

We add two additional projects p_1^* and p_2^* to the set of projects. Their cost is $c(p_1^*) = c(p_2^*) = b$, with the budget limit still being $b = \sum_{p \in \{p_1, \dots, p_{|Z|}, p_+, p_-\}} c(p)$.

Project p_1^* is approved by agent 1, and project p_2^* by agent 2. With this construction, for any budget allocation π , we have:

$$\begin{aligned}\Phi^{relcost}(\mathbf{I}, \mathbf{A}, (\pi), t_+) &= \frac{sat^{cost}(\pi \cap A_1^1)}{\max\{sat^{cost}(P) \mid P \subseteq A_1^1 \text{ and } c(P) \leq b\}} \\ &= \frac{sat^{cost}(\pi \cap A_1^1)}{c(p_1^*)} \\ &= \frac{\Phi^{cost}(\mathbf{I}, \mathbf{A}, (\pi), t_+)}{b}.\end{aligned}$$

The same holds for t_- instead of t_+ .

This implies that any solution formed from a budget allocation that does not contain p_1^* and p_2^* satisfies $\Phi^{relcost}$ -GINI in the new instance if and only if it satisfies Φ^{cost} -GINI in the original instance. We claim that, furthermore, a solution in which either p_1^* or p_2^* is selected cannot satisfy $\Phi^{relcost}$ -GINI. By construction, the only feasible solution in which p_1^* is selected (respectively p_2^*) is $(\{p_1^*\})$ (respectively $(\{p_2^*\})$), which clearly violates $\Phi^{relcost}$ -GINI. Hence, the modified reduction shows that $\Phi^{relcost}$ -GINI SATISFACTION is coNP-complete. \square

7.4.2 Among Exhaustive Solutions

Exhaustiveness is often a requirement in PB processes (see sections 2.1 and 3.4.1). We thus study the following variation of Φ -GINI SATISFACTION in which it is enforced.

EXHAUSTIVE Φ -GINI SATISFACTION	
Input:	A k -PPB instance $\mathbf{I} = (I_1, \dots, I_k)$, a PPB profile \mathbf{A} , and a solution $\pi = (\pi_1, \dots, \pi_k)$ for \mathbf{I} that is exhaustive.
Question:	Does π satisfy Φ -GINI with respect to \mathfrak{S}_{Ex} , the set of all the solutions π' such that every $\pi_j \in \pi'$ is feasible and exhaustive for I_j ?

In the following, we prove that the computational complexity does not change when requiring exhaustive solutions.

Proposition 7.4.3. *The problems EXHAUSTIVE Φ^{cost} -GINI SATISFACTION and EXHAUSTIVE $\Phi^{relcost}$ -GINI SATISFACTION are weakly coNP-complete, even if there are only one round and two agents.*

PROOF. The proof is similar to that of Proposition 7.4.1. We start with EXHAUSTIVE Φ^{cost} -GINI SATISFACTION and show that EXHAUSTIVE Φ^{cost} -GINI VIOLATION is NP-complete. As before, the problem EXHAUSTIVE Φ^{cost} -GINI VIOLATION asks

whether, given an exhaustive solution π , there exists another exhaustive solution π' with a lower Φ^{cost} -GINI coefficient.

It is clear that EXHAUSTIVE Φ^{cost} -GINI VIOLATION is in NP, as we can just guess a feasible and exhaustive solution, and check, in polynomial time, if it has a lower Φ^{cost} -GINI coefficient than the input solution π . Let us now prove that it is NP-hard. Once again, we use the SUBSET-SUM problem, see the proof of Proposition 7.3.3 for the definition of the problem.

Consider an instance of the SUBSET-SUM problem $Z \subseteq \mathbb{Z}$ and let $\zeta = \sum_{z \in Z} |z|$. We will now construct an instance of EXHAUSTIVE Φ^{cost} -GINI VIOLATION. First, consider the 1-PPB instance $\mathbf{I} = (I_1)$ where $I_1 = \langle \mathcal{P}, c, b \rangle$. The budget limit is $b = 4\zeta$. The set of projects is $\mathcal{P} = \{p_z \mid z \in Z\} \cup \{p_+\} \cup P^*$, with $P^* = \{p_1^*, \dots, p_{b-1}^*\}$. It is such that:

- ▶ For every $z \in Z$, there is a project p_z with $c(p_z) = 4|z|$;
- ▶ There is a additional project p_+ with $c(p_+) = 1$,
- ▶ There are $b - 1$ additional projects $\{p_1^*, \dots, p_{b-1}^*\}$, all of cost 1.

There are two agents 1 and 2 of type t_+ and t_- respectively. The PPB profile $\mathbf{A} = (\mathbf{A}^1)$ is composed of the following two approval ballots:

$$\begin{aligned} A_1^1 &= \{p_z \mid z \in Z \text{ and } z > 0\} \cup \{p_+\} \cup P^*, \\ A_2^1 &= \{p_z \mid z \in Z \text{ and } z < 0\} \cup P^*. \end{aligned}$$

Finally, the solution under consideration is $\pi = (\pi_1)$ where $\pi_1 = \{p_+\} \cup P^*$. It is clear that $c(\pi_1) = b$, meaning that π is exhaustive. We also have:

$$\begin{aligned} \Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi, t_+) &= b, \\ \Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi, t_-) &= b - 1. \end{aligned}$$

We claim that there is an exhaustive solution π' with a lower Φ^{cost} -GINI coefficient than π if and only if Z is a positive instance of the SUBSET-SUM problem.

First assume that Z is indeed a positive instance of the SUBSET-SUM problem. Let then $Z' \subseteq Z$ be a set such that $\sum_{z \in Z'} z = 0$, and let $P' = \{p_z \mid z \in Z'\}$ be the corresponding set of projects. We claim that the solution $\pi' = (\pi')$, where

$$\pi' = P' \cup \{p_1^*, \dots, p_{b-c(P')}^*\},$$

is exhaustive and has a lower Φ^{cost} -GINI coefficient than π .

As $c(\pi') = b$, π' is clearly exhaustive. Let us delve into the Φ^{cost} -GINI coefficients now. First, for π , we have:

$$\text{GINI}_{\Phi^{cost}}(\mathbf{I}, \mathbf{A}, \pi) = 1 - \frac{3(b-1) + b}{2(b-1+b)} = 1 - \frac{4b-3}{4b-2},$$

which is clearly larger than 0. On the other hand, for π' , we have:

$$\begin{aligned}\Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi', t_+) &= \sum_{p \in A_1^1 \cap (\pi' \setminus P^*)} c(p) + \sum_{p^* \in \pi' \cap P^*} c(p^*) \\ &= \sum_{p \in A_2^1 \cap (\pi' \setminus P^*)} c(p) + \sum_{p_i^* \in \pi' \cap P^*} c(p^*) \\ &= \Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi', t_-),\end{aligned}$$

where the second line is obtained from the first line by the fact that:

$$\sum_{z \in Z', z > 0} z = \sum_{z \in Z', z \leq 0} z.$$

We thus have $\text{GINI}_{\Phi^{cost}}(\mathbf{I}, \mathbf{A}, \pi') = 0 < \text{GINI}_{\Phi^{cost}}(\mathbf{I}, \mathbf{A}, \pi)$. This concludes the first part of the proof.

Now, assume there is an exhaustive solution $\pi' = (\pi')$ with a lower Φ^{cost} -GINI than π . We will find a corresponding solution for the SUBSET-SUM problem.

First of all, observe that since π' is exhaustive, it must contain at least one project from $\mathcal{P} \setminus P^*$. Indeed, by construction P^* is not exhaustive. Moreover, still by construction, we have:

$$\Phi^{cost}(\mathbf{I}, \mathbf{A}, (\pi' \cap P^*), t_+) = \sum_{p^* \in \pi' \cap P^*} c(p^*) = \Phi^{cost}(\mathbf{I}, \mathbf{A}, (\pi' \cap P^*), t_-). \quad (7.1)$$

We will distinguish two cases. First, assume that the following holds:

$$\Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi', t_+) = \Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi', t_-). \quad (7.2)$$

We claim that π' cannot contain p_+ in this case. Indeed, if it is the case that $p_+ \in \pi'$, then $\Phi^{cost}(\mathbf{I}, \mathbf{A}, (\pi' \setminus P^*), t_+)$ would be odd, while $\Phi^{cost}(\mathbf{I}, \mathbf{A}, (\pi' \setminus P^*), t_-)$ is always even. Together with (7.1), this contradicts assumption (7.2). Therefore, π' can only contain projects in $\{p_z \mid z \in Z\} \cup P^*$. Hence, it follows from (7.1) and (7.2) that:

$$\begin{aligned}\Phi^{cost}(\mathbf{I}, \mathbf{A}, (\pi' \setminus P^*), t_+) &= \sum_{p \in (\pi' \setminus P^*) \cap A_1^1} c(p) \\ &= \sum_{p \in (\pi' \setminus P^*) \cap A_2^1} c(p) \\ &= \Phi^{cost}(\mathbf{I}, \mathbf{A}, (\pi' \setminus P^*), t_-).\end{aligned}$$

Thus, for $Z' = \{z \in Z \mid p_z \in \pi'\}$, we have $\sum_{z \in Z'} z = 0$. Since π' is exhaustive, we know that Z' is non-empty. Z is thus a positive SUBSET-SUM instance.

Now, assume that the equality (7.2) does not hold, that is, that we have:

$$\Phi^{cost}(\mathbf{I}, \mathbf{A}, \boldsymbol{\pi}', t_+) \neq \Phi^{cost}(\mathbf{I}, \mathbf{A}, \boldsymbol{\pi}', t_-).$$

We show that in this case, the Φ^{cost} -GINI coefficient of $\boldsymbol{\pi}'$ is not lower than that of $\boldsymbol{\pi}$. This would be a contradiction.

As noted above $\boldsymbol{\pi}'$ must contain at least one project in $\mathcal{P} \setminus P^*$. If $\boldsymbol{\pi}'$ does not contain any projects from $\{p_z \mid z \in Z\}$, then we must have $\boldsymbol{\pi}' = \boldsymbol{\pi}$. It would then be impossible for $\boldsymbol{\pi}'$ to have a lower Φ^{cost} -GINI coefficient than $\boldsymbol{\pi}$. So, assume $\boldsymbol{\pi}'$ contains at least one project from $\{p_z \mid z \in Z\}$.

As before, and because of Fact 7.2.12, we can focus on the best case solution $\boldsymbol{\pi}'$ in terms of the Φ^{cost} -GINI coefficient. In the best case, $\boldsymbol{\pi}'$ can have:

$$\Phi^{cost}(\mathbf{I}, \mathbf{A}, \boldsymbol{\pi}', t_+) = b \quad \text{and} \quad \Phi^{cost}(\mathbf{I}, \mathbf{A}, \boldsymbol{\pi}', t_-) = b - 1,$$

as any other distribution of cost-satisfaction must have a bigger difference in satisfaction, or a lower total satisfaction. In both cases, the Gini coefficient would increase (because of Fact 7.2.12). Since this is the same distribution of cost-satisfaction as in $\boldsymbol{\pi}$, we have $\text{GINI}_{\Phi}^{cost}(\mathbf{I}, \mathbf{A}, \boldsymbol{\pi}) = \text{GINI}_{\Phi}^{cost}(\mathbf{I}, \mathbf{A}, \boldsymbol{\pi}')$. This sets the contradiction.

Finally, we look into $\Phi^{relcost}$. Observe that for both agents there is a feasible solution that gives them a cost-satisfaction of b . Therefore, we have:

$$\begin{aligned} \text{relsat}_{\text{sat}^{cost}}(\boldsymbol{\pi} \cap A_1^1) &= \frac{\text{sat}^{cost}(\boldsymbol{\pi} \cap A_1^1)}{b}, \\ \text{relsat}_{\text{sat}^{cost}}(\boldsymbol{\pi} \cap A_2^1) &= \frac{\text{sat}^{cost}(\boldsymbol{\pi} \cap A_2^1)}{b}. \end{aligned}$$

This implies that a solution satisfies $\Phi^{relcost}$ -GINI if and only if it satisfies Φ^{cost} -GINI (as the Gini coefficient is independent of multiplicative factors). Hence the reduction described above also shows that the problem EXHAUSTIVE $\Phi^{relcost}$ -GINI SATISFACTION is coNP-complete. \square

A similar proof also works for the share-based evaluation function. Some computations are different however.

Proposition 7.4.4. *The problem EXHAUSTIVE Φ^{share} -GINI SATISFACTION is weakly coNP-complete, even if there are only one round and two agents.*

PROOF. As before, we consider the co-problem EXHAUSTIVE Φ^{share} -GINI VIOLATION and show that it is weakly NP-complete. The problem is clearly in NP, we

show NP-hardness via a reduction from the problem SUBSET-SUM (defined in the proof of Proposition 7.3.3).

Consider an instance of the SUBSET-SUM problem $Z \subseteq \mathbb{Z}$ and let $\zeta = \sum_{z \in Z} |z|$. We will now construct an instance of EXHAUSTIVE Φ^{share} -GINI VIOLATION. First, consider the 1-PPB instance $\mathbf{I} = (I_1)$ where $I_1 = \langle \mathcal{P}, c, b \rangle$. The budget limit is then $b = 4\zeta$. The set of projects is $\mathcal{P} = \{p_z \mid z \in Z\} \cup \{p_+\} \cup P^*$ with $P^* = \{p_1^*, \dots, p_{b-1}^*\}$. It is such that:

- ▶ For every $z \in Z$, there is a project p_z with $c(p_z) = 4|z|$;
- ▶ There is a additional project p_+ with $c(p_+) = 1$;
- ▶ There are $b - 1$ additional projects $P^* = \{p_1^*, \dots, p_{b-1}^*\}$, all of cost 1.

There are two agents 1 and 2 of type t_+ and t_- respectively. The PPB profile $\mathbf{A} = (\mathbf{A}^1)$ is composed of the following two approval ballots:

$$\begin{aligned} A_1^1 &= \{p_z \mid z \in Z \text{ and } z > 0\} \cup \{p_+\} \cup P^*, \\ A_2^1 &= \{p_z \mid z \in Z \text{ and } z < 0\} \cup P^*. \end{aligned}$$

Finally, the solution under consideration is $\boldsymbol{\pi} = (\pi_1)$ where $\pi_1 = \{p_+\} \cup P^*$. It is clear that $c(\boldsymbol{\pi}) = b$, hence $\boldsymbol{\pi}$ is exhaustive. Further, observe that:

$$\begin{aligned} \Phi^{share}(\mathbf{I}, \mathbf{A}, \boldsymbol{\pi}, t_+) &= 1 + \frac{1}{2}(b - 1) = \frac{1}{2}b + \frac{1}{2}, \\ \Phi^{share}(\mathbf{I}, \mathbf{A}, \boldsymbol{\pi}, t_-) &= \frac{1}{2}(b - 1) = \frac{1}{2}b - \frac{1}{2} = \Phi^{share}(\mathbf{I}, \mathbf{A}, \boldsymbol{\pi}, t_+) - 1. \end{aligned}$$

We claim that there exists an exhaustive solution $\boldsymbol{\pi}'$ with a lower Φ^{share} -GINI coefficient than $\boldsymbol{\pi}$ if and only if Z is a positive SUBSET-SUM instance.

First, assume that Z is indeed a positive instance of the SUBSET-SUM problem. Let then $Z' \subseteq Z$ be such that $\sum_{z \in Z'} z = 0$, and let $P' = \{p_z \mid z \in Z'\}$ be the corresponding set of projects. We claim that the solution $\boldsymbol{\pi}' = (\pi')$ where

$$\boldsymbol{\pi}' = P' \cup \{p_1^*, \dots, p_{b-c(P')}^*\}$$

is exhaustive and has a lower Φ^{share} -GINI coefficient than $\boldsymbol{\pi}$.

As $c(\boldsymbol{\pi}') = b$, $\boldsymbol{\pi}'$ clearly is exhaustive. Let us now investigate the respective Φ^{share} -GINI coefficient. First, for $\boldsymbol{\pi}$, we have:

$$\text{GINI}_{\Phi}^{share}(\mathbf{I}, \mathbf{A}, \boldsymbol{\pi}) = 1 - \frac{3\left(\frac{1}{2}b - \frac{1}{2}\right) + \frac{1}{2}b + \frac{1}{2}}{2\left(\frac{1}{2}b - \frac{1}{2} + \frac{1}{2}b + \frac{1}{2}\right)} = 1 - \frac{2b - 1}{2b},$$

which is clearly larger than 0. On the other hand, for π' , we have:

$$\begin{aligned}\Phi^{share}(\mathbf{I}, \mathbf{A}, \pi', t_+) &= \sum_{p \in A_1^1 \cap (\pi' \setminus P^*)} c(p) + \frac{1}{2} \sum_{p^* \in \pi' \cap P^*} c(p^*) \\ &= \sum_{p \in A_2^1 \cap (\pi' \setminus P^*)} c(p) + \frac{1}{2} \sum_{p^* \in \pi' \cap P^*} c(p^*) \\ &= \Phi^{share}(\mathbf{I}, \mathbf{A}, \pi', t_-).\end{aligned}$$

Since the two types evaluate π' the same, we immediately have $\text{GINI}_{\Phi}^{share}(\mathbf{I}, \mathbf{A}, \pi') = 0 < \text{GINI}_{\Phi}^{share}(\mathbf{I}, \mathbf{A}, \pi)$. This concludes the first part of the proof.

Now, assume there is an exhaustive solution $\pi' = (\pi')$ with a lower Φ^{share} -GINI coefficient than π . First of all, observe that any exhaustive allocation must contain at least one project from $\mathcal{P} \setminus P^*$, because, by construction, P^* is not exhaustive. Furthermore, still by construction, we also have:

$$\Phi^{share}(\mathbf{I}, \mathbf{A}, (\pi' \cap P^*), t_+) = \frac{1}{2} \sum_{p^* \in \pi' \cap P^*} c(p^*) = \Phi^{share}(\mathbf{I}, \mathbf{A}, (\pi' \cap P^*), t_-). \quad (7.3)$$

We will then distinguish two cases. First, assume that the following holds:

$$\Phi^{share}(\mathbf{I}, \mathbf{A}, \pi', t_+) = \Phi^{share}(\mathbf{I}, \mathbf{A}, \pi', t_-). \quad (7.4)$$

In this case, π' does not contain p_+ . Indeed, if $p_+ \in \pi'$, then $\Phi^{share}(\mathbf{I}, \mathbf{A}, (\pi' \setminus P^*), t_+)$ would be odd, while $\Phi^{share}(\mathbf{I}, \mathbf{A}, (\pi' \setminus P^*), t_-)$ is always even. Given, (7.3), this would contradict (7.4). Therefore, we know that $\pi' \subseteq \{p_z \mid z \in Z\} \cup P^*$. Given (7.3) and (7.4), we thus know that:

$$\begin{aligned}\Phi^{share}(\mathbf{I}, \mathbf{A}, (\pi' \setminus P^*), t_+) &= \sum_{p \in (\pi' \setminus P^*) \cap A_1^1} c(p) \\ &= \sum_{p \in (\pi' \setminus P^*) \cap A_2^1} c(p) \\ &= \Phi^{share}(\mathbf{I}, \mathbf{A}, (\pi' \setminus P^*), t_-).\end{aligned}$$

Overall, for $Z' = \{z \in Z \mid p_z \in \pi'\}$, it must be the case that $\sum_{z \in Z'} z = 0$. Since π' is exhaustive, we know that Z' is non-empty. We have thus proven that Z' is a solution of the SUBSET-SUM problem for the instance Z .

Now, assume that the equality (7.4) does not hold, *i.e.*, that we have:

$$\Phi^{share}(\mathbf{I}, \mathbf{A}, \pi', t_+) \neq \Phi^{share}(\mathbf{I}, \mathbf{A}, \pi', t_-). \quad (7.5)$$

We show that in this case, the Φ^{share} -GINI coefficient of π' is not lower than that of π , a contradiction.

Given (7.5), it must be the case that $|\Phi^{share}(\mathbf{I}, \mathbf{A}, \pi', t_+) - \Phi^{share}(\mathbf{I}, \mathbf{A}, \pi', t_-)|$ is at least one as there is no project with cost less than one. Furthermore, by definition of the share, we know that $\Phi^{share}(\mathbf{I}, \mathbf{A}, \pi', t_+) + \Phi^{share}(\mathbf{I}, \mathbf{A}, \pi', t_-) = b$. Then, the smallest Φ^{share} -GINI coefficient for π' is reached if one type evaluates π' at $1/2(b+1)$ and the other at $1/2(b-1)$. This is the same distribution as for π , hence, $\text{GINI}_{\Phi}^{share}(\mathbf{I}, \mathbf{A}, \pi) = \text{GINI}_{\Phi}^{share}(\mathbf{I}, \mathbf{A}, \pi')$. This sets the contradiction. \square

This concludes our computational analysis of the fairness criteria based on the Gini coefficient. In the remainder of this chapter, we focus on EQUAL- Φ -CONV.

7.5 Converging Towards Fairness: EQUAL- Φ -CONV

The last section of our formal analysis is dedicated to EQUAL- Φ -CONV. Because EQUAL- Φ -CONV deals with infinite sequences, a concept that does not fit the standard framework of computational complexity, our analysis will only focus on existence guarantees. The section is organised based on the evaluation functions under consideration.

This section is also rather technical. For a quick summary, keep in mind that for both Φ^{cost} and Φ^{share} we can satisfy EQUAL- Φ -CONV only in very limited cases (two or three agents), while for $\Phi^{relcost}$ we can satisfy it if there are two *types* of agents.

7.5.1 For the Cost Evaluation Function Φ^{cost}

We start with the cost evaluation function Φ^{cost} and first show that for two agents, EQUAL- Φ^{cost} -CONV can always be guaranteed (under mild additional assumptions).

As the structure of the proof will be reused abundantly in this section, it can be worth keeping the general steps in mind. We will show that in every round we can favour one agent over the other, and thus, by each time favouring the worse-off agent, we can converge to EQUAL- Φ^{cost} .

Proposition 7.5.1. *Consider an ∞ -PPB instance \mathbf{I} such that there exists a constant $B^* \in \mathbb{N}$ with $b_j \leq B^*$ for every round $j \in \mathbb{N}_{>0}$, and a PBB profile \mathbf{A} with two agents. Furthermore, assume that for every round $j \in \mathbb{N}_{>0}$, each agent approves of at least one project of cost less than b_j . Then, there is a non-empty and feasible infinite solution π for \mathbf{I} that satisfies EQUAL- Φ^{cost} -CONV.*

PROOF. Call the agents 1 and 2 and assume they belong to types t_1 and t_2 respectively (as EQUAL- Φ^{cost} is trivially satisfied if there is only one type). We claim that

there exists a solution π such that for every round j , we can guarantee:

$$\Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t_1) - B^* \leq \Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t_2) \leq \Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t_1) + B^* \quad (7.6)$$

Let us prove the claim by induction. For the first round, it is clear that whichever non-empty budget allocation has been chosen, we have $0 \leq \Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[1]}, t) \leq B^*$ for any type $t \in \{t_1, t_2\}$. The claim would thus hold in the first round.

Now, assume the claim holds for round $j - 1$. Let $\pi_{[j-1]} = (\pi_1, \dots, \pi_{j-1})$. Without loss of generality, assume that type t_1 has higher cost-satisfaction than type t_2 in round $j - 1$, i.e., that $\Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j-1]}, t_2) \leq \Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j-1]}, t_1)$ holds. Let p be a project approved by agent 2 in round j such that $c(p) \leq b_j$. Such a project exists by assumption. Let $\pi_{[j]} = (\pi_1, \dots, \pi_{j-1}, \pi_j)$ for $\pi_j = \{p\}$. We clearly have:

$$sat^{cost}(\pi_j \cap A_1^j) \leq sat^{cost}(\pi_j \cap A_2^j) \leq B^*.$$

These inequalities, together with the induction hypothesis and the assumption that $\Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j-1]}, t_2) \leq \Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j-1]}, t_1)$ directly proves that our claim holds for round j . The induction is thus concluded.

Let π be a solution satisfying (7.6) for all rounds $j \in \mathbb{N}_{>0}$. It should be clear that for any round j , we also have $\Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t_1) + \Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t_2) \geq \sum_{j'=1}^j c(\pi_{j'})$. Thus, it follows from (7.6) that:

$$\lim_{j \rightarrow +\infty} \Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t_1) = \lim_{j \rightarrow +\infty} \Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t_2) = +\infty.$$

Therefore, for any round $j \in \mathbb{N}_{>0}$, we have:

$$\frac{\Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t_1) - B^*}{\Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t_1)} \leq \frac{\Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t_2)}{\Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t_1)} \leq \frac{\Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t_1) + B^*}{\Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t_1)}.$$

The proposition directly follows from this statement. \square

Unfortunately, this result cannot be generalized, not even for three agents.

Proposition 7.5.2. *There exists an ∞ -PPB instance \mathbf{I} for which there is a constant $B^* \in \mathbb{N}$ with $b_j \leq B^*$ for every round j , and a PPB profile \mathbf{A} with three agents such that no infinite solution for \mathbf{I} satisfies EQUAL- Φ^{cost} -CONV, even in the unit-cost setting and if for every round $j \in \mathbb{N}_{>0}$, each agent approve of at least one project of cost less than b_j .*

PROOF. Let \mathbf{I} be a ∞ -PPB instance. Assume $b_j = 1$ for every round $j \in \mathbb{N}_{>0}$, and $c(p) = 1$ for all projects $p \in \mathbb{P}$. There are three agents 1, 2 and 3. Agent

1 has type t_1 , and agents 2 and 3 have type t_2 . For every round $j \in \mathbb{N}_{>0}$, $\mathcal{P}_j = \{p_1^j, p_2^j\}$. In round j , agent 1 approves of p_1^j and p_2^j , agent 2 of p_1^j , and agent 3 of p_2^j . Then, for every non-empty feasible solution π and every round j , we have $\Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t_1) = j$ and $\Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t_2) = j/2$ as $|t_2| = 2$. Therefore, we have:

$$\lim_{j \rightarrow +\infty} \frac{\Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t_1)}{\Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t_2)} = \frac{1}{2}.$$

This is a clear violation of EQUAL- Φ^{cost} -CONV. \square

This counterexample can be avoided if we impose some restrictions on the ballots the agents can submit. For instance, if ballots are feasible and exhaustive, then for three agents we can always find a solution that converges to EQUAL- Φ^{cost} .

Proposition 7.5.3. *Consider an ∞ -PPB instance \mathbf{I} such that there exists a constant $B^* \in \mathbb{N}$ with $b_j \leq B^*$ for every round $j \in \mathbb{N}_{>0}$, and a PPB profile \mathbf{A} with three agents in which, for all rounds, the ballots of the agents are feasible and exhaustive. Then, there is a non-empty and feasible infinite solution π for \mathbf{I} that satisfies EQUAL- Φ^{cost} -CONV.*

PROOF. The proof is rather long. To simplify its delivery, we structure it around two claims, distinguishing between whether there are two or three types.

Claim 7.5.4. *The statement of Proposition 7.5.3 holds when there are exactly two types.*

Proof: Assume first that there are only two types, t_1 and t_2 . Without loss of generality, we can assume that agent 1 has type t_1 and agents 2 and 3 have type t_2 . We claim that there exists a solution π such that for every round $j \in \mathbb{N}_{>0}$, we can guarantee:

$$\Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t_1) - B^* \leq \Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t_2) \leq \Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t_1) + B^*. \quad (7.7)$$

Note that proving that this holds, would imply that π satisfies EQUAL- Φ^{cost} -CONV, in an analogous way to how we proved Proposition 7.5.1.

Let us prove the claim by induction. For the first round, it is clear that whichever non-empty budget allocation has been chosen, we have $0 \leq \Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t) \leq B^*$, for any type $t \in \{t_1, t_2\}$. The claim thus holds for the first round.

Now, assume the claim holds for round $j - 1$ and let $(\pi_1, \dots, \pi_{j-1})$ be a solution that satisfies (7.7) for round $j - 1$. We are trying to construct a budget allocation π_j such that the solution $\pi_{[j]} = (\pi_1, \dots, \pi_{j-1}, \pi_j)$ satisfies (7.7) for round j .

Assume first that $\Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j-1]}, t_1) \leq \Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j-1]}, t_2)$. Then, we can just set $\pi_j = A_1^j$. Since ballots are assumed to be feasible, this is a feasible budget allocation for I_j . Doing so would guarantee that:

$$\frac{sat^{cost}(A_2^j \cap \pi_j) + sat^{cost}(A_3^j \cap \pi_j)}{2} \leq sat^{cost}(A_1^j \cap \pi) \leq B^*.$$

Together with the induction hypothesis, this implies inequality (7.7).

Assume now that $\Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j-1]}, t_2) \leq \Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j-1]}, t_1)$. Let us make some further case distinctions based on the ballots of the agents.

- ▶ Consider the case in which the ballots are such that $A_1^j = A_2^j = A_3^j$, then the difference in cost-satisfaction would remain the same regardless of which π_j is selected. The induction hypothesis would then directly imply inequality (7.7) for round j and any π_j .
- ▶ Now, assume that $A_1^j = A_2^j = A_3^j$ does not hold. Since the ballots are assumed to be exhaustive, it cannot be the case that $A_i^j \subsetneq A_{i'}^j$ for any two agents $i, i' \in \{1, 2, 3\}$. There must therefore be a project p that is not approved by all three agents. We claim moreover that there is a project p^* in $A_2^j \cup A_3^j$ that is not approved by agent 1. Indeed, if no such p^* exists, then project p must be approved by agent 1, and only them (otherwise it would contradict the existence of p). Assume without loss of generality that agent 2 does not approve p . This implies that $A_1^j \cap A_2^j$ is not exhaustive, since $p^* \in A_1^j \setminus A_2^j$. Therefore $A_1^j \cap A_2^j \subsetneq A_2^j$. In other words, there is a project p^* as desired that is not approved by agent 1 but that is approved by either agent 2 or agent 3. Now, fix $\pi_j = \{p^*\}$. By definition, we have:

$$0 = sat^{cost}(A_1^j \cap \pi) \leq \frac{sat^{cost}(A_2^j \cap \pi_j) + sat^{cost}(A_3^j \cap \pi_j)}{2} \leq B^*.$$

Together with the induction hypothesis, this implies (7.7) for round j and the solution $(\pi_1, \dots, \pi_{j-1}, \pi_j)$.

This concludes the case when there are only two types. ◇

We now prove that the statement also holds if there are three types.

Claim 7.5.5. *The statement of Proposition 7.5.3 holds when there are exactly three types.*

Proof: Assume now that there are three types t_1, t_2 and t_3 . We can assume without loss of generality that agent $i \in \{1, 2, 3\}$ belongs to type t_i .

We claim that there exists a solution π such that for every round j , we can guarantee for any two types $t, t' \in \{t_1, t_2, t_3\}$ that:

$$|\Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t) - \Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t')| \leq 2B^*. \quad (7.8)$$

It is straightforward to check that EQUAL- Φ^{cost} -CONV would follow from this equation, by an analogous argument to the one made in Proposition 7.5.1.

Let us prove the claim by induction. For the first round, it is clear that whichever non-empty budget allocation is chosen, we would have $0 \leq \Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[1]}, t) \leq B^*$ for any type $t \in \{t_1, t_2, t_3\}$. The claim thus holds for the first round.

Now, assume the claim holds for round $j - 1$ and let $(\pi_1, \dots, \pi_{j-1})$ be a solution that satisfies (7.8) for round $j - 1$. We will construct a budget allocation π_j such that the solution $\pi_{[j]} = (\pi_1, \dots, \pi_{j-1}, \pi_j)$ satisfies (7.8) for round j . Assume without loss of generality that cost evaluations of the types are ordered as follows:

$$\Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j-1]}, t_1) \leq \Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j-1]}, t_2) \leq \Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j-1]}, t_3). \quad (7.9)$$

We observe that the induction hypothesis implies that:

$$\Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j-1]}, t_3) \leq \Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j-1]}, t_1) + 2B^*.$$

Therefore, we must have either:

$$\Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j-1]}, t_3) \leq \Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j-1]}, t_2) + B^*. \quad (7.10)$$

or:

$$\Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j-1]}, t_2) \leq \Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j-1]}, t_1) + B^*, \quad (7.11)$$

We will distinguish between these two cases.

- Assume first that (7.10) holds. Then, let $\pi_j = A_1^j$. From this, we know that the following two hold:

$$\begin{aligned} sat^{cost}(A_2^j \cap \pi_j) &\leq sat^{cost}(A_1^j \cap \pi_j) \leq B^*, \\ sat^{cost}(A_3^j \cap \pi_j) &\leq sat^{cost}(A_1^j \cap \pi_j) \leq B^*. \end{aligned}$$

Together with (7.9) this implies, that (7.8) holds for $t = t_1$ and $t' \in \{t_2, t_3\}$ at round j . It remains to show that it also holds for $t = t_2$ and $t' = t_3$ at round j .

Assume first that $sat^{cost}(A_2^j \cap \pi_j) \leq sat^{cost}(A_3^j \cap \pi_j)$. In this case, we can derive (7.8) at round j for $t = t_2$ and $t' = t_3$ from the fact that $sat^{cost}(A_2^j \cap \pi_j) \leq B^*$ and $\Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j-1]}, t_2) \leq \Phi^{cost}(\mathbf{I}, \mathbf{A}, \pi_{[j-1]}, t_3)$ hold.

Assume otherwise that $\text{sat}^{\text{cost}}(A_3^j \cap \pi_j) \leq \text{sat}^{\text{cost}}(A_2^j \cap \pi_j)$. Then, inequality (7.8) at round j for $t = t_2$ and $t' = t_3$ follows from inequality (7.10) together with the fact that $\text{sat}^{\text{cost}}(A_2^j \cap \pi_j) \leq B^*$ holds.

This concludes the branch of the proof in which we assumed (7.10) to hold.

- Now, assume that (7.11) holds. By the same argument as the one presented in Claim 7.5.3, we can find a project p^* such that $p^* \in A_1^j \cup A_2^j$ but $p^* \notin A_3^j$. From there, we set $\pi_j = \{p^*\}$, and distinguish, again, between two cases.

Assume first that $p^* \in A_1^j$. Then, it must be that:

$$\text{sat}^{\text{cost}}(A_3^j \cap \pi_j) \leq \text{sat}^{\text{cost}}(A_2^j \cap \pi_j) \leq \text{sat}^{\text{cost}}(A_1^j \cap \pi_j) \leq B^*.$$

Together with (7.9), this entails (7.8) at round j for all types $t, t' \in \{t_1, t_2, t_3\}$.

Then, assume that $p^* \in A_2^j \setminus A_1^j$. In this case, the following must hold:

$$\text{sat}^{\text{cost}}(A_1^j \cap \pi_j) = \text{sat}^{\text{cost}}(A_3^j \cap \pi_j) = 0 \quad \text{and} \quad \text{sat}^{\text{cost}}(A_2^j \cap \pi_j) \leq B^*.$$

Thanks to the induction hypothesis, this implies (7.8) at round j for $t = t_1$ and $t' = t_3$, and for $t = t_2$ and $t' = t_3$. Finally, using (7.11) and again $\text{sat}^{\text{cost}}(A_2^j \cap \pi_j) \leq B^*$ we can show that (7.8) also holds at round j for types $t = t_1$ and $t' = t_2$.

This concludes the case when there are three types. ◇

We have now prove the statement for both two and three types. Since the case when there is a single type is trivial, the proof is complete. □

Restricting the ballot format allowed us to obtain a possibility result for three agents. However, by increasing the number of agents we again encounter an impossibility, even with these restricted ballots.

Proposition 7.5.6. *There exists an ∞ -PPB instance \mathbf{I} such that there exists a constant $B^* \in \mathbb{N}$ with $b_j \leq B^*$ for every round j , and a PPB profile \mathbf{A} with eight agents such that no infinite solution for \mathbf{I} satisfies EQUAL- Φ^{cost} -CONV, even if there are only two types and for every round $j \in \mathbb{N}_{>0}$, each agent's ballot is exhaustive.*

PROOF. Let \mathbf{I} be a ∞ -PPB instance. In every round $j \in \mathbb{N}_{>0}$, we have $b_j = 10$ and $\mathcal{P}_j = \{p_1^j, \dots, p_6^j\}$. The cost function is such that for any round $j \in \mathbb{N}_{>0}$, we have $c(p_1^j) = c(p_2^j) = c(p_3^j) = 5$, and $c(p_4^j) = c(p_5^j) = c(p_6^j) = 3$.

We consider a PPB profile \mathbf{A} with eight agents $1, \dots, 8$ such that agents 1, 2 and 3 belong to type t_1 , and agents 4, 5, 6, 7 and 8 to type t_2 . The ballots are such

that, for every round j :

$$\begin{aligned} A_1^j &= \{p_1^j, p_4^j\} & A_2^j &= \{p_2^j, p_5^j\} & A_3^j &= \{p_3^j, p_6^j\} \\ A_4^j &= \{p_1^j, p_2^j\} & A_5^j &= \{p_1^j, p_3^j\} & A_6^j &= \{p_2^j, p_3^j\} \\ A_7^j &= \{p_4^j, p_5^j, p_6^j\} & A_8^j &= \{p_4^j, p_5^j, p_6^j\} \end{aligned}$$

We can check that at any round $j \in \mathbb{N}_{>0}$, for each project $p \in \mathcal{P}_j$, the cost evaluation of $\{p\}$ for type t_2 is higher than that of type t_1 .

For any round $j \in \mathbb{N}_{>0}$, given the budget allocation $\{p_1^j\}$ and, we have:

$$\begin{aligned} \Phi_{marg}^{cost}(\mathbf{I}, \mathbf{A}, (\{p_1^j\}), t_1, j) &= \frac{5 + 0 + 0}{3} \\ &< \frac{5 + 5 + 0 + 0 + 0}{5} = \Phi_{marg}^{cost}(\mathbf{I}, \mathbf{A}, (\{p_1^j\}), t_2, j). \end{aligned}$$

The same applies for $\{p_2\}$ or $\{p_3\}$ instead of $\{p_1\}$.

Now, for $\{p_4\}$ we have:

$$\begin{aligned} \Phi_{marg}^{cost}(\mathbf{I}, \mathbf{A}, (\{p_4^j\}), t_1, j) &= \frac{3 + 0 + 0}{3} \\ &< \frac{0 + 0 + 0 + 3 + 3}{5} = \Phi_{marg}^{cost}(\mathbf{I}, \mathbf{A}, (\{p_4^j\}), t_2, j). \end{aligned}$$

Again, the same applies for $\{p_5\}$ and $\{p_6\}$.

Since Φ^{cost} is additive with respect to the round and to the projects inside a budget allocation, this directly implies that there are no non-empty solution that satisfy EQUAL- Φ^{cost} -CONV in this PPB instance. \square

Our analysis of EQUAL- Φ^{cost} -CONV is now complete. The reader may find it interesting to know that in his Master's thesis, [Klein Goldewijk \(2022\)](#) generalised Proposition 7.5.3, proving that EQUAL- Φ^{cost} -CONV can always be satisfied for four agents and three types, when ballots are feasible exhaustive. He also reduced the gap left by Proposition 7.5.6 by proving an analogous result for 7 agents. Note that a gap still exists.

We will now turn to the case of EQUAL- Φ -CONV for the share evaluation function.

7.5.2 For the Share Evaluation Function Φ^{share}

In the following we discuss EQUAL- Φ^{share} -CONV. The result are very similar to the ones for EQUAL- Φ^{cost} -CONV. We will thus not provide many details.

Essentially, by a similar argument as the one used to prove Proposition 7.5.1, we can show that EQUAL- Φ^{share} -CONV can be achieved for two agents. Unfortunately, we

cannot go far beyond this, even if we assume ballots to be exhaustive. We present a counterexample below.

Example 7.5.7. Consider again the same ∞ -PPB instance as given in the proof of Proposition 7.5.6. We claim that for every project, selecting it would lead to a higher share for type t_2 than that of type t_1 . The computations are as follows.

At any round $j \in \mathbb{N}_{>0}$, if project p_1^j is selected, we have:

$$\begin{aligned}\Phi_{\text{marg}}^{\text{share}}(\mathbf{I}, \mathbf{A}, (\{p_1^j\}), t_1, j) &= \frac{1}{3} \cdot \frac{5}{3} = \frac{5}{9} \\ &< \frac{2}{3} = \frac{1}{5} \left(\frac{5}{3} + \frac{5}{3} \right) = \Phi_{\text{marg}}^{\text{share}}(\mathbf{I}, \mathbf{A}, (\{p_1^j\}), t_2, j).\end{aligned}$$

The same applies for projects p_2^j and p_3^j .

Now, still for any round $j \in \mathbb{N}_{>0}$, if project p_4^j is selected, we have:

$$\begin{aligned}\Phi_{\text{marg}}^{\text{share}}(\mathbf{I}, \mathbf{A}, (\{p_4^j\}), t_1, j) &= \frac{1}{3} \cdot \frac{3}{3} = \frac{1}{3} \\ &< \frac{2}{5} = \frac{1}{5} \left(\frac{3}{3} + \frac{3}{3} \right) = \Phi_{\text{marg}}^{\text{share}}(\mathbf{I}, \mathbf{A}, (\{p_4^j\}), t_2, j).\end{aligned}$$

The same applies for projects p_5^j and p_6^j .

It follows that, in this example, EQUAL- Φ^{share} -CONV cannot be satisfied. \triangle

This concludes our investigation of EQUAL- Φ^{share} -CONV. Similar open questions as for EQUAL- Φ^{cost} -CONV remain, including the case of 3 agents with feasible and exhaustive ballots. In the remainder of this section, we will investigate the case of the relative cost evaluation function.

7.5.3 For the Relative Cost Evaluation Function Φ^{relcost}

Results are more positive when it comes to relative satisfaction. Indeed, we can guarantee EQUAL- Φ^{relcost} -CONV for any PPB instance and PPB profile with two types of agents. Note that this thus applies to a much larger class of instances than what we obtained for EQUAL- Φ^{cost} -CONV and EQUAL- Φ^{share} -CONV for which we had guarantees for two agents only, and impossibility for 2 types and eight agents.

The proof is rather technical. We will start by proving an important lemma stating that with two types, we can always favour one type over the other.

Lemma 7.5.8. *Let \mathbf{I} be a k -PPB instance, \mathbf{A} a PPB profile in which for any round $j \in \{1, \dots, k\}$ and any agent $i \in \mathcal{N}$, the ballot A_i^j is non-empty and such that $c(A_i^j) < b$. Consider the solution $\boldsymbol{\pi} = (\pi_1, \pi_2, \dots)$ for \mathbf{I} . If there are only two types t_1 and t_2 , then, in every round $j \in \{1, \dots, k\}$ there are two feasible budget allocations π_j^1 and π_j^2 such that, if $\pi_j = \pi_j^1$ in $\boldsymbol{\pi}$, we have:*

$$\Phi_{\text{marg}}^{\text{relcost}}(\mathbf{I}, \mathbf{A}, \boldsymbol{\pi}, t_1, j) > 0 \quad \text{and} \quad \Phi_{\text{marg}}^{\text{relcost}}(\mathbf{I}, \mathbf{A}, \boldsymbol{\pi}, t_1, j) \geq \Phi_{\text{marg}}^{\text{relcost}}(\mathbf{I}, \mathbf{A}, \boldsymbol{\pi}, t_2, j),$$

and, if $\pi_j = \pi_j^2$ in π , we have:

$$\Phi_{\text{marg}}^{\text{relcost}}(\mathbf{I}, \mathbf{A}, \pi, t_2, j) > 0 \quad \text{and} \quad \Phi_{\text{marg}}^{\text{relcost}}(\mathbf{I}, \mathbf{A}, \pi, t_2, j) \geq \Phi_{\text{marg}}^{\text{relcost}}(\mathbf{I}, \mathbf{A}, \pi, t_1, j).$$

PROOF. Consider a round $j \in \{1, \dots, k\}$ corresponding to the instance $I_j = \langle \mathcal{P}_j, c, b_j \rangle$ and the profile \mathbf{A}^j . Consider $j - 1$ budget allocations π_1, \dots, π_{j-1} for the first $j - 1$ rounds of \mathbf{I} . We show that there exists $\pi_j \in \text{FEAS}(I_j)$ such that for the solution $\pi = (\pi_1, \dots, \pi_{j-1}, \pi_j)$ we have:

$$\Phi_{\text{marg}}^{\text{relcost}}(\mathbf{I}, \mathbf{A}, \pi, t_1, j) > 0 \quad \text{and} \quad \Phi_{\text{marg}}^{\text{relcost}}(\mathbf{I}, \mathbf{A}, \pi, t_2, j) \geq \Phi_{\text{marg}}^{\text{relcost}}(\mathbf{I}, \mathbf{A}, \pi, t_1, j),$$

The second part of the statement—when we favour t_2 over t_1 —then follows from an analogous argument.

Before delving into the detail, let us introduce one abbreviation that will make the proof easier to read. For a given type $t \in \mathcal{T}$ and budget allocation $\pi \in \text{FEAS}(I_j)$, we will use $\Phi_{\text{marg}}^{\text{relcost}}(\pi, t)$ defined as:

$$\Phi_{\text{marg}}^{\text{relcost}}(\pi, t) = \Phi_{\text{marg}}^{\text{relcost}}(\mathbf{I}, \mathbf{A}, (\pi_1, \dots, \pi_{j-1}, \pi), t, j).$$

Recall that all ballots in \mathbf{A}^j are feasible, *i.e.*, we have $A_i^j \in \text{FEAS}(I_j)$ for any agent $i \in \mathcal{N}$. For any budget allocation $\pi \in \text{FEAS}(I_j)$ and type $t \in \mathcal{T}$, we introduce the number of agents of type t who submit π as their ballot, defined as:

$$n_{\pi}^t = |\{i \in \mathcal{N} \mid A_i^j = \pi\}|.$$

Consider a budget allocation $\pi^* \in \text{FEAS}(I_j)$ such that $n_{\pi^*}^{t_1} \neq 0$. It is clear that if $\Phi_{\text{marg}}^{\text{relcost}}(\pi^*, t_1) \geq \Phi_{\text{marg}}^{\text{relcost}}(\pi^*, t_2)$, then the statement of the lemma holds and the proof would be concluded. Therefore, assume that we have:

$$\Phi_{\text{marg}}^{\text{relcost}}(\pi^*, t_1) < \Phi_{\text{marg}}^{\text{relcost}}(\pi^*, t_2). \quad (7.12)$$

Now, observe that by definition, for any type $t \in \mathcal{T}$, we have:

$$\Phi_{\text{marg}}^{\text{relcost}}(\pi^*, t) = \frac{1}{|t|} \sum_{\pi \in \text{FEAS}(I_j)} n_{\pi}^t \cdot \alpha(\pi, \pi^*), \quad (7.13)$$

where $\alpha(\pi, \pi^*)$ is the relative overlap of π and π^* , defined as:

$$\alpha(\pi, \pi^*) = \frac{c(\pi \cap \pi^*)}{c(\pi)}.$$

It should be clear that by definition, $0 \leq \alpha(\pi, \pi^*) \leq 1$ holds for any $\pi \in \text{FEAS}(I_j)$.

Now, plugging (7.13) into inequality (7.12) leads to:

$$\begin{aligned} \frac{1}{|t_1|} \sum_{\pi \in \text{FEAS}(I_j)} n_{\pi}^{t_1} \cdot \alpha(\pi, \pi^*) &< \frac{1}{|t_2|} \sum_{\pi \in \text{FEAS}(I_j)} n_{\pi}^{t_2} \cdot \alpha(\pi, \pi^*) \\ \Leftrightarrow \frac{\sum_{\pi \in \text{FEAS}(I_j)} n_{\pi}^{t_1} \cdot \alpha(\pi, \pi^*)}{\sum_{\pi \in \text{FEAS}(I_j)} n_{\pi}^{t_2} \cdot \alpha(\pi, \pi^*)} &< \frac{|t_1|}{|t_2|}. \end{aligned} \quad (7.14)$$

Let $\mathcal{A}^1 = \{\pi \in \text{FEAS}(I_j) \mid \alpha(\pi, \pi^*) = 1\}$ be the set of budget allocations π for which $\alpha(\pi, \pi^*) = 1$. Inequality (7.14) can then be rewritten as:

$$\frac{\sum_{\pi \in \mathcal{A}^1} n_{\pi}^{t_1} + \sum_{\pi \in \text{FEAS}(I_j) \setminus \mathcal{A}^1} n_{\pi}^{t_1} \cdot \alpha(\pi, \pi^*)}{\sum_{\pi \in \mathcal{A}^1} n_{\pi}^{t_2} + \sum_{\pi \in \text{FEAS}(I_j) \setminus \mathcal{A}^1} n_{\pi}^{t_2} \cdot \alpha(\pi, \pi^*)} < \frac{|t_1|}{|t_2|}.$$

For this to hold, we must have either:

$$\frac{\sum_{\pi \in \mathcal{A}^1} n_{\pi}^{t_1}}{\sum_{\pi \in \mathcal{A}^1} n_{\pi}^{t_2}} < \frac{|t_1|}{|t_2|} \quad \text{or} \quad \frac{\sum_{\pi \in \text{FEAS}(I_j) \setminus \mathcal{A}^1} n_{\pi}^{t_1} \cdot \alpha(\pi, \pi^*)}{\sum_{\pi \in \text{FEAS}(I_j) \setminus \mathcal{A}^1} n_{\pi}^{t_2} \cdot \alpha(\pi, \pi^*)} < \frac{|t_1|}{|t_2|}. \quad (7.15)$$

The above two inequalities are derived from the general rule that for any positive scalar $a, b, c, d, x, y \in \mathbb{R}_{>0}$, if we have $(a+b)/(c+d) < x/y$, then we either $a/c < x/y$ or $b/d < x/y$. Indeed, The $(a+b)/(c+d) < x/y$ can be rewritten as $ay + by < cx + dx$, which necessarily only holds if either $ay < cx$ or $by < dx$.

Since the left-hand side of the second inequality in (7.15) again involves summations, we can iterate the same general rule there. For this inequality to hold, there must then be sufficiently many budget allocations $\pi \in \text{FEAS}(I_j) \setminus \mathcal{A}^1$ with $n_{\pi}^{t_1}/n_{\pi}^{t_2} < |t_1|/|t_2|$. As we know that $\alpha(\pi, \pi^*) < 1$, dropping $\alpha(\pi, \pi^*)$ only increases the influence of these budget allocations. Therefore, in the two cases in (7.15), we can find a set of budget allocations \mathcal{A} with $\mathcal{A}^1 \subseteq \mathcal{A} \subseteq \text{FEAS}(I_j)$ and such that the following holds:

$$\frac{\sum_{\pi \in \mathcal{A}} n_{\pi}^{t_1}}{\sum_{\pi \in \mathcal{A}} n_{\pi}^{t_2}} < \frac{|t_1|}{|t_2|}. \quad (7.16)$$

On the other hand, since for any type t , we have $|t| = \sum_{\pi \in \text{FEAS}(I_j)} n_{\pi}^t$, inequality (7.16) can be rewritten as:

$$\frac{\sum_{\pi \in \mathcal{A}} n_{\pi}^{t_1}}{\sum_{\pi \in \mathcal{A}} n_{\pi}^{t_2}} < \frac{\sum_{\pi \in \text{FEAS}(I_j)} n_{\pi}^{t_1}}{\sum_{\pi \in \text{FEAS}(I_j)} n_{\pi}^{t_2}}.$$

Given that $\mathcal{A} \subseteq \text{FEAS}(I_j)$, we immediately have $\sum_{\pi \in \mathcal{A}} n_{\pi}^{t_2} \leq \sum_{\pi \in \text{FEAS}(I_j)} n_{\pi}^{t_2}$. Thus, for the above to hold, it must be that $\sum_{\pi \in \mathcal{A}} n_{\pi}^{t_1} < \sum_{\pi \in \text{FEAS}(I_j)} n_{\pi}^{t_1}$. There must then exist a budget allocation $\pi^0 \in \text{FEAS}(I_j) \setminus \mathcal{A}$ such that $n_{\pi^0}^{t_1} \neq 0$. Since $\pi^0 \notin \mathcal{A}$, we have $\alpha(\pi^0, \pi^*) < 1$, which implies $\pi^0 \neq \pi^*$.

Consider then the budget allocation $\pi^{*'} = \pi^0 \setminus \pi^*$. Since $\alpha(\pi^0, \pi^*) < 1$, we must have $\pi^{*'} \neq \emptyset$. Furthermore, $\Phi_{\text{marg}}^{\text{relcost}}(\pi^{*'}, t_1) > 0$ also holds as we know there is an agent in t_1 that submitted π^0 as their ballot. Therefore, if $\Phi_{\text{marg}}^{\text{relcost}}(\pi^{*'}, t_1) \geq \Phi_{\text{marg}}^{\text{relcost}}(\pi^{*'}, t_2)$ holds, then the lemma holds as well and the proof is concluded. Hence, we assume that:

$$\Phi_{\text{marg}}^{\text{relcost}}(\pi^{*'}, t_1) < \Phi_{\text{marg}}^{\text{relcost}}(\pi^{*'}, t_2).$$

Then, we have

$$\Phi_{\text{marg}}^{\text{relcost}}(\pi^*, t_1) + \Phi_{\text{marg}}^{\text{relcost}}(\pi^{*'}, t_1) < \Phi_{\text{marg}}^{\text{relcost}}(\pi^*, t_2) + \Phi_{\text{marg}}^{\text{relcost}}(\pi^{*'}, t_2). \quad (7.17)$$

Analogously to what we proved earlier, it follows from (7.13) and (7.17) that:

$$\frac{\sum_{\pi \in \text{FEAS}(I_j)} n_{\pi}^{t_1} (\alpha(\pi, \pi^*) + \alpha(\pi, \pi^{*'}))}{\sum_{\pi \in \text{FEAS}(I_j)} n_{\pi}^{t_2} (\alpha(\pi, \pi^*) + \alpha(\pi, \pi^{*'}))} < \frac{|t_1|}{|t_2|}.$$

Now, because π^* and $\pi^{*'}$ are disjoint, we know all factors $\alpha(\pi, \pi^*) + \alpha(\pi, \pi^{*'})$ are smaller or equal 1. Following the same reasoning we detailed above, we can conclude that there must be a set of budget allocations $\mathcal{A}' \subseteq \text{FEAS}(I_j)$ such that $\alpha(\pi, \pi^*) + \alpha(\pi, \pi^{*'}) = 1$ implies $\pi \in \mathcal{A}'$, and such that the following holds:

$$\frac{\sum_{\pi \in \mathcal{A}'} n_{\pi}^{t_1}}{\sum_{\pi \in \mathcal{A}'} n_{\pi}^{t_2}} < \frac{|t_1|}{|t_2|}. \quad (7.18)$$

It follows again from $|t| = \sum_{\pi \in \text{FEAS}(I_j)} n_{\pi}^t$ and (7.18) that there must be another budget allocation π^1 such that $n_{\pi^1}^{t_1} > 0$ and $\alpha(\pi^1, \pi^*) + \alpha(\pi^1, \pi^{*'}) < 1$. This implies that $\pi^1 \notin \{\pi^*, \pi^{*'}\}$. Then, as before, we have $\Phi_{\text{marg}}^{\text{relcost}}(\pi^1 \setminus (\pi^* \cup \pi^{*'}), t_1) > 0$. Hence, if

$$\Phi_{\text{marg}}^{\text{relcost}}(\pi^1 \setminus (\pi^* \cup \pi^{*'}), t_1) \geq \Phi_{\text{marg}}^{\text{relcost}}(\pi^1 \setminus (\pi^* \cup \pi^{*'}), t_2)$$

holds, then the lemma holds as well. Otherwise, we can iterate the construction.

As there are only finitely many budget allocations, this construction must lead, after finitely many steps, to an allocation π such that:

$$\Phi_{\text{marg}}^{\text{relcost}}(\pi, t_1) > 0 \quad \text{and} \quad \Phi_{\text{marg}}^{\text{relcost}}(\pi, t_1) \geq \Phi_{\text{marg}}^{\text{relcost}}(\pi, t_2).$$

The proof is thus complete. □

Thanks to this lemma, and using a similar line of reasoning as in Proposition 7.5.1, we can show that, for two types, we can always find a solution that converges to equal relative satisfaction.

Theorem 7.5.9. Consider an ∞ -PPB instance \mathbf{I} such that there exists a constant $B^* \in \mathbb{N}$ with $b_j \leq B^*$ for every round $j \in \mathbb{N}_{>0}$, and a PPB profile \mathbf{A} in which for all rounds, the ballots of the agents are non-empty and feasible. If there are only two types, then there is a non-empty and feasible infinite solution π for \mathbf{I} that satisfies EQUAL- $\Phi^{relcost}$ -CONV.

PROOF. Let us call the two types t_1 and t_2 . We claim that there exists a solution $\pi = (\pi_1, \pi_2, \dots)$ such that for every round j , we can guarantee:

$$\Phi^{relcost}(\mathbf{I}, \mathbf{A}, \pi, t_1) - 1 \leq \Phi^{relcost}(\mathbf{I}, \mathbf{A}, \pi, t_2) \leq \Phi^{relcost}(\mathbf{I}, \mathbf{A}, \pi, t_1) + 1.$$

We will prove this by induction. For the first round and for both types $t \in \{t_1, t_2\}$, we clearly have $0 \leq \Phi^{relcost}(\mathbf{I}, \mathbf{A}, \pi, t) \leq 1$. The claim thus holds then.

Now assume the claim holds for round $j - 1$. Without loss of generality, assume that:

$$\Phi^{relcost}(\mathbf{I}, \mathbf{A}, \pi_{[j-1]}, t_2) \leq \Phi^{relcost}(\mathbf{I}, \mathbf{A}, \pi_{[j-1]}, t_1). \quad (7.19)$$

At round j we can select a budget allocation π_j in π that satisfies:

$$\Phi_{marg}^{relcost}(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t_1, j) \leq \Phi_{marg}^{relcost}(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t_2, j). \quad (7.20)$$

The existence of such a π_j is guaranteed by Lemma 7.5.8. Then, for both types $t \in \{t_1, t_2\}$, we have:

$$\Phi^{relcost}(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t) = \Phi^{relcost}(\mathbf{I}, \mathbf{A}, \pi_{[j-1]}, t) + \Phi_{marg}^{relcost}(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t_1, j).$$

Then, from (7.20) and the induction hypothesis, it follows that:

$$\Phi^{relcost}(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t_1) - 1 \leq \Phi^{relcost}(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t_2).$$

On the other hand, from (7.19) and the fact that $\Phi_{marg}^{relcost}(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t_2, j) \leq 1$ holds by definition, we obtain:

$$\Phi^{relcost}(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t_2) \leq \Phi^{relcost}(\mathbf{I}, \mathbf{A}, \pi_{[j]}, t_1) + 1.$$

Now, in each round, we know from Lemma 7.5.8 that we can always select a budget allocation that improves the relative satisfaction of the type t that enjoys the lowest relative satisfaction by at least $1/B^* \cdot |t|$ (this is the lowest non-zero relative cost evaluation of a type). Selecting such a budget allocation at each round leads to the relative cost evaluation of both types to go towards infinity, while their difference to always be less than 1. The solution π thus constructed satisfies EQUAL- $\Phi^{relcost}$ -CONV. \square

It is important to mention that the proofs of Lemma 7.5.8 and of Theorem 7.5.9 are

both constructive, in the sense that they show how to compute the relevant solutions. However, the constructions do not guarantee the solution to be exhaustive. To achieve this, an additional ballot restriction is necessary.

Proposition 7.5.10. *Consider an ∞ -PPB instance \mathbf{I} and a PBB profile \mathbf{A} that satisfies all the conditions of Theorem 7.5.9. If there are exactly two types, then, there exists a non-empty feasible solution $\pi = (\pi_1, \pi_2, \dots)$ for \mathbf{I} that:*

- (i) *satisfies EQUAL- Φ^{relcost} -CONV, and,*
- (ii) *ensures that for each round $j \in \mathbb{N}_{>0}$, there is an agent $i \in \mathcal{N}$ such that $A_i^j \subseteq \pi_j$.*

In particular, if all ballots are exhaustive, then every budget allocation in π is exhaustive.

PROOF. The idea of the proof is that for every round $j \in \mathbb{N}_{>0}$, we can find two budget allocations π_j^1 and π_j^2 as described in Lemma 7.5.8 and such that there are two agents $i_1, i_2 \in \mathcal{N}$ with $A_{i_1}^j \setminus \pi_j^1 = \emptyset$ and $A_{i_2}^j \setminus \pi_j^2 = \emptyset$ by applying Lemma 7.5.8 several times.

Let $I_j = \langle \mathcal{P}_j, c, b_j \rangle$ be the instance in round j and A^j the profile in round j . We use the same notation as in Lemma 7.5.8, so that we use $\Phi_{\text{marg}}^{\text{relcost}}(\pi, t)$ instead of $\Phi_{\text{marg}}^{\text{relcost}}(\mathbf{I}, \mathbf{A}, (\pi_1, \dots, \pi_{j-1}, \pi), t, j)$. We claim that, under the given assumptions, there exist two feasible budget allocations $\pi_j^1, \pi_j^2 \in \text{FEAS}(I_j)$ such that:

$$\begin{aligned} \Phi_{\text{marg}}^{\text{relcost}}(\pi_j^1, t_1) > 0 & \quad \text{and} & \quad \Phi_{\text{marg}}^{\text{relcost}}(\pi_j^1, t_1) \geq \Phi_{\text{marg}}^{\text{relcost}}(\pi_j^1, t_2), \\ \Phi_{\text{marg}}^{\text{relcost}}(\pi_j^2, t_2) > 0 & \quad \text{and} & \quad \Phi_{\text{marg}}^{\text{relcost}}(\pi_j^2, t_2) \geq \Phi_{\text{marg}}^{\text{relcost}}(\pi_j^2, t_1), \end{aligned}$$

and for which there are two agents $i_1, i_2 \in \mathcal{N}$ with $A_{i_1}^j \setminus \pi_j^1 = \emptyset$ and $A_{i_2}^j \setminus \pi_j^2 = \emptyset$. We will only prove the existence of π_j^1 , the existence of π_j^2 follows from an analogous argument.

From Lemma 7.5.8, we know that there exists $\pi_j^1 \in \text{FEAS}(I_j)$, such that:

$$\Phi_{\text{marg}}^{\text{relcost}}(\pi_j^1, t_1) > 0 \quad \text{and} \quad \Phi_{\text{marg}}^{\text{relcost}}(\pi_j^1, t_1) \geq \Phi_{\text{marg}}^{\text{relcost}}(\pi_j^1, t_2).$$

If there is no agent $i \in \mathcal{N}$ such that $A_i^j \setminus \pi_j^1 = \emptyset$, we can consider the instance $I_j^1 = \langle \mathcal{P}_j^1, c, b_j \rangle$ where $\mathcal{P}_j^1 = \mathcal{P}_j \setminus \pi_j^1$, and the profile $\mathbf{A}^{1j} = (A_1^{1j}, \dots, A_n^{1j})$ such that for all agents $i \in \mathcal{N}$ we have $A_i^{1j} = A_i^j \setminus \pi_j^1$. By assumption, $A_i^{1j} \neq \emptyset$ for all agents $i \in \mathcal{N}$. Therefore, in I_j^1 and \mathbf{A}^{1j} , there are two types, all ballots are non-empty and satisfy the budget constraint. We can thus apply Lemma 7.5.8 once again. We obtain a budget allocation $\pi_j^{1'} \in \text{FEAS}(I_j^1)$ such that:

$$\Phi_{\text{marg}}^{\text{relcost}}(\pi_j^{1'}, t_1) > 0 \quad \text{and} \quad \Phi_{\text{marg}}^{\text{relcost}}(\pi_j^{1'}, t_1) \geq \Phi_{\text{marg}}^{\text{relcost}}(\pi_j^{1'}, t_2).$$

By definition, we know that $\pi_j^1 \cap \pi_j^{1'} = \emptyset$. Therefore, we have $\Phi_{\text{marg}}^{\text{relcost}}(\pi_j^1 \cup \pi_j^{1'}, t_1) > 0$ and:

$$\begin{aligned} \Phi_{\text{marg}}^{\text{relcost}}(\pi_j^1 \cup \pi_j^{1'}, t_1) &= \Phi_{\text{marg}}^{\text{relcost}}(\pi_j^1, t_1) + \Phi_{\text{marg}}^{\text{relcost}}(\pi_j^{1'}, t_1) \\ &\geq \Phi_{\text{marg}}^{\text{relcost}}(\pi_j^1, t_2) + \Phi_{\text{marg}}^{\text{relcost}}(\pi_j^{1'}, t_2) \\ &= \Phi_{\text{marg}}^{\text{relcost}}(\pi_j^1 \cup \pi_j^{1'}, t_2). \end{aligned}$$

Now, if there is no agent i such that $A_i^j \setminus \pi_j^1 \cup \pi_j^{1'} = \emptyset$, we can apply Lemma 7.5.8 again and again, until we have $A_i^j \setminus (\pi_j^1 \cup \pi_j^{1'} \cup \dots) = \emptyset$ for some agent $i \in \mathcal{N}$. This thus proves the claim.

Finally, note that if all ballots are exhaustive, then $A_i^j \setminus \pi = \emptyset$ can only hold for an exhaustive budget allocation π . \square

Whether Theorem 7.5.9 and Proposition 7.5.10 can be extended to larger numbers of types remains an important open question. Under additional assumptions, Klein Goldewijk (2022) proved that $\text{EQUAL-}\Phi^{\text{relcost}}\text{-CONV}$ is always satisfiable when there are three types. This is the strongest result we are aware of regarding $\text{EQUAL-}\Phi^{\text{relcost}}\text{-CONV}$.

7.6 Summary

With this chapter, we presented the first formal study of a long-term model for participatory budgeting. In addition to the conceptual contribution of introducing such a model, we investigated several notions of fairness within this framework. We started from a notion that represents a situation of perfect fairness: $\text{EQUAL-}\Phi$. Because it is impossible to satisfy in general, we turned to two of its relaxations that provide fairness guarantees on a larger set of instances. First, we explored the idea of optimising for fairness through the study of $\Phi\text{-GINI}$. Our computational analysis revealed that the latter was hard, computationally speaking, to satisfy. Subsequently, we looked into the idea of providing fairness, but only in an infinite horizon. Our study of $\text{EQUAL-}\Phi\text{-CONV}$ demonstrated that this was possible, maybe not in general, but at least for a large set of instances. We summarise all our findings in Table 7.6.1 for $\text{EQUAL-}\Phi$ and $\Phi\text{-GINI}$, and in Table 7.6.2 for $\text{EQUAL-}\Phi\text{-CONV}$.

	Existence Guarantees	Complexity
$\text{EQUAL-}\Phi^{\text{cost}}$ $\text{EQUAL-}\Phi^{\text{relcost}}$	\times (even for $n = 2$) Proposition 7.3.1	Strongly NP-complete even with exhaustiveness Proposition 7.3.2 and Corollary 7.3.4
$\text{EQUAL-}\Phi^{\text{share}}$	\times (even for $n = 2$) Proposition 7.3.1	Weakly NP-complete even with exhaustiveness Proposition 7.3.3 and Corollary 7.3.4
$\Phi^{\text{cost}}\text{-GINI}$ $\Phi^{\text{relcost}}\text{-GINI}$ $\Phi^{\text{share}}\text{-GINI}$	\checkmark By definition	Weakly coNP-complete even with exhaustiveness Propositions 7.4.1, 7.4.2, 7.4.3, and 7.4.4

Table 7.6.1: Summary of the results presented in this chapter about $\text{EQUAL-}\Phi$ and $\Phi\text{-GINI}$. For the complexity results, they refer to the problems $\text{EQUAL-}\Phi$ SATISFIABILITY and $\Phi\text{-GINI}$ SATISFIABILITY. The reference to exhaustiveness indicates that the results hold even if exhaustiveness is required.

	Existence Guarantees	
$\text{EQUAL-}\Phi^{\text{cost}}\text{-CONV}$	\checkmark for $n = 2$ \times for $n = 3$ \checkmark for $n = 3$ and exhaustive ballots \times for $n \geq 8$ and exhaustive ballots	Proposition 7.5.1 Proposition 7.5.2 Proposition 7.5.3 Proposition 7.5.6
$\text{EQUAL-}\Phi^{\text{relcost}}\text{-CONV}$	\checkmark for $ \mathcal{T} = 2$ and feasible ballots	Theorem 7.5.9
$\text{EQUAL-}\Phi^{\text{share}}\text{-CONV}$	\checkmark for $n = 2$ \times for $n \geq 8$ and exhaustive ballots	Section 7.5.3 Example 7.5.7

Table 7.6.2: Summary of the results presented in this chapter about $\text{EQUAL-}\Phi\text{-CONV}$. A check mark \checkmark indicates that for all instances with the specified number of agents or types there always exists a solution satisfying the fairness criteria, with some potential additional assumptions on the ballots.

Chapter 8

An End-to-End Model for Participatory Budgeting

As already described in Section 1.1.2, real-life implementations of PB processes usually span several months and include many steps. These steps can be roughly grouped into the following three stages:

- ▶ Recovering all the projects that can be potentially implemented through the PB process, and selecting a shortlist of them that will carry over to the next step;
- ▶ Collecting the opinions of the participants regarding the shortlisted projects in order to determine the ones that will actually be implemented;
- ▶ Monitoring the actual implementation of the projects and reporting to the participants about it.

From a social choice perspective, the second stage is the most exciting—how to collect and aggregate opinions is actually the research question at the very core of the discipline. It is thus probably not surprising that it has been the focus of almost all the social choice literature on PB (see Chapter 3), including all the previous chapters. Does that mean that this is the only stage of PB that is worth studying? I am claiming that the answer to this question is no, and that there are many interesting research questions when one considers the PB process in its entirety, from one end to the other. The present chapter is here to convince you of that.

In what follows we will develop an *end-to-end* model of PB. Despite the name, this model will only cover the first two stages that we discussed above. Indeed, the third stage does not really call for a formal analysis, at least according to the meaning of formal that we are using in this thesis. Where does the social choice challenge reside in the first stage? By looking deeper into how potential projects are collected by the organising body—say a municipality—we can learn that in most cases they come from

the citizens themselves who submit proposals directly to the municipality (Shah, 2007; Wampler, McNulty and Touchton, 2021). Citizens are thus also involved during the first stage of a PB process, submitting their opinions on the projects that should be considered by the municipality. Later on, the municipality selects a shortlist of the proposals by taking said opinions into account. What we just described is a perfect example of a social choice problem. Our end-to-end model will thus formalise the first stage of PB, where proposals are collected and shortlisted, and the second stage where the final budget allocation is selected. It is hopefully clear that the second stage corresponds to the standard model of PB introduced in Chapter 2.

Our first point of focus in this new model will be the shortlisting stage. Interpreting it as multi-winner voting scenario—where the proposals submitted by the agents represent the ballots they cast, and the set of shortlisted proposals is the outcome—we will discuss different ways of determining the shortlist, different *shortlisting rules*. It is important to keep in mind that there are no formal constraints on the outcome of the shortlisting stage. Indeed, any subset of proposals is an admissible outcome (except for the empty set maybe). Because of this, there is a lot of room to develop shortlisting rules. This raises the question of what makes some shortlists more desirable than others. To inform our answer to this question, we will once again look into what is happening in practice. We identify four distinct objectives.

- (i) A first round of review usually removes the proposals that are simply infeasible, typically because of legal issues.⁵⁷ This specific objective cannot really be incorporated into our analysis and we assume that all projects considered in our formal model are implementable.
- (ii) Another goal of the shortlisting stage is to reduce the number of proposals entering the second stage. For instance, if we look at the PB exercises in Lisbon, around 30% of the projects were shortlisted (Allegretti and Antunes, 2014). In Toronto, this number was as low as 10% (Murray, 2019).
- (iii) What we call the shortlisting stage often takes the form of rounds of public meetings where the proposals are discussed. During these meetings, citizens will typically develop their proposals, helped by other citizens and employees of the municipality. It is common for official organisers to avoid that projects that are too similar get proposed.⁵⁸ This constitutes our third objective of the shortlisting stage: avoiding redundancy in the shortlist.
- (iv) As mentioned already before, one of the core objectives of PB is to provide citizens with a platform to express their opinion, and to witness direct impact of

⁵⁷For instance, in the 2017 PB process in Paris, one of the proposals was to demolish the Sacré-Cœur, a church in the centre of Paris. This proposal was rejected by the municipality of Paris because it falls outside of its jurisdiction. For more details, see francetvinfo.fr/france/ile-de-france/paris/affreux-disproportionne-un-parisien-propose-a-la-mairie-de-raser-le-sacre-coeur_2068737.html (in French).

⁵⁸This has, for instance, been witnessed at PB meetings in Amsterdam.

the democratic process. It is necessarily not possible to guarantee every citizen to be satisfied with, at least part of, the outcome of the voting stage—essentially because there is only a limited amount of money available. However, when shortlisting projects, there is no hard constraint on what can be selected. It is thus easier to ensure that everyone will have an impact on the shortlist. This is our fourth objective.

Identifying suitable objectives is a necessary first step, but is not a sufficient one. What is needed now is to formalise these requirements, and to operationalise them through well-defined shortlisting rule. The first half of this chapter will aim at answering the following question:

What are good shortlisting rules for the first stage of participatory budgeting?

Later in this chapter, we will introduce several shortlisting rules, presenting different ways the shortlist can be constructed. The actual answer to the question above will be given shortly after, when we will study their respective merits. This will be achieved using the standard toolbox of the computational social choice scientist: the axiomatic method—to assess how shortlisting rules fare with respect to the desirable properties we described above—and the algorithmic approach—to assess their practical usability.

Our second point of focus concerns the interaction between the two stages of our model. More specifically, we will investigate the incentives for the agents to submit, or not, proposals during the shortlisting stage, because of the potential impact on the final budget allocation. Let us exemplify what we mean here.

Consider the case of Sophie. She is a very active citizen who would gladly see a fountain in the middle of the main square of the town. After attending some of the public meetings for the local PB process, she is convinced that this is the perfect opportunity for her project to come to life. Unfortunately, she is not the only one to have ideas about potential improvements of the main square. Adrian—another citizen who shares with Sophie her love for water-based ornamental structures—wants the middle of the square to be occupied by a small lotus-covered pond. On the other hand, Jan, a now retired chess master, would love to be able to go to the square on Sunday afternoons to play chess with the members of his chess club. Overall, the following three projects may be proposed during the shortlisting stage.



Since she attentively attended the public meetings, Sophie is aware of the proposals that will be submitted by Adrian and Jan. Now comes a dilemma: Given what she knows about the other proposals, and what she expects will happen when the voting stage comes, should Sophie submit her fountain proposal or not? Indeed, it is not hard to imagine that voters who enjoy water-based ornamental structures could split between the lotus pond and the fountain (provided that they both get shortlisted), resulting in few votes for each so that the chess boards would be implemented. Assuming that she likes the lotus pond better than the chess boards, Sophie would thus have an incentive not to propose her plan for a fountain, so that Adrian’s lotus pond is selected by a larger set of voters, and would then be implemented.

This very stylised example puts in the spotlight the potential strategic behaviours agents can engage into when voting. The fact that voters can behave strategically, and submit opinions that do not reflect their true inner preferences is generally considered non-desirable. One aspect of the research in social choice has been to look into ways for preventing such behaviours to emerge (Zwicker, 2016; Meir, 2018). Coming back to our example, it is interesting to note that the proposals regarding the lotus pond and the fountain are in some sense very similar. In particular, a shortlisting rule that enforces the fourth objective of the shortlisting stage—avoiding to shortlist projects that are too similar—would potentially prevent Sophie’s strategic move as only one of the two proposals would get shortlisted. This raises the following question:

Are some shortlisting rules better at preventing strategic behaviour than others?

Elements of responses to this question will be provided in the second half of this chapter, when we will delve into *first-stage strategy-proofness*. As a start, we will formally define the idea of strategic behaviour as we discussed above. Once the conceptual work will have been done, we will study how immune to manipulation different shortlisting rules are.

It is now time to look more closely into the points we raised above. After quickly discussing the literature specifically relevant for this chapter, we will introduce the formal model (Section 8.1), and some shortlisting rules (Section 8.2). A comprehensive example of our model will then follow (Section 8.3). Next, we will delve into the more technical parts, first studying the general properties of the shortlisting rules (Section 8.4), and then looking into first-stage strategy-proofness (Section 8.5). We will then summarise our findings (Section 8.6).

Additional Related Work. In addition to the literature review on PB we presented in Chapter 3, we provide here references that are more specific to this chapter.

We study the shortlisting stage by taking inspiration from the literature on multi-winner voting (Lackner and Skowron, 2023). Note that the term “shortlisting” is used there in two different senses: either to emphasise that choosing a set of k candidates is but a first step in making a final decision (Faliszewski, Skowron, Slinko and Talmon, 2017), or to refer to the problem of electing a set of variable size (Kilgour, 2016; Duddy,

Piggins and Zwicker, 2016; Faliszewski, Slinko and Talmon, 2020; Lackner and Maly, 2021). Only the latter is formally related to our approach to shortlisting for PB. Note that the idea of basing shortlisting on clustering techniques (Jain and Dubes, 1988), as we present later in this chapter—the to-be-called median-based shortlisting rule—was already developed by Lackner and Maly (2021).

8.1 The Formal End-to-End Model

We will essentially adopt the notation that we introduced already in Chapter 2. However, we are not yet equipped with notation for the first stage. It will be introduced in the following.

8.1.1 Additional Notation and Terminology

For the end-to-end model, we need to go one step further in our terminology relating to projects. We will denote by $\mathbb{P} = \{p_1, \dots, p_m\}$ the (finite) set of all *conceivable* projects.⁵⁹ Those are projects that can be submitted in the first stage. As we did in the previous chapter, we extend the cost function so that, $c : \mathbb{P} \rightarrow \mathbb{R}_{>0}$ now maps any conceivable project $p \in \mathbb{P}$ to its cost $c(p) \in \mathbb{R}_{>0}$. As usual, the total cost of any set $P \subseteq \mathbb{P}$ is written $c(P) = \sum_{p \in P} c(p)$. We still use $b \in \mathbb{R}_{>0}$ to denote the budget limit and assume without loss of generality that for every project $p \in \mathbb{P}$, we have $c(p) \leq b$.

Throughout this chapter, we are going to require means for breaking ties, both between alternative projects and between alternative sets of projects. We will do so using tie-breaking rules. Those are functions—typically denoted by UNTIE —taking as input a set of tied projects, and returning a single one of them.

The most natural way of breaking ties, when no other information is available, is to do so through the index of the projects. This is captured by the *canonical tie-breaking rule* CANONUNTIE , defined for any $P \subseteq \mathbb{P}$ as:

$$\text{CANONUNTIE}(P) = \arg \min_{p_j \in P} j.$$

We also use tie-breaking rules to transform weak orders over projects into strict orders. Let UNTIE be an arbitrary tie-breaking rule. Take any weak order \succsim on \mathbb{P} . Then for every indifference class $P \subseteq \mathbb{P}$ of \succsim , we break ties as follows: $p = \text{UNTIE}(P)$ is the first project, then comes $\text{UNTIE}(P \setminus \{p\})$, and so forth. Overloading notation, we denote by $\text{UNTIE}(\succsim)$ the strict order thus obtained.

Finally, we extend the canonical tie-breaking rule CANONUNTIE to non-empty sets $\mathfrak{P} \subseteq 2^{\mathbb{P}}$ of sets of projects in a lexicographic manner: $\text{CANONUNTIE}(\mathfrak{P})$ is the unique

⁵⁹Note here that we use the same notation as we used to denote the set of all the projects appearing in a perpetual PB instance (in Chapter 7) These two concepts are somehow similar as they define a set containing “all the projects”.

set $P \in \mathfrak{P}$ such that $\text{CANONUNTIE}(P \triangle P') \in P$ for all $P' \in \mathfrak{P} \setminus \{P\}$.⁶⁰ Thus, we require that, amongst all the projects on which P and P' differ, the one with the lowest index must belong to P . This formalises the way the usual lexicographic tie-breaking is lifted to sets of words (as in a dictionary for instance).

When no tie-breaking rule is specified, it is assumed that ties are broken with respect to CANONUNTIE .

8.1.2 The Shortlisting Stage

In the first stage, agents are asked to propose projects. A *shortlisting instance* is a tuple $\langle \mathbb{P}, c, b \rangle$. Because of bounded rationality, an agent may not be able to conceive of all the projects they would approve of if only she were aware of them. We denote by $C_i \subseteq \mathbb{P}$ the set of projects that agent i can conceive of—their *awareness set*—and we call the vector $\mathbf{C} = (C_1, \dots, C_n)$ the *awareness profile*. We do assume that agent i knows the cost of the projects in C_i as well as the budget limit b .

We denote by $P_i \subseteq C_i$ the set of projects agent $i \in \mathcal{N}$ chooses to actually propose, and we call the resulting vector $\mathbf{P} = (P_1, \dots, P_n)$ a *shortlisting profile*. We use (\mathbf{P}_{-i}, P'_i) to denote the profile we obtain when, starting from profile \mathbf{P} , agent i changes their proposal to P'_i .

A *shortlisting rule* SHORT maps any given shortlisting instance $I = \langle \mathbb{P}, c, b \rangle$ and shortlisting profile \mathbf{P} to a shortlist, *i.e.*, a set $\text{SHORT}(I, \mathbf{P}) \subseteq \bigcup \mathbf{P}$ of shortlisted projects, where $\bigcup \mathbf{P} = P_1 \cup \dots \cup P_n$.

8.1.3 The Allocation Stage

In the second stage, agents vote on the shortlisted projects to decide which ones should get funded. This stage correspond to the standard model that we introduced in Chapter 2. We will adopt the same notation here. To differentiate between the shortlisting and the allocation stage, will be refer to what we used to call PB rules as *allocation rules*. We will sometimes refer to allocation rules we have already defined (in Chapter 2). We will consider their resolute variants, where we break ties according to CANONUNTIE as defined earlier.

8.1.4 Agent Preferences

Later on in this chapter, we will discuss the incentives of the agents in our end-to-end model. To do so, we need ways of discussing their preferences. In the following we present what we assume to be the internal preference model that the agents follow.

Consider a shortlisting instance $\langle \mathbb{P}, c, b \rangle$. We make the assumption that agent $i \in \mathcal{N}$ has preferences over all individual projects in \mathbb{P} —including those they are unaware of—and that those preferences take the form of a strict linear order \triangleright_i over \mathbb{P} . It is

⁶⁰ \triangle is the symmetric difference between sets, defined for any S and S' as $S \triangle S' = (S \setminus S') \cup (S' \setminus S)$.

important to keep in mind that we do not assume that i is aware of \triangleright_i in full. For any subset of projects $P \subseteq \mathbb{P}$, we denote by $\triangleright_{i|P}$ the restriction of \triangleright_i to P . Our second assumption is that for any subset of projects $P \subseteq \mathbb{P}$, agent i is able to determine an *ideal set* of projects, denoted $top_i(P)$, that is determined through the use of the greedy selection procedure that we already introduced in Chapter 2. Formally, we define:

$$top_i(P) = \text{GREED}(\langle P, c, b \rangle, \triangleright_{i|P}).$$

This approach permits us to model what constitutes a *truthful* vote by an agent for varying shortlists P . We call the vector $top(P) = (top_1(P), \dots, top_n(P))$ the *ideal profile* given P .

When investigating the potential strategic behaviour of the agents, we will need to compare different budget allocations from the perspective of the agents. We assume that agents derive preferences over budget allocations from their ideal sets through the use of *completion principles* (Lang and Xia, 2016). A completion principle is a method that, given an ideal point (or top element), generates a weak order over subsets of projects.⁶¹ For any ideal set $top \subseteq \mathbb{P}$, we denote by \succsim_{top} the weak order over subsets of projects induced by a given completion principle (we omit the latter to simplify the notation; it will always be clear from the context). We will denote by \succ_{top} the strict part of \succsim_{top} , and \sim_{top} its indifference part. For instance, if we follow the *cardinality-based* completion principle, then, given an ideal set $top \subseteq \mathbb{P}$, we have $P \succsim_{top} P'$ for every two subsets of projects $P, P' \subseteq \mathbb{P}$ such that $|P \cap top| \geq |P' \cap top|$. Under the *cost-based* completion principle, we have $P \succsim_{top} P'$ for every two subsets of projects $P, P' \subseteq \mathbb{P}$ such that $c(P \cap top) \geq c(P' \cap top)$.

Instead of stating our results for specific completion principles, we will phrase them so that they apply to all completion principles behaving in certain ways. In the following we introduce the different properties we will need. A completion principle generating \succsim_{top} from top is said to satisfy:

- **Top-First** if the ideal point top strictly dominates any other subset of projects: for all $P \subseteq \mathbb{P}$, we have $top \succ_{top} P$;
- **Top-Sufficiency** if the empty set is strictly dominated by any non-empty subset of top : for all $P \subseteq top$, if $P \neq \emptyset$, then we have $P \succ_{top} \emptyset$;
- **Top-Necessity** if any subset of projects that does not intersect with top is treated the same way as the empty set: for all $P \subseteq \mathbb{P}$, if $P \cap top = \emptyset$, then we have $P \sim_{top} \emptyset$;
- **Cost-Neutral Monotonicity** if selecting more projects from top is strictly better than fewer, as long as they all have the same cost: for all $P, P' \subseteq \mathbb{P}$ such that $P \Delta P' \subseteq top$, if $|P \cap top| > |P' \cap top|$ and $c(p) = c(p')$ for any two projects $p, p' \in P \Delta P'$, then we must have $P \succ_{top} P'$.

⁶¹One could also work with partial orders here. All of our definitions would carry over seamlessly in this case. Their interpretation would however differ.

So, a completion principle is top-first if top is indeed the best outcome. It is top-sufficient if it is sufficient to have some projects from top to be better than the empty set. It is top-necessary if it is necessary to have some projects from top to be better than the empty set. Finally, it is cost-neutral monotonic if having more projects from top is better than having less, even if those are different projects, provided that they all have the same cost.

As a warm-up, the reader can check that both the cardinality- and the cost-based completion principles satisfy all of the above properties.

We finally provide our last definition (for this section). For any weak order \succsim_{top} and any family of subsets of projects $\mathfrak{P} \subseteq 2^{\mathcal{P}}$, we use $undom(\succsim_{top}, \mathfrak{P})$ to denote the set of subsets of projects that are undominated in \mathfrak{P} according to \succsim_{top} .

8.2 Shortlisting Rules

As we have already seen, several allocation rules have been defined in the literature. This is however not the case for shortlisting rules. In this section, we therefore propose several of them.

Our first shortlisting rule is what arguably is the simplest of them, the *nomination (shortlisting) rule*. Following this rule, every agent acts as a nominator, *i.e.*, someone whose proposals are always all accepted.

Definition 8.2.1 (Nomination Rule). *The nomination rule NOMSHORT returns, for every shortlisting instance $I = \langle \mathbb{P}, c, b \rangle$ and shortlisting profile \mathbf{P} , the shortlist:*

$$\text{NOMSHORT}(I, \mathbf{P}) = \bigcup \mathbf{P}.$$

Although very natural, the nomination shortlisting rule is not effective in reducing the number of projects. So let us go through some more examples of shortlisting rules.

8.2.1 The Equal Representation Shortlisting Rule

Since the budget limit is not a hard constraint at the shortlisting stage, one of its objectives could be to ensure that every participant has a say in the decision. Building on this idea, we introduce the *k-equal representation shortlisting rule*, a Thiele rule⁶² that maximises the minimum number of selected projects per agent. Here the parameter k determines the maximum cost of the shortlist selected by the rule, expressed as a multiplier of b .

⁶²Thiele rules are multi-winner voting rules that have received substantial attention (Lackner and Skowron, 2023). We briefly sketch their definition here. Let $\mathbf{w} = (w_1, w_2, \dots)$ be an infinite weight vector. Assume we are aiming at selecting exactly $k \in \mathbb{N}_{>0}$ projects. Given a shortlisting instance $I = \langle \mathbb{P}, c, b \rangle$ and a shortlisting profile $\mathbf{P} = (P_1, \dots, P_n)$, the \mathbf{w} -Thiele method is a multi-winner voting rule that selects k -sized subsets of projects P with maximum weight, where the weight of P is defined as $\sum_{i \in \mathcal{N}} \sum_{j=1}^{|P \cap P_i|} w_j$.

Definition 8.2.2 (*k*-Equal Representation Shortlisting Rule). Let $k \in \mathbb{N}$. The *k*-equal representation shortlisting rule REPRSHORT_k is defined for every shortlisting instance $I = \langle \mathbb{P}, c, b \rangle$ and every shortlisting profile $\mathbf{P} = (P_1, \dots, P_n)$ as:

$$\text{REPRSHORT}_k(I, \mathbf{P}) = \text{CANONUNTIE} \left(\arg \max_{\substack{P \subseteq \bigcup \mathbf{P} \\ c(P) \leq k \cdot b}} \sum_{i \in \mathcal{N}} \sum_{\ell=0}^{|P_i \cap P|} \frac{1}{n^\ell} \right).$$

The choice of the weight $1/n^\ell$ in the definition of the rule ensures that the rule will always select a project proposed by the agents with the smallest number of selected projects, who still have unselected projects.

The *k*-equal representation shortlisting rule can be seen as a Thiele rule where the *j*-th weight is defined as $w_j = \sum_{\ell=0}^j 1/n^\ell$. Note that in particular this implies that the weight is dependent on the number of agents.

Let us explore the computational complexity of the rule now.

Proposition 8.2.3. Let $k \in \mathbb{N}_{>0}$. There is no algorithm running in polynomial-time that computes, given a shortlisting instance I and profile \mathbf{P} , the outcome of the *k*-equal representation shortlisting rule, unless $\text{P} = \text{NP}$.

PROOF. Note that for $k' \in \mathbb{N}_{>0}$, if all projects have cost b/k' , computing the outcome of the *k*-equal representation shortlisting rule amounts to finding a committee of size k' with a Thiele rule with weights $(1, 1/n, 1/n^2, \dots)$ for a multi-winner election (Janson, 2016). Interestingly, the reduction presented by Aziz, Gaspers, Gudmundsson, Mackenzie, Mattei and Walsh (2015) to show that proportional approval voting is NP-complete works for all Thiele rules with decreasing weights. Since this is the case here, their reduction also applies. \square

8.2.2 Median-Based Shortlisting Rules

One criterion frequently used for shortlisting in real PB processes is the similarity between the proposals. Since only few projects will be shortlisted, it would be particularly inefficient to shortlist two very similar ones. In the following we rationalise this decision process by introducing a shortlisting rule that clusters the projects and selects representative projects for each cluster.

We assume that projects are embedded in a metric space, the distance between two projects being given. Using this metric space, we will try and cluster the proposals submitted during the shortlisting stage.

Formally speaking, we call *distance* any metric over \mathbb{P} . For a distance δ , let $\text{med}(P)$ be the the *geometric median* of $P \subseteq \mathbb{P}$ defined by:

$$\text{med}(P) = \text{CANONUNTIE} \left(\left\{ \{p\} \mid p \in \arg \min_{p^* \in P} \sum_{p' \in P} \delta(p^*, p') \right\} \right).$$

A partition of P , denoted by $V = \{V_1, \dots, V_p\}$, is a (k, ℓ) -Voronoi partition with respect to the distance δ , if the representatives of V cost no more than $k \cdot b$ in total:

$$\sum_{V_j \in V} c(\text{med}(V_j)) \leq k \cdot b,$$

and for every distinct $V_j, V_{j'} \in V$ and every project $p \in V_j$, we have:

- ▶ $\delta(p, \text{med}(V_j)) \leq \delta(p, \text{med}(V_{j'}))$, i.e., every project is in the cluster of its closest geometric median; and
- ▶ $\delta(p, \text{med}(V_j)) \leq \ell$, i.e., p is within distance ℓ of $\text{med}(V_j)$.

Thus, the parameter k bounds the total cost of the representatives, and the parameter ℓ bounds the maximum distance within a cluster. We denote by $\mathcal{V}_{\delta, k, \ell}(P)$ be the set of all (k, ℓ) -Voronoi partitions of P with respect to distance δ and parameters k and ℓ .

With all these definitions in mind, we are ready to define the class of k -median shortlisting rules. As before, k parametrised the maximum cost of the shortlist.

Definition 8.2.4 (*k*-Median Shortlisting Rules). *Let $k \in \mathbb{N}$. The k -median shortlisting rule $\text{MEDIANSHORT}_{k, \delta}$ with respect to the distance δ is such that for every shortlisting instance I and profile \mathbf{P} , we have:*

$$\text{MEDIANSHORT}_{k, \delta}(I, \mathbf{P}) = \text{CANONUNTIE} \left(\left\{ \bigcup_{V_j \in V} \text{med}(V_j) \mid V \in \mathcal{V}_{\delta, k, \ell^*}(\bigcup \mathbf{P}) \right\} \right)$$

where ℓ^* is the smallest ℓ such that $\mathcal{V}_{\delta, k, \ell}(\bigcup \mathbf{P}) \neq \emptyset$.

Note that we chose to minimise ℓ in our definition. One can similarly try to minimise k , or both ℓ and k , instead.

It is finally worth saying a few words about the computation complexity of these shortlisting rules. It is straightforward to show that, unless $P = NP$, there cannot be an algorithm that runs in polynomial time and that computes the outcome of a k -median shortlisting rule, for any value of k and suitable distance δ . Indeed, when δ is the Euclidean distance over \mathbb{R}^2 , our formulation coincide with the k -median problem, known to be NP-hard (Kariv and Hakimi, 1979). Several other results have been published, including approximation algorithms (Kanungo, Mount, Netanyahu, Piatko, Silverman and Wu, 2004) and fixed-parameters analyses (Cohen-Addad, Gupta, Kumar, Lee and Li, 2019), and can be used to cope with intractability.

8.3 End-to-End Example

We have now introduced all the components of our model. Before getting to the more technical analysis, let us give an example to clarify the whole setting.

Consider the following shortlisting instance $I = \langle \mathbb{P}, c, b \rangle$ with nine projects, $\mathbb{P} = \{p_1, \dots, p_9\}$. Suppose for simplicity that for every project $p \in \mathbb{P}$ we have $c(p) = 1$, *i.e.*, we are in the unit-cost setting. The budget limit is $b = 3$. Consider five agents as described below.

	Preferences over the Projects	Awareness set C_i	Ideal Set Based on C_i
Agent 1	$p_4 \triangleright p_5 \triangleright p_1 \triangleright p_2 \triangleright \dots$	$\{p_1, p_2, p_4, p_5\}$	$\{p_1, p_4, p_5\}$
Agent 2	$p_1 \triangleright p_2 \triangleright p_6 \triangleright p_4 \triangleright \dots$	$\{p_2, p_6\}$	$\{p_2, p_6\}$
Agent 3	$p_1 \triangleright p_2 \triangleright p_7 \triangleright p_4 \triangleright \dots$	$\{p_2, p_7\}$	$\{p_2, p_7\}$
Agent 4	$p_1 \triangleright p_3 \triangleright p_8 \triangleright p_5 \triangleright \dots$	$\{p_3, p_8\}$	$\{p_3, p_8\}$
Agent 5	$p_1 \triangleright p_3 \triangleright p_9 \triangleright p_5 \triangleright \dots$	$\{p_3, p_9\}$	$\{p_3, p_9\}$

Assuming agents are truthful they will propose projects according to their ideal sets, computed given their awareness sets. The truthful shortlisting profile would then be:

$$\mathbf{P} = (\{p_1, p_4, p_5\}, \{p_2, p_6\}, \{p_2, p_7\}, \{p_3, p_8\}, \{p_3, p_9\}).$$

Thus, if the nomination shortlisting rule **NOMSHORT** is used, the shortlist would be \mathbb{P} . In case the 1-equal representation rule **REPRSHORT₁** is used, it would be $\{p_1, p_2, p_3\}$.

Suppose the set of shortlisted projects is $\mathcal{P} = \mathbb{P}$. Agents are now aware of all the shortlisted projects. They recompute their ideal sets given the new information. Still assuming that agents behave truthfully, the profile for the allocation stage is:

$$\mathbf{A} = (\{p_1, p_4, p_5\}, \{p_1, p_2, p_6\}, \{p_1, p_2, p_7\}, \{p_1, p_3, p_8\}, \{p_1, p_3, p_9\}).$$

This corresponds to the vector of the ideal sets computed by each agent with respect to \mathcal{P} , and their respective preferences over the projects. With such a profile, if the allocation rule **GREEDCOST** is used, the final budget allocation would be $\pi = \{p_1, p_2, p_3\}$.

8.4 Axioms for Shortlisting Rules

We now assess the axiomatic merits of the shortlisting rules we have introduced.

The first axiom we define is *non-wastefulness*. It requires that no amount of the budget should be wasted because not enough projects were shortlisted.

Definition 8.4.1 (Non-Wastefulness). *A shortlisting rule **SHORT** is non-wasteful if for every shortlisting instance $I = \langle \mathbb{P}, c, b \rangle$ and profile \mathbf{P} , one of the following two holds:*

$$c(\text{SHORT}(I, \mathbf{P})) \geq b \quad \text{or} \quad \text{SHORT}(I, \mathbf{P}) = \bigcup \mathbf{P}.$$

This axioms can be interpreted as an efficiency requirement ensuring that no money is wasted because of the shortlisting rule.

We believe that another important property of a shortlisting rule is that every agent is represented in the outcome. This is particularly relevant in the shortlisting stage since any subset of \mathbb{P} is theoretically admissible.

Definition 8.4.2 (Representation Efficiency). *For a given shortlisting instance $I = \langle \mathbb{P}, c, b \rangle$ and a given shortlisting profile \mathbf{P} , a set of projects $P \subseteq \mathbb{P}$ is representatively dominated if there is a set $P' \subseteq \mathbb{P}$ with $c(P') \leq c(P)$, and $|P' \cap P_i| \geq |P \cap P_i|$ for all $i \in \mathcal{N}$, with a strict inequality for at least one agent.*

A shortlisting rule SHORT is representatively efficient if for every shortlisting instance $I = \langle \mathbb{P}, c, b \rangle$, and every shortlisting profile \mathbf{P} , $\text{SHORT}(I, \mathbf{P})$ is not representatively dominated by any other subset of projects.

A set of projects P is thus representatively dominated by another one P' if P' does not cost more than P , and, for every agent, at least as many projects that they submitted have been selected in P' as in P , and strictly more for at least one of them.

This axiom provides guarantees that the shortlisting rule is aiming to achieve some kind of representation. Note however that the guarantee is not very strong and can lead to large disparities between the agents: some could have all their proposals shortlisted, and some others none, in a shortlist that is still representatively efficient.

These are the two axioms with respect to which we will analyse the shortlisting rules. This analysis is presented below.

We will start with the nomination shortlisting rule, that trivially satisfies both non-wastefulness and representation efficiency.

Proposition 8.4.3. *The nomination shortlisting rule is non-wasteful and representatively efficient.*

PROOF. With the nomination shortlisting rule NOMSHORT , for every shortlisting instance I and profile \mathbf{P} , we have $\text{NOMSHORT}(I, \mathbf{P}) = \bigcup \mathbf{P}$. Thus, the second condition of non-wastefulness is always trivially satisfied.

Given that every project is shortlisted, the shortlist $\text{NOMSHORT}(I, \mathbf{P})$ cannot be representatively dominated. NOMSHORT is thus representatively efficient. \square

We now move to the k -equal representation shortlisting rule, showing that it is both non-wasteful and representatively efficient, as long as k is at least 2.

Proposition 8.4.4. *For every $k \geq 2$, the k -equal representation shortlisting rule is non-wasteful, but it is not for $k = 1$. Moreover, for every $k \geq 1$, the k -equal representation shortlisting rule is representatively efficient.*

PROOF. Let us first prove that for every $k \geq 2$, the k -equal representation shortlisting rule SHORT is non-wasteful. Suppose it is not, then, there would exist a shortlisting instance $I = \langle \mathbf{P}, c, b \rangle$ and a shortlisting profile \mathbf{P} such that:

$$c(\text{REPRSHORT}_k(I, \mathbf{P})) < b \quad \text{and} \quad \text{REPRSHORT}_k(I, \mathbf{P}) \neq \bigcup \mathbf{P}.$$

From this, we know that there exists a project $p \in \bigcup \mathbf{P}$ that has not been shortlisted, *i.e.*, such that $p \notin \text{REPRSHORT}_k(I, \mathbf{P})$. Thus, the representation score of the set $\text{REPRSHORT}_k(I, \mathbf{P}) \cup \{p\}$ is higher than that of $\text{SHORT}(I, \mathbf{P})$. Moreover, for any $k \geq 2$, the facts that $c(p) \leq b$ and $c(\text{REPRSHORT}_k(I, \mathbf{P})) < b$ together imply that:

$$c(\text{SHORT}(I, \mathbf{P}) \cup \{p\}) \leq 2 \cdot b \leq k \cdot b.$$

Overall, if such a project p exists, then $\text{REPRSHORT}_k(I, \mathbf{P}) \cup \{p\}$ is an admissible outcome of REPRSHORT_k with a higher total weight than $\text{REPRSHORT}_k(I, \mathbf{P})$. This contradicts the definition of REPRSHORT_k , and thus proves that it is non-wasteful.

Note that for $k = 1$, the definition of REPRSHORT_k implies that the cost of the shortlist will not be more than b (Definition 8.2.2), and non-wastefulness requires the same cost to be at least b (Definition 8.4.1). Since projects are indivisible, it is clearly not always possible to shortlist a set of projects of cost exactly b .

We now show that for every $k \geq 1$, the k -equal representation shortlisting rule is representatively efficient. The proof is actually trivial, it is immediately derived from the choice of the weight $1/n$ in the definition of the rule. Indeed, since $1/n > 0$, a representatively dominated shortlist would always have a lower total score and would thus not be selected. \square

Now comes the turn of the median shortlisting rules. We prove that these rules are non-wasteful, but not representatively efficient.

Proposition 8.4.5. *Let δ be an arbitrary distance over \mathbb{P} . The following facts hold:*

- For every $k \geq 2$, the shortlisting rule $\text{MEDIANSHORT}_{k,\delta}$ is non-wasteful;
- There exists no $k \in \mathbb{N}_{>0}$ such that $\text{MEDIANSHORT}_{k,\delta}$ is representatively efficient.

PROOF. The proof that for $k \geq 2$, the $\text{MEDIANSHORT}_{k,\delta}$ is non-wasteful is similar to that of Proposition 8.4.4. To see why, note that for every shortlisting instance I and profile \mathbf{P} , there will never be an unselected project $p \in \bigcup \mathbf{P} \setminus \text{SHORT}(I, \mathbf{P})$ such that $c(\text{SHORT}(I, \mathbf{P}) \cup \{p\}) \leq k \cdot b$. Indeed, if such a p exists, selecting it would always lead to a smaller within-cluster distance, simply by including p as its own cluster (since by the definition of a metric, the distance between p and any other project $p' \in \mathbb{P} \setminus \{p\}$ is non-zero). We can thus reach the same contradiction that we reached in the proof of Proposition 8.4.4.

It is also easy to see that $\text{MEDIANSHORT}_{k,\delta}$ is not efficiently representative. We do not provide a formal proof here as the correctness of the statement should be intuitively clear. It is derived from the fact that for any shortlisting profiles \mathbf{P} and \mathbf{P}' , such that $\bigcup \mathbf{P} = \bigcup \mathbf{P}'$, the outcome of $\text{MEDIANSHORT}_{k,\delta}$ would be the

same. This means that $\text{MEDIANSHORT}_{k,\delta}$ is completely oblivious of the identity of the agents, making it fail representation efficiency. \square

The axiomatic analysis of the shortlisting rules is now complete. We will move on to the next focus point of this chapter: the interactions between the two stages.

8.5 First-Stage Strategy-Proofness

We now turn to the analysis of strategic interactions during the shortlisting stage. Remember our motivational example in the introduction, we wondered whether Sophie should propose her project about the fountain or not, because of its impact on the final decision, taken during the second stage. We have hinted at reasons why it would actually be better for her not to. Throughout this section, we will study such strategic behaviour and investigate whether it can be prevented or not.

One of the main challenges to formalise the concept of strategic behaviour during the first stage, is that agents actually reason about the outcome of the process—the final budget allocation—that is only decided one stage later. Let us then take the time to discuss the information available to an agent willing to strategise, the *manipulator*.

In the classical voting framework (Zwicker, 2016), it is assumed that the potential manipulator has access to all the other ballots before submitting their own. In our setting, when considering a manipulator choosing which proposal to submit during the first stage, the same assumption is reasonable with respect to the proposals submitted by the other agents during the first stage, but not with respect to the ballots the other agents are going to submit during the second stage, only *after* the shortlist will have been determined. Indeed, the set of actions for the second stage depends on the proposal of the manipulator in the first stage. We thus need to reason about the outcome of the second stage given the profile that the manipulator *expects* to occur.

We explore three possibilities. In the first two cases, a manipulator in the first stage is unsure what will happen during the second stage, but assumes that either the worst scenario will be realised (*pessimistic manipulation*) or the best one (*optimistic manipulation*). In the third case, they know the other agents' true preferences and trusts they will vote accordingly (*anticipatory manipulation*).

Because there are no reasons to assume that a potential manipulator would only behave strategically in the first stage, and not in the second stage, we also need the concept of a best response in the second stage. For that, we introduce some further notation. For a given allocation rule F , allocation instance $I = \langle \mathcal{P}, c, b \rangle$, profile \mathbf{A} , and agent $i \in \mathcal{N}$, let $A_i^*(I, \mathbf{A}, F)$ be the best response of i to \mathbf{A} , defined such that:

$$A_i^*(I, \mathbf{A}, F) = \text{CANONUNTIE}(\{A'_i \in \mathcal{P} \mid F(I, (\mathbf{A}_{-i}, A'_i)) = \text{CANONUNTIE}(\mathfrak{P}^*)\}),$$

where $\mathfrak{P}^* = \text{undom}(\succsim_{\text{top}_i(\mathcal{P})}, \{F(I, (\mathbf{A}_{-i}, P)) \mid P \subseteq \mathcal{P}\})$ and \succsim is generated given an arbitrary completion principle.

Let us unravel a bit this definition. $\succsim_{top_i(\mathcal{P})}$ is the weak order over subsets of projects that is induced by the completion principle in use, based on $top_i(\mathcal{P})$, the ideal set of agent i . Then, \mathfrak{A}^* is the set of undominated budget allocations returned by the allocation rule F , for any approval ballot P , agent i can submit in the second stage, where domination is defined with respect to $\succsim_{top_i(\mathcal{P})}$. Because we need a single outcome, we break ties between the, potentially, several such undominated budget allocation. Then, $A_i^*(I, \mathbf{A}, F)$ is the ballot that achieved this aforementioned budget allocation and that is selected by the tie-breaking rule (since the relevant may be reachable via several ballots). The intuition is that $A_i^*(I, \mathbf{A}, F)$ is the best ballot agent i can submit in the second stage given I , \mathbf{A} , and F , and our assumptions on the preferences of the agents. When clear from the context, we omit I , \mathbf{A} , and/or F from the notation $A_i^*(I, \mathbf{A}, F)$.

We are now ready to properly formalise all we described above. This is the aim of the following definition.

Definition 8.5.1 (Successful Manipulation). *Let SHORT be a shortlisting rule, F an allocation rule, $I_1 = \langle \mathbb{P}, c, b \rangle$ a shortlisting instance, \mathbf{P} a shortlisting profile, and $P'_i \subseteq \mathbb{P}$ an alternative proposal for agent $i \in \mathcal{N}$. Consider the shortlists $\mathcal{P} = \text{SHORT}(I_1, \mathbf{P})$ and $\mathcal{P}' = \text{SHORT}(I_1, (\mathbf{P}_{-i}, P'_i))$, determining the allocation instances $I_2 = \langle \mathcal{P}, c, b \rangle$ and $I'_2 = \langle \mathcal{P}', c, b \rangle$.*

For any two profiles, \mathbf{A} for \mathcal{P} and \mathbf{A}' for \mathcal{P}' , we simplify the notation by defining the two following abbreviations:

$$F^*(I_2, \mathbf{A}) = F(I_2, (\mathbf{A}_{-i}, A_i^*(I_2, \mathbf{A}))),$$

$$F^*(I'_2, \mathbf{A}') = F(I'_2, (\mathbf{A}'_{-i}, A_i^*(I'_2, \mathbf{A}'))).$$

To clarify, $F^(I_2, \mathbf{A})$ is thus the final budget allocation for the instance I_2 and profile \mathbf{A} in which agent i is playing their best response. The case of $F^*(I'_2, \mathbf{A}')$ is analogous, for the instance I'_2 and profile \mathbf{A}' .*

Then, for a given completion principle generating $\succsim_{top_i(\mathcal{P} \cup \mathcal{P}')}$, we say that:

- ▶ P'_i is a successful pessimistic manipulation if, for all profiles \mathbf{A} on \mathcal{P} and \mathbf{A}' on \mathcal{P}' , it is the case that $F^*(I'_2, \mathbf{A}') \succsim_{top_i(\mathcal{P} \cup \mathcal{P}')} F^*(I_2, \mathbf{A})$, with a strict preference for at least one pair $(\mathbf{A}, \mathbf{A}')$.
- ▶ P'_i is a successful optimistic manipulation if, for at least one profile \mathbf{A} on \mathcal{P} and one profile \mathbf{A}' on \mathcal{P}' , it is the case that $F^*(I'_2, \mathbf{A}') \succ_{top_i(\mathcal{P} \cup \mathcal{P}')} F^*(I_2, \mathbf{A})$.
- ▶ P'_i is a successful anticipatory manipulation if, for the two profiles $\mathbf{A} = \mathbf{top}(\mathcal{P})$ and $\mathbf{A}' = \mathbf{top}(\mathcal{P}')$, it is the case that $F^*(I'_2, \mathbf{A}') \succ_{top_i(\mathcal{P} \cup \mathcal{P}')} F^*(I_2, \mathbf{A})$.

Thus, a pessimist is pessimistic with respect to the advantages they can gain from manipulating: assuming the best if she is truthful and the worst otherwise. For optimists it is the other way around. Finally, an anticipatory manipulator knows everyone's preferences on both \mathcal{P} and \mathcal{P}' and uses them to predict their votes for the second stage.

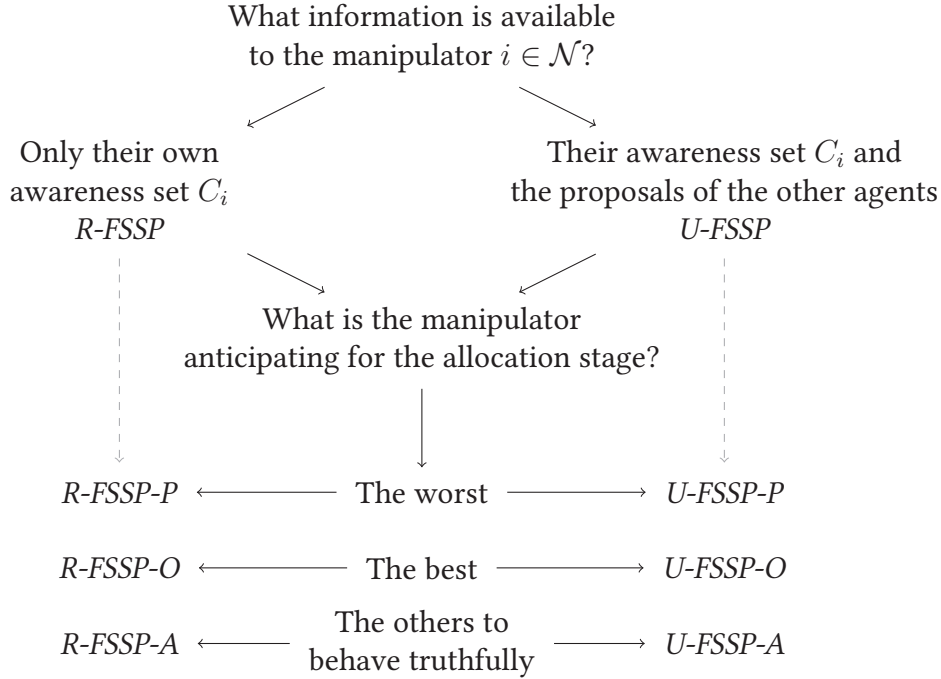


Figure 8.5.1: Explanation of first-stage strategy-proofness concepts.

We are looking for rules that do not allow for successful manipulation, *i.e.*, that are *first-stage strategy-proof* (FSSP). We distinguish two cases: either the manipulator is *restricted* to their awareness set (R-FSSP) or they can also propose any of the projects proposed by others (during the shortlisting stage), *i.e.*, they are *unrestricted* (U-FSSP).

Definition 8.5.2 (First-Stage Strategy-Proofness). *For a given completion principle, a pair $\langle \text{SHORT}, \text{F} \rangle$ consisting of a shortlisting rule SHORT and an allocation rule F is said to be restricted-FSSP (R-FSSP) with respect to a given type of manipulation if for every shortlisting instance $\langle \mathbb{P}, c, b \rangle$, every awareness profile $\mathcal{C} = (C_1, \dots, C_n)$, every shortlisting profile $\mathbf{P} = (P_1, \dots, P_n)$ where $P_{i'} \subseteq C_{i'}$ for all $i' \in \mathcal{N}$, and every agent $i \in \mathcal{N}$, there is no $P'_i \subseteq C_i$ such that submitting P'_i instead of $\text{top}_i(C_i)$ is a successful manipulation for i .*

In case $P'_i \subseteq C_i \cup \bigcup \mathbf{P}$ and we consider $\text{top}_i(C_i \cup \bigcup \mathbf{P})$ instead of $\text{top}_i(C_i)$ in the above, we say that $\langle \text{SHORT}, \text{F} \rangle$ is unrestricted-FSSP (U-FSSP).

Thus, in the unrestricted case, agents are assumed to *first* gain access to everyone's proposals and *then* decide whether or not to vote truthfully.

We introduce some further abbreviations. Let FSSP-P stand for FSSP with respect to pessimistic manipulation attempts, FSSP-O for FSSP with respect to optimistic manipulation attempts, and FSSP-A for FSSP with respect to anticipatory manipulation attempts. We have thus introduced six different FSSP concepts in total. A simplified overview is given in Figure 8.5.1 to clarify everything.

It should be clear, at least from the text around the definitions that there are some links between the different FSSP concepts we have introduced. The following result summarises how the different notions introduced relate to each other, where \mathfrak{X} implying \mathfrak{X}' means that any pair $\langle \text{SHORT}, \text{F} \rangle$ satisfying \mathfrak{X} also satisfies \mathfrak{X}' .

Proposition 8.5.3. *The following implications hold for any given completion principle:*

- ▶ *R-FSSP-O implies R-FSSP-A and R-FSSP-P.*
- ▶ *U-FSSP-O implies U-FSSP-A and U-FSSP-P.*
- ▶ *R-FSSP implies U-FSSP for all types of manipulation.*

PROOF. The first two claims are immediately derived from the relevant definitions. To see that the last of these claims is also true, observe that U-FSSP is a special case of R-FSSP, namely when the manipulator can conceive of all the proposed projects, *i.e.*, when $C_i = \bigcup \mathbf{P}$. \square

Interestingly, the link between pessimistic and anticipatory manipulations is not clear. Although if a successful pessimistic manipulation exists it ensures that an anticipative manipulation would lead to a weakly better outcome, nothing guarantees that this outcome would be strictly better for the manipulator.

8.5.1 Awareness-Restricted Manipulation

We start by proving an impossibility theorem stating that no pair of reasonable rules can be first-stage strategy-proof when manipulators are restricted to their awareness sets. By “reasonable rule” we mean a non-wasteful shortlisting rule, followed by a *determined* allocation rule.

Definition 8.5.4 (Determined). *An allocation rule F is determined if, for every allocation instance $I = \langle \mathcal{P}, c, b \rangle$, and every profile \mathbf{A} , we have $\text{F}(I, \mathbf{A}) \neq \emptyset$.*

Theorem 8.5.5. *Every pair $\langle \text{SHORT}, \text{F} \rangle$ of a non-wasteful shortlisting rule SHORT and a determined allocation rule F is neither R-FSSP-P nor R-FSSP-A (and thus also not R-FSSP-O), for any completion principle that is top-first.*

PROOF. We provide a proof for R-FSSP-P, but the same proof also goes through for R-FSSP-A. The claim for R-FSSP-O then follows from Proposition 8.5.3.

Let $I = \langle \mathbb{P}, c, b \rangle$ be the shortlisting instance with two conceivable projects p_1 and p_2 , both of cost 1, and a budget limit $b = 1$. Suppose there are two agents. The preferences of the first agent are such that $p_2 \triangleright_1 p_1$, their awareness set is

$C_1 = \{p_1\}$. For the second agent, we have $p_1 \triangleright_2 p_2$, and $C_2 = \{p_2\}$. Overall, each agent is aware only of the project they like less. The truthful shortlisting profile is then $\mathbf{P} = (\{p_1\}, \{p_2\})$.

Assuming that SHORT is non-wasteful, we know that $|\text{SHORT}(I, \mathbf{P})| \geq 1$. There are thus three possible cases for $\text{SHORT}(I, \mathbf{P})$: to shortlist either just p_1 , just p_2 , or both p_1 and p_2 . Let us go through each of them independently.

In case $\text{SHORT}(I, \mathbf{P}) = \{p_1\}$, whichever way the agents vote in the allocation stage, as F is assumed to be determined, the final budget allocation must be $\{p_1\}$. Now, if agent 1 manipulates by not proposing any project during the first stage, only project $\{p_2\}$ will get shortlisted (and this has to happen since SHORT is non-wasteful). In that case, $\{p_2\}$ will also be the final budget allocation, given that F is determined. Since $\{p_2\}$ is the ideal point of agent 1 for the set of projects $\{p_1, p_2\}$, they would strictly prefer $\{p_2\}$ over $\{p_1\}$ for any completion principle that is top-first. So agent 1 has an incentive to pessimistically manipulate.

The case of $\text{SHORT}(I, \mathbf{P}) = \{p_2\}$ is perfectly analogous to the previous one: the final budget allocation under truthful behaviour would be $\{p_2\}$, but then agent 2 has an incentive to pessimistically manipulate by not submitting p_2 during the first stage so that the final outcome would be $\{p_1\}$.

Finally, consider the case $\text{SHORT}(I, \mathbf{P}) = \{p_1, p_2\}$. Suppose the final budget allocation is $\{p_1\}$ in case both agents vote truthfully. Then, just as in the first case, agent 1 has an incentive to submit an empty set of proposals instead, as that guarantees a final budget allocation of $\{p_2\}$. In the analogous case where the final budget allocation is $\{p_2\}$, agent 2 would pessimistically manipulate.

Overall, there always is an agent who has an incentive to pessimistically manipulate. $\langle \text{SHORT}, F \rangle$ is thus not R-FSSP-P. \square

Note that the scenario used in the proof shows that unrestricted-FSSP does *not* imply restricted-FSSP. Indeed, under U-FSSP, no agent would have an incentive to manipulate in this scenario, as they would have all the information they need to submit an optimal truthful proposal.

Regarding the specific shortlisting rules we have introduced, we can now derive the following corollary.

Corollary 8.5.6. *Let F be an allocation rule that is exhaustive, $k \geq 2$, and δ an arbitrary distance over \mathbb{P} . Then, none of the pairs $\langle \text{REPRSHORT}_k, F \rangle$, $\langle \text{MEDIANSHORT}_{k,\delta} \rangle$ or $\langle \text{NOMSHORT}, F \rangle$ are R-FSSP-P, R-FSSP-A, or R-FSSP-O, for any completion principle that is top-first.*

PROOF. The proof is immediately derived from Theorem 8.5.5 and the fact that the relevant shortlisting rules are non-wasteful (Propositions 8.4.3, 8.4.4 and 8.4.5). \square

8.5.2 Unrestricted Manipulation

We now turn to the case where the manipulator gains awareness by looking at the projects already submitted for the first stage.

Let us start with the nomination rule. We will show that it is immune to pessimistic manipulation when used with allocation rules that are *unanimous*, a new axiom we introduce below.

Definition 8.5.7 (Unanimity). *An allocation rule F is unanimous if, for every allocation instance $I = \langle \mathcal{P}, c, b \rangle$ and every feasible subset of projects $A \in \text{FEAS}(I)$, it is the case that for the profile $\mathbf{A} = (A, \dots, A)$, we have:*

$$F(I, \mathbf{A}) \supseteq A.$$

This axiom states that if every agent submits the same feasible ballot, then the set of projects in this ballot should be part of the outcome. This is a rather weak axiom and every allocation rule we have defined satisfies it.⁶³

Let us now state our result for the nomination shortlisting rule.

Proposition 8.5.8. *The pair $\langle \text{NOMSHORT}, F \rangle$ where F is an allocation rule that is unanimous is U-FSSP-P for every completion principle that is top-first.*

PROOF. Let $I = \langle \mathbb{P}, c, b \rangle$ be a shortlisting instance, \mathbf{C} an awareness profile, and \mathbf{P} a shortlisting profile. Consider an agent $i^* \in \mathcal{N}$ and let $P_{i^*} = \text{top}_{i^*}(C_{i^*} \cup \bigcup \mathbf{P}_{-i^*})$. Denote $\mathcal{P} = \text{NOMSHORT}(I, (\mathbf{P}_{-i^*}, P_{i^*}))$, and observe that $P_{i^*} \subseteq \mathcal{P}$ because of the definition of NOMSHORT. Moreover, from the definition of NOMSHORT, we know that if agent i^* submits P'_{i^*} instead of P_{i^*} , with $P'_{i^*} \neq P_{i^*}$ the shortlist will become $\mathcal{P}' = P'_{i^*} \cup \left(\bigcup_{i \in \mathcal{N} \setminus \{i^*\}} P_i \right)$. This implies that:

$$\mathcal{P} \cap P_{i^*} \supseteq \mathcal{P}' \cap P_{i^*}, \quad (8.1)$$

keeping in mind that P_{i^*} is the ideal set of i^* for $C_{i^*} \cup \bigcup \mathbf{P}$. Thus, under $(\mathbf{P}_{-i^*}, P'_{i^*})$, it cannot be that more projects from $\text{top}_{i^*}(C_{i^*} \cup \bigcup \mathbf{P})$ are shortlisted than under \mathbf{P} .

We focus on the second stage now. Consider first the profile \mathbf{A} in which all agents submit $\text{top}_{i^*}(\mathcal{P})$. Clearly, submitting $\text{top}_{i^*}(\mathcal{P})$ is a best response for i^* in this case, so $A_{i^*}^*(\langle \mathcal{P}, c, b \rangle, \mathbf{A}) = \text{top}_{i^*}(\mathcal{P})$. Since F is unanimous, we thus have $F(\langle \mathcal{P}, c, b \rangle, \mathbf{A}) = \text{top}_{i^*}(\mathcal{P})$. Now, since $\mathcal{P} \cup \mathcal{P}' \subseteq C_{i^*} \cup \bigcup \mathbf{P}$ and $P_{i^*} \subseteq \mathcal{P}$, we know that:

$$\text{top}_{i^*}(\mathcal{P} \cup \mathcal{P}') = P_{i^*} = \text{top}_{i^*}(\mathcal{P}).$$

⁶³Note how this definition differs from that of *project-wise unanimity* we introduced in Definition 4.4.14. In particular, unanimity is significantly weaker than the latter since it applies only to unanimous profiles.

Given that we assumed the completion principle to be top-first, it is clear that no budget allocation $\pi \in \text{FEAS}(\langle \mathcal{P}', c, b \rangle)$ will be strictly preferred to $F(\langle \mathcal{P}, c, b \rangle, \mathbf{A})$ by i^* . This directly implies that submitting P'_{i^*} cannot be a successful pessimistic manipulation for i^* . \square

Because anticipative manipulation is defined for one very specific profile of the second stage, it is harder to get general results. Still, we can show that any pair consisting of the nomination shortlisting rule **NOMSHORT** and one of the allocation rules we have introduced satisfy **U-FSSP-A**.

Proposition 8.5.9. *The pair $\langle \text{NOMSHORT}, F \rangle$ is not **U-FSSP-A**, and thus not **U-FSSP-O**, when F is one of **GREEDCARD**, **GREEDCOST**, **MAXCARD**, **MAXCOST**, **SEQPHRAG**, **MAXIMINSUPP**, $\text{MES}[\text{sat}^{\text{card}}]$, or $\text{MES}[\text{sat}^{\text{cost}}]$, for every completion principle that satisfies cost-neutral monotonicity.*

PROOF. Recall our initial example presented in Section 8.3. In case agents submit their proposal for the first stage truthfully, the shortlist under **NOMSHORT** is $\mathcal{P} = \mathbb{P} = \{p_1, \dots, p_9\}$. In this case, the truthful profile for the allocation stage is:

$$\mathbf{A} = (\{p_1, p_4, p_5\}, \{p_1, p_2, p_6\}, \{p_1, p_2, p_7\}, \{p_1, p_3, p_8\}, \{p_1, p_3, p_9\}).$$

For the allocation instance $I = \langle \mathcal{P}, c, b \rangle$, the different rules we are considering produce the following outcomes:

$$\begin{aligned} \text{GREEDCARD}(I, \mathbf{A}) &= \text{GREEDCOST}(I, \mathbf{A}) = \text{MAXCARD}(I, \mathbf{A}) = \text{MAXCOST}(I, \mathbf{A}) \\ &= \text{SEQPHRAG}(I, \mathbf{A}) = \text{MAXIMINSUPP}(I, \mathbf{A}) = \{p_1, p_2, p_3\}, \\ \text{MES}[\text{sat}^{\text{card}}](I, \mathbf{A}) &= \text{MES}[\text{sat}^{\text{cost}}](I, \mathbf{A}) = \{p_1\}. \end{aligned}$$

Suppose now that agent 1 submits $P'_1 = \{p_4, p_5\}$ instead of $P_1 = \{p_1, p_4, p_5\}$ in the shortlisting stage. The shortlist computed by **NOMSHORT** then becomes $\mathcal{P}' = \mathbb{P} \setminus \{p_1\}$. After the agents have recomputed their ideal set, and assuming that they behave truthfully, the profile in the second stage would then be:

$$\mathbf{A}' = (\{p_4, p_5\}, \{p_2, p_4, p_6\}, \{p_2, p_4, p_7\}, \{p_3, p_5, p_8\}, \{p_3, p_5, p_9\}).$$

For the new allocation instance $I' = \langle \mathcal{P}', c, b \rangle$, the different rules we are considering produce the following outcomes:

$$\begin{aligned} \text{GREEDCARD}(I, \mathbf{A}') &= \text{GREEDCOST}(I, \mathbf{A}') = \text{MAXCARD}(I, \mathbf{A}') \\ &= \text{MAXCOST}(I, \mathbf{A}') = \text{SEQPHRAG}(I, \mathbf{A}') \\ &= \text{MAXIMINSUPP}(I, \mathbf{A}') = \{p_2, p_4, p_5\}, \\ \text{MES}[\text{sat}^{\text{card}}](I, \mathbf{A}') &= \text{MES}[\text{sat}^{\text{cost}}](I, \mathbf{A}') = \{p_4, p_5\}. \end{aligned}$$

Let us now check that P'_1 is a successful anticipative manipulation for agent 1. Their ideal set across the two scenarios is $top_1(\mathcal{P} \cup \mathcal{P}') = \{p_1, p_4, p_5\}$. Thus in the first case only p_1 is in the intersection of the outcomes of the rules with $top_1(\mathcal{P} \cup \mathcal{P}')$. In the second case—when agent 1 submits P'_1 —this intersection includes both p_4 and p_5 . Given that we assumed the completion principle be cost-neutral monotonic, the anticipative manipulation of agent 1 is thus successful.

The statement for U-FSSP-A then follows from Proposition 8.5.3. \square

Moving on to other shortlisting rules, we can show that they are not immune to manipulation. We will prove that this is the case when paired with either *unanimous* or *determined* allocation rules.

We first, prove this for the k -equal-representation shortlisting rule.

Proposition 8.5.10. *For all $k \in \mathbb{N}_{>0}$, the pair $\langle \text{REPRSHORT}_k, F \rangle$ where F is a unanimous allocation rule, is neither U-FSSP-P nor U-FSSP-O, for any completion principle that satisfies top-necessity and top-sufficiency.*

Moreover, if F is determined, then the pair $\langle \text{REPRSHORT}_1, F \rangle$ is not U-FSSP-A for any completion principle that satisfies top-sufficiency.

PROOF. We first prove the claim for $k = 1$ and then explain how to generalise to any $k \in \mathbb{N}_{>0}$ (only for U-FSSP-P and U-FSSP-O). Let $I = \langle \mathbb{P}, c, b \rangle$ be a shortlisting instance with $\mathbb{P} = \{p_1, \dots, p_4\}$, $c(p_2) = 2$, $c(p) = 1$ for all $p \in \mathbb{P} \setminus \{p_2\}$, and $b = 2$. We consider three agents with the following preferences.

	Preferences over the Projects	Awareness set C_i	Ideal Set Based on C_i
Agent 1	$p_3 \triangleright p_4 \triangleright p_2 \triangleright p_1$	$\{p_1, p_2, p_3, p_4\}$	$\{p_3, p_4\}$
Agent 2	$p_1 \triangleright p_2 \triangleright p_3 \triangleright p_4$	$\{p_1, p_2\}$	$\{p_1, p_3\}$
Agent 3	$p_2 \triangleright p_1 \triangleright p_3 \triangleright p_4$	$\{p_2\}$	$\{p_2\}$

Consider the 1-equal-representation shortlisting rule. Under the truthful profile $\mathbf{P} = (\{p_3, p_4\}, \{p_1, p_2\}, \{p_2\})$, the shortlist would be $\mathcal{P} = \{p_2\}$.

Assume now that agent 1 submits $\{p_1, p_3\}$ instead of $\{p_3, p_4\}$, leading to the shortlisting profile $\mathbf{P}' = (\{p_3, p_4\})$. Then, because of the tie-breaking mechanism, the outcome of the first stage becomes $\mathcal{P}' = \{p_1, p_3\}$.

The ideal set of agent 1 across both scenarios is $top_1(\mathcal{P} \cup \mathcal{P}') = \{p_1, p_3\} = \mathcal{P}'$ which implies that $top_1(\mathcal{P} \cup \mathcal{P}') \cap \mathcal{P} = \emptyset$. Thus, we have the following two facts:

- ▶ $\pi \cap top_1(\mathcal{P} \cup \mathcal{P}') = \emptyset$ for every $\pi \in \text{FEAS}(\langle \mathcal{P}, c, b \rangle)$;
- ▶ $\pi' \cap top_1(\mathcal{P} \cup \mathcal{P}') \neq \emptyset$ for every $\pi' \in \text{FEAS}(\langle \mathcal{P}', c, b \rangle)$ such that $\pi' \neq \emptyset$.

So every budget allocation in $\text{FEAS}(\langle \mathcal{P}', c, b \rangle)$ that is non-empty will be strictly preferred by agent 1 to all of the ones in $\text{FEAS}(\langle \mathcal{P}, c, b \rangle)$ (since the completion principle is top-sufficient). Moreover, agent 1 is indifferent between the empty budget allocation and any one from $\text{FEAS}(\langle \mathcal{P}', c, b \rangle)$ (since the completion principle is top-necessary). Given that F is unanimous, any budget allocation from $\text{FEAS}(\langle \mathcal{P}', c, b \rangle)$ can be achieved (by the corresponding unanimous profile). This means that agent 1's manipulation is pessimistically successful.

In case F is determined, for any profile \mathbf{A} for the allocation stage, we have that $F(\langle \mathcal{P}, c, b \rangle, \mathbf{A}) \cap \text{top}_1(\mathcal{P} \cup \mathcal{P}') \neq \emptyset$. This proves that agent 1's manipulation would also be a successful anticipative manipulation, given a completion principle that is top-sufficient.

We can generalise this to all $k > 1$ for the case of U-FSSP-P. To do so, add $3(k - 1)$ agents in groups of 3. Each group can conceive and approve of two new projects. It is easy to check that all the new projects will always be shortlisted, so we are back to the scenario above.


Note that we only talked about U-FSSP-P and U-FSSP-A. The result for U-FSSP-O is immediately derived from Proposition 8.5.3. \square

We finally consider the case of the median-based shortlisting rules.

Proposition 8.5.11. *Let δ be the Euclidean distance over \mathbb{R}^2 , for all $k \in \mathbb{N}_{>0}$, the pair $\langle \text{MEDIANSHORT}_{k,\delta}, F \rangle$, where F is a unanimous allocation rule, is neither U-FSSP-P nor U-FSSP-O, for any completion principle that satisfies top-necessity and top-sufficiency.*

Moreover, if F is determined, then the pair $\langle \text{MEDIANSHORT}_{1,\delta}, F \rangle$ is not U-FSSP-A, for any completion principle that satisfies top-sufficiency.

PROOF. We first prove the claim for $k = 1$ and then explain how to generalise it to all $k \in \mathbb{N}_{>0}$ (only for U-FSSP-P and U-FSSP-O). Consider the shortlisting instance $I = \langle \mathbb{P}, c, b \rangle$ with $\mathbb{P} = \{p_1, \dots, p_6\}$, all projects having cost 1, and a budget limit of $b = 3$. Suppose the distance δ is the usual distance in the plane, with the projects being positioned as in the figure below.



We assume that two agents are involved in the process. The first agent is such that $p_1 \triangleright_1 p_2 \triangleright_1 p_3 \triangleright_1 p_5 \triangleright_1 \dots$ and $C_1 = \{p_1, p_2, p_3, p_5\}$. The second agent is such

that $p_4 \triangleright_2 p_6 \triangleright_2 p_7 \triangleright_2 \dots$ and $C_2 = \{p_4, p_6, p_7\}$. The truthful shortlisting profile is then $\mathbf{P} = (\{p_1, p_2, p_3\}, \{p_4, p_6, p_7\})$. All projects have thus been submitted to the first stage, except for p_5 . The shortlisting rule $\text{MEDIANSHORT}_{1,\delta}$ would thus consider the following three clusters: $\{p_1, p_2, p_3, p_4\}$, $\{p_6\}$, and $\{p_7\}$, and the shortlist would be $\mathcal{P} = \{p_4, p_6, p_7\}$.

Now, assume that agent 1 submits $\{p_1, p_2, p_5\}$ instead of $\{p_1, p_2, p_3\}$, leading to the shortlisting profile $\mathbf{P}' = (\{p_1, p_2, p_5\}, \{p_4, p_6, p_7\})$. Then, the clusters will be $\{p_1\}$, $\{p_2, p_4\}$, and $\{p_5, p_6, p_7\}$ (for a suitable tie-breaking between Voronoi partitions). The shortlist would then be $\mathcal{P}' = \{p_1, p_2, p_5\}$.

Interestingly, we have $\text{top}_1(\mathcal{P} \cup \mathcal{P}') = \{p_1, p_2, p_5\} = \mathcal{P}'$. We have thus reached a similar construction as the one in the proof of Proposition 8.5.10. Every non-empty budget allocation in $\text{FEAS}(\langle \mathcal{P}', c, b \rangle)$ is strictly preferred by agent 1 to any budget allocation in $\text{FEAS}(\langle \mathcal{P}, c, b \rangle)$. The empty budget allocation is weakly preferred by agent 1 to any budget allocation in $\text{FEAS}(\langle \mathcal{P}, c, b \rangle)$. Given that F is unanimous, agent 1's manipulation is pessimistically successful. The same applies in the case of U-FSSP-A when F is determined.

We can extend this to all $k > 1$ for the case of U-FSSP-P. To do so, add $k - 1$ agents, all knowing and approving of three new projects of cost 1. All the new projects are placed uniformly on a circle with centre p_4 and a radius large enough so that all new projects will be in their own cluster. Then all the new projects will be shortlisted and we are back to the original case for $k = 1$.

Note that we only talked about U-FSSP-P and U-FSSP-P. The result for U-FSSP-O is immediately derived from Proposition 8.5.3. \square

Note that the previous statement can be generalised to many distances.

Interestingly both of the above statements are about unrestricted FSSP, though the successful manipulations presented in the proofs only make use of the awareness set of the manipulator. That is because in both proofs, the manipulator is initially aware of their top projects and manipulates by restraining from submitting some. This hints at some potentially interesting refinements of the FSSP axioms that can be worth studying. For instance, by restraining the type of ballots a manipulator can submit, or said differently, the type of strategic behaviours they can engage into.

8.6 Summary

The aim of this chapter was to capture more accurately the different stages that occur in real-life PB processes. Starting from the standard model of PB, we introduced a preliminary stage in which voters can submit proposals that will then be shortlisted to form an allocation instance, *i.e.*, an instance of the standard model. We explored two lines of work within this *end-to-end* model: one that relates to the creation of the shortlist, and another one that relates to the interactions between the two stages.

On the shortlisting side, we presented three different shortlisting rules, tailored to fulfil different objectives for the shortlist. The three main objectives we identified are the following: reducing the number of shortlisted projects, representing all the agents in the shortlist, and, avoiding having overly similar projects in the shortlist. The shortlisting rules REPRSHORT and MEDIANSHORT performs well on the first objective. They are also good candidates for the second and third objective, respectively. They also enjoy interesting axiomatic properties, as we have seen.

On the interaction side, we focused on the strategic behaviour that can emerge during the shortlisting stage, due to the two-stage nature of the process. Taking our time to reflect on what exactly it would mean for an agent—who has in mind the budget allocation determined in the second stage—to engage in strategic behaviour during the first stage, we introduced six notions of *first-stage strategy-proofness*. For each of the shortlisting rules we introduced, we then successively studied under which conditions they are immune to manipulation, or not. Our results are summarised in Table 8.6.1. The main take-home message here is that, overall, shortlisting rules are not immune to manipulation, and that there can always be an incentive for an agent not to propose to build a fountain in the centre of the main square.

	$\langle \text{NOMSHORT}, F \rangle$	$\langle \text{REPRSHORT}_k, F \rangle$	$\langle \text{MEDIANSHORT}_{k,\delta}, F \rangle$
R-FSSP-P, R-FSSP-A, and R-FSSP-O	✗	✗ ($k \geq 2$)	✗ ($k \geq 2$)
Result Statement	Theorem 8.5.5	Theorem 8.5.5	Theorem 8.5.5
Condition on F	Determined	Determined	Determined
Condition on the Completion Principle	Top-First	Top-First	Top-First
U-FSSP-P	✓	✗ ($k \geq 1$)	✗ ($k \geq 1$)
Result Statement	Proposition 8.5.8	Proposition 8.5.10	Proposition 8.5.11
Condition on F	Unanimous	Unanimous or Determined	Unanimous or Determined
Condition on the Completion Principle	Top-First	Top-Necessity and Top-Sufficiency	Top-Necessity and Top-Sufficiency
U-FSSP-O	✗	✗ ($k \geq 1$)	✗ ($k \geq 1$)
Result Statement	Proposition 8.5.9	Proposition 8.5.10	Proposition 8.5.11
Condition on F	Only for Specific Rules	Unanimous or Determined	Unanimous or Determined
Condition on the Completion Principle	Cost-Neutral Monotonicity	Top-Necessity and Top-Sufficiency	Top-Necessity and Top-Sufficiency
U-FSSP-A	✗	✗ ($k = 1$)	✗ ($k = 1$)
Result Statement	Proposition 8.5.9	Proposition 8.5.10	Proposition 8.5.11
Condition on F	Only for Specific Rules	Determined	Determined
Condition on the Completion Principle	Cost-Neutral Monotonicity	Top-Sufficiency	Top-Sufficiency

Table 8.6.1: Summary of the results on first-stage strategy-proofness for different shortlisting rules. For each FSSP axiom (or set of axioms), we indicate whether it is satisfied or not in general, and under which conditions on the allocation rule F and the completion principle the results applies. F is an arbitrary PB rule, and δ is the Euclidean distance over \mathbb{R}^2 .

Final Words

Chapter 9

Conclusion

The goal of this thesis was to address the problem of the aggregation of preferences in the context of participatory budgeting. We specifically emphasised the need for a formal approach that accounts for the wide variety of PB processes that are implemented throughout the world. We adopted a multi-faceted approach to achieve this, and saw a multitude of ways to assess the merits of aggregation rules for PB.

We will now present a recollection of all we discussed throughout the thesis (Section 9.1) and will then present a more general concluding discussion (Section 9.2).

9.1 Closing out the Thesis

After introducing the thesis, we directly dived into the formal analysis of PB. The first step was to set the scene. We did so by introducing the standard formal model of PB in Chapter 2. We then provided a long and comprehensive survey of the literature on PB in Chapter 3. This allowed us to precisely position the work presented in this thesis with regard to the field of research as a whole. Thus prepared, we jumped into the technical contributions of the thesis.

We started with Part **Two** where we presented investigations of the standard model of PB using non-standard methods.

In Chapter 4 we presented an analysis of fairness in PB in terms of equality of resources, a new approach radically different from the usual satisfaction-based fairness theory that has been developed in the literature. Doing so allowed us to circumvent the drawbacks of satisfaction-based fairness—notably the difficulty of accessing the satisfaction of a voter—while still presenting a viable approach to fairness in PB. The results indeed suggest that several of the measures of fairness we introduced are of interest and deserve further investigation. On top of that, we were able to find rules—

such as $\text{MES}[\text{sat}^{\text{cost}}]$ —that seem to perform well both in terms of equality of resources and in terms of satisfaction.

We continued our analysis of the standard model of PB in Chapter 5 where we adopted an epistemic view of the aggregation problem of PB. There, we reviewed the epistemic merits of different PB rules in terms of whether they could be seen as maximum likelihood estimators or not. Our analysis demonstrated that the most prominent PB rules are not good epistemic aggregators in that sense. At the same time, we were able to provide two rules that are maximum likelihood estimators, thus presenting a more positive picture.

Then came Part Three where we investigated different variations of the standard model of PB.

We started with Chapter 6 where we looked into a model of multi-constraint PB. The problem we faced there was to find PB rules that could easily accommodate additional constraints, *i.e.*, that would be robust to small variations of the model. To do so, we proposed to embed PB into judgment aggregation (JA) and to use JA rules to handle the aggregation. The main obstacle when using this approach is the potential blow-up in computational complexity. Focusing first on that point, we presented several specific ways of translating PB instances into JA instances that enabled us to use this approach efficiently. Specifically, they ensured that we could efficiently use JA rules for the aggregation problem of PB in several cases: PB with multiple resources, PB with dependencies between the projects, and PB with types and quotas. We then turned to the assessment of the qualities of JA rules when used for PB purposes. We investigated their merits with respect to the typical monotonicity axioms that have been introduced in the literature on PB. We showed that JA rules fared similarly to the standard PB rules with respect to these axioms.

In Chapter 7 we adopted a long-term perspective for our study of PB. We investigated how to define and enforce fairness when considering not only one-shot instances, but sequences of instances. Focusing on types of voters, we looked for combinations of evaluation functions and fairness criteria that would be satisfiable. We showed that perfect fairness cannot always be achieved, but that one could either optimise for fairness—which comes at a high computational cost—or converge towards fairness in the long run—which can only be guaranteed for limited numbers of agents or types.

Our last technical chapter, Chapter 8, concluded Part Three. In this chapter we studied an end-to-end model for PB, that is, a model accounting for both the stage during which voters submit proposals—the shortlisting stage—and the stage during which they vote on the shortlisted proposals—the allocation stage. Studying this model in a logical order, we started with the shortlisting stage. We introduced three shortlisting rules, designed to implement different objectives that we identified for this stage: reducing the number of proposals, giving everyone a chance to voice their wishes, and avoiding shortlisting similar proposals. We then explored the model and its two stages as a whole. We zoomed in on the problem of the incentives of the

agents, with a special point of focus on the interactions between the two stages. Our main findings indicate that it is hard to prevent agents from behaving strategically during the shortlisting stage.

Getting back to the main goal of the thesis, we wondered how to solve the aggregation problems in PB scenarios. Throughout the thesis, we have seen many ways of doing so, though we were not able to single out a specific method that fared well against all the benchmarks we studied. The main added value of this thesis is that it lays the ground for a multi-faceted analysis of PB rules through the numerous new takes on PB rules we developed. Many new perspectives for the formal study of PB have thus been opened up.

9.2 Opening up New Perspectives

As for much academic work, this thesis opens up many more research questions than it answers. I am not interested in providing a list of open questions, or future directions directly derived from the content of the thesis here. I trust the interested reader to be able to devise these themselves. Overall, each chapter presented an approach to PB that had never been investigated before; there are thus plentiful interesting directions to deepen this work. Instead, I want to use this last section of the thesis to discuss broader themes that I believe deserve to be explored in more depth.

Investigating ballot formats. Throughout this thesis we have only focused on the case of PB with approval ballots. These ballots unfortunately suffer an important drawback: their limited expressivity for a framework in which alternatives have different cost. As we have seen in the part of the survey dedicated to ballot formats (Section 3.1), other kinds of ballots have been considered in the literature, ballots that do not suffer this drawback. Still, these other ballot formats are either under-studied, or suffer other drawbacks in terms of ease of use. There is thus a need for devising new ballot formats achieving a balance between expressivity and ease of use (as done by, e.g., [Baumeister, Boes, Laußmann and Rey, 2023](#)), that should be complemented by a more systematic analysis of the impact of the ballot format on the outcome (in the line of [Fairstein, Benadè and Gal, 2023](#)).

Extending the literature on fairness. As we detailed when discussing new ways of approaching fairness for PB in Chapter 4, the question of fairness in PB is far from being settled. There are many open problems to be solved there (status of the core for approval ballots, natural rule satisfying FJR, *etc.*). A question that is still under-studied and close to my heart is that of fairness in extended PB settings. For instance, our analysis on multi-constraint PB models developed in Chapter 6 did not touch on the question of fairness at all. Developing a fairness theory for such models, and in general for complex voting domains is a fascinating research agenda.

Deepening the axiomatic analysis. The corpus of axioms that have been introduced in the literature on PB is still rather slim. There is a crucial need to develop that side of the literature to have other means to compare rules than fairness guarantees. There is still to date no consistent study of standard PB rules in terms of the monotonicity axioms we introduced in Section 3.4.2 and studied in Chapter 6 for instance. Investigating how to adapt the characterisation results from the multi-winner voting literature (Skowron, Faliszewski and Slinko, 2019; Lackner and Skowron, 2021) is another interesting direction.

Looking beyond the voting stage. Chapter 8 is to date the only formal analysis of a multi-stage PB model. Even though I said I would not discuss future directions directly linked to the thesis, I believe that this one calls for an exception as real-world implementations of PB processes all rely on several stages. Many other interesting questions that arise from a more holistic view of the PB process can be answered with social choice methods. This includes, for example, which incentives voters have when planning a project that they want to propose, *i.e.*, whether it is beneficial to make a project as cheap as possible or to merge two similar projects.

Stepping outside the Western world. The (computational) social choice literature so far has almost exclusively used PB processes in the Western world as examples. However, these processes do not represent the diversity of actual implementations of PB processes around the world. For instance, generally only a small percentage of a municipalities budget is allocated to PB in Western countries, and most projects funded through PB processes are small “quality of life” improvements that are not essential to the functioning of the city. In contrast, for example in early implementations of PB in Brazil, significant parts of a city’s budget was spent through PB processes, and many projects funded through PB processes addressed crucial parts of life, such as access to basic health care (Cabannes, 2004; Wampler, McNulty and Touchton, 2021). The objectives of the PB processes are thus radically different and call for different ways of determining the winning budget allocation. The weight given to fairness (in the sense of proportionality requirements) may be lowered for instance, or more deliberative processes can be considered. There is definitely a need for a deeper analysis here.

Bibliography

- Rebecca Abers. 2000. *Inventing Local Democracy: Grassroots Politics in Brazil*. Lynne Rienner Publishers. (Cited on page 5)
- 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 page 74)
- Marianne Akian, Ravindra Bapat, and Stéphane Gaubert. 2006. Max-plus algebra. In *Handbook of Linear Algebra*, Leslie Hogben (Ed.). Chapman and Hall, Chapter 35. (Cited on page 147)
- Giovanni Allegretti and Sofia Antunes. 2014. The Lisbon Participatory Budget: Results and Perspectives on an Experience in Slow but Continuous Transformation. *Field Actions Science Reports* (2014). Special Issue 11. (Cited on pages 140 and 218)
- Sanjeev Arora and Boaz Barak. 2009. *Computational Complexity: A Modern Approach*. Cambridge University Press. (Cited on pages 8, 22, and 159)
- Kenneth J. Arrow. 1951. *Social Choice and Individual Values*. Cowles Foundation. (Cited on pages 4 and 8)
- Kenneth J. Arrow, Amartya Sen, and Kotaro Suzumura (Eds.). 2002. *Handbook of Social Choice and Welfare*. Vol. 1. North-Holland. (Cited on pages 4 and 8)
- Kenneth J. Arrow, Amartya Sen, and Kotaro Suzumura (Eds.). 2011. *Handbook of Social Choice and Welfare*. Vol. 2. North-Holland. (Cited on pages 4 and 8)
- Haris Aziz, Anna Bogomolnaia, and Hervé Moulin. 2019. Fair Mixing: the Case of Dichotomous Preferences. In *Proceedings of the 20th ACM Conference on Economics and Computation (ACM-EC)*. 753–781. (Cited on page 74)

- Haris Aziz, Markus Brill, Vincent Conitzer, Edith Elkind, Rupert Freeman, and Toby Walsh. 2017. Justified Representation in Approval-Based Committee Voting. *Social Choice and Welfare* 48, 2 (2017), 461–485. (Cited on pages 36, 37, 38, 39, 93, and 94)
- Haris Aziz, Edith Elkind, Shenwei Huang, Martin Lackner, Luis Sanchez-Fernandez, and Piotr Skowron. 2018. On the Complexity of Extended and Proportional Justified Representation. In *Proceedings of the 32rd AAAI Conference on Artificial Intelligence (AAAI)*. 902–909. (Cited on pages 37 and 40)
- Haris Aziz and Aditya Ganguly. 2021. Participatory Funding Coordination: Model, Axioms and Rules. In *Proceedings of the 7th International Conference on Algorithmic Decision Theory (ADT)*. 409–423. (Cited on pages 25, 55, 68, and 73)
- 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 page 225)
- Haris Aziz, Sujit Gujar, Manisha Padala, Mashbat Suzuki, and Jeremy Vollen. 2022. Coordinating Monetary Contributions in Participatory Budgeting. *arXiv preprint arXiv:2206.05966* (2022). (Cited on pages 25, 68, and 73)
- Haris Aziz and Barton E. Lee. 2021. Proportionally Representative Participatory Budgeting with Ordinal Preferences. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*. 5110–5118. (Cited on pages 25, 28, 32, 36, 51, 52, 54, 57, 58, and 73)
- Haris Aziz, Barton E. Lee, and Nimrod Talmon. 2018. Proportionally Representative Participatory Budgeting: Axioms and Algorithms. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 23–31. (Cited on pages 25, 34, 35, 39, 44, 45, 46, 53, 57, 58, 59, 66, 87, and 130)
- John Bartholdi, Craig A. Tovey, and Michael A. Trick. 1989. Voting Schemes for which it can be Difficult to Tell who Won the Election. *Social Choice and Welfare* 6 (1989), 157–165. (Cited on pages 8 and 9)
- Dorothea Baumeister, Linus Boes, and Johanna Hillebrand. 2021. Complexity of Manipulative Interference in Participatory Budgeting. In *Proceedings of the 7th International Conference on Algorithmic Decision Theory (ADT)*. 424–439. (Cited on pages 25 and 68)
- Dorothea Baumeister, Linus Boes, and Christian Laußmann. 2022. Time-Constrained Participatory Budgeting Under Uncertain Project Costs. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*. 74–80. (Cited on pages 25, 55, and 72)

- Dorothea Baumeister, Linus Boes, Christian Laußmann, and Simon Rey. 2023. Bounded Approval Ballots: Balancing Expressiveness and Simplicity for Multiwinner Elections. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. (Cited on page 247)
- Dorothea Baumeister, Linus Boes, and Tessa Seeger. 2020. Irresolute Approval-based Budgeting. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 1774–1776. (Cited on pages 25, 33, and 62)
- Rachel Ben-Eliyahu and Rina Dechter. 1996. On Computing Minimal Models. *Annals of Mathematics and Artificial Intelligence* 18, 1 (1996), 3–27. (Cited on page 153)
- Gerdus Benadè, Swaprava Nath, Ariel D. Procaccia, and Nisarg Shah. 2021. Preference Elicitation for Participatory Budgeting. *Management Science* 67, 5 (2021), 2813–2827. (Cited on pages 25, 28, and 29)
- Jonathan Bendor, Daniel Diermeier, David A Siegel, and Michael Ting. 2011. *A Behavioral Theory of Elections*. Princeton University Press. (Cited on page 26)
- Duncan Black. 1948. On the Rationale of Group Decision-Making. *Journal of Political Economy* 56, 1 (1948), 23–34. (Cited on page 8)
- Charles Blackorby and David Donaldson. 1978. Measures of Relative Equality and their Meaning in Terms of Social Welfare. *Journal of Economic Theory* 18, 1 (1978), 59–80. (Cited on page 185)
- Hans L. Bodlaender. 1998. A Partial k -Arboretum of Graphs with Bounded Treewidth. *Theoretical Computer Science* 209, 1-2 (1998), 1–45. (Cited on page 154)
- Hans L. Bodlaender and Ton Kloks. 1996. Efficient and Constructive Algorithms for the Pathwidth and Treewidth of Graphs. *Journal of Algorithms* 21, 2 (1996), 358–402. (Cited on page 154)
- Linus Boes, Rachael Colley, Umberto Grandi, Jerome Lang, and Arianna Novaro. 2021. Collective Discrete Optimisation as Judgment Aggregation. *arXiv preprint arXiv:2112.00574* (2021). (Cited on page 74)
- Anna Bogomolnaia, Hervé Moulin, and Richard Stong. 2005. Collective Choice Under Dichotomous Preferences. *Journal of Economic Theory* 122, 2 (2005), 165–184. (Cited on page 74)
- Jean-Charles de Borda. 1781. *Mémoire sur les Élections au Scrutin*. Histoire de l’Académie Royale des Sciences. (Cited on page 8)
- Sylvain Bouveret. 2007. *Allocation et Partage Équitables de Ressources Indivisibles : Modélisation, Complexité et Algorithmique*. Ph.D. Dissertation. Supaéro and University of Toulouse. (Cited on page 185)

- Luc Bovens and Wlodek Rabinowicz. 2006. Democratic Answers to Complex Questions—An Epistemic Perspective. *Synthese* 150, 1 (2006), 131–153. (Cited on page 121)
- Irem Bozbay, Franz Dietrich, and Hans Peters. 2014. Judgment Aggregation in Search for the Truth. *Games and Economic Behavior* 87 (2014), 571–590. (Cited on page 121)
- Florian Brandl, Felix Brandt, Dominik Peters, and Christian Stricker. 2021. Distribution Rules Under Dichotomous Preferences: Two Out of Three Ain’t Bad. In *Proceedings of the 22nd ACM Conference on Economics and Computation (ACM-EC)*. 158–179. (Cited on page 74)
- Felix Brandt. 2018. Collective Choice Lotteries: Dealing with Randomization in Economic Design. In *The Future of Economic Design*, Jean-François Laslier, Hervé Moulin, Remzi Sanver, and William S. Zwicker (Eds.). Springer-Verlag. (Cited on page 74)
- Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (Eds.). 2016a. *Handbook of Computational Social Choice*. Cambridge University Press. (Cited on pages 8 and 74)
- Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia. 2016b. Introduction to Computational Social Choice. In *Handbook of Computational Social Choice*, Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (Eds.). Cambridge University Press, Chapter 1. (Cited on page 74)
- Robert Brederick, Piotr Faliszewski, Andrzej Kaczmarczyk, and Rolf Niedermeier. 2019. An Experimental View on Committees Providing Justified Representation. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*. 109–115. (Cited on page 37)
- 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)*. (Cited on pages 17, 21, 25, 31, 32, 34, 42, 43, 44, 45, 46, 50, 51, 57, 58, and 78)
- Markus Brill, Rupert Freeman, Svante Janson, and Martin Lackner. 2017. Phragmén’s Voting Methods and Justified Representation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*. 406–413. (Cited on page 20)
- Eric Budish. 2011. The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes. *Journal of Political Economy* 119, 6 (2011), 1061–1103. (Cited on page 83)
- Yves Cabannes. 2004. Participatory Budgeting: A Significant Contribution to Participatory Democracy. *Environment and Urbanization* 16, 1 (2004), 27–46. (Cited on pages 6, 7, and 248)

- Marco Cadoli, Francesco M. Donini, Paolo Liberatore, and Marco Schaerf. 2002. Pre-processing of Intractable Problems. *Information and Computation* 176, 2 (2002), 89–120. (Cited on page 157)
- Ioannis Caragiannis, George Christodoulou, and Nicos Protopapas. 2022. Truthful Aggregation of Budget Proposals with Proportionality Guarantees. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*. 4917–4924. (Cited on page 74)
- Ioannis Caragiannis, Ariel D. Procaccia, and Nisarg Shah. 2013. When do Noisy Votes Reveal the Truth? In *Proceedings of the 14th ACM Conference on Electronic Commerce (ACM-EC)*. 143–160. (Cited on page 121)
- Ioannis Caragiannis, Ariel D. Procaccia, and Nisarg Shah. 2014. Modal Ranking: A Uniquely Robust Voting Rule. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*. 616–622. (Cited on page 121)
- Federica Ceron, Stéphane Gonzalez, and Adriana Navarro-Ramos. 2022. Axiomatic Characterizations of the Knapsack and Greedy Participatory Budgeting Methods. *Working paper* (2022). (Cited on page 65)
- Alfonso Cevallos and Alistair Stewart. 2021. A Verifiably Secure and Proportional Committee Election Rule. In *Proceedings of the 3rd ACM Conference on Advances in Financial Technologies (AFT)*. 29–42. (Cited on page 35)
- Jiehua Chen, Martin Lackner, and Jan Maly. 2022. Participatory Budgeting with Donations and Diversity Constraints. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*. 9323–9330. (Cited on pages 25, 68, 71, and 73)
- Yu Cheng, Zhihao Jiang, Kamesh Munagala, and Kangning Wang. 2020. Group Fairness in Committee Selection. *ACM Transactions on Economics and Computation (TEAC)* 8, 4 (2020), 1–18. (Cited on page 47)
- 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 page 141)
- City of Amsterdam. 2022. Oost Begroot. <https://www.amsterdam.nl/stadsdelen/oost/oost-begroot/>. Last accessed on 16 May 2023. (Cited on pages 7, 70, and 140)
- Vincet Cohen-Addad, Anupam Gupta, Amit Kumar, Euiwoong Lee, and Jason Li. 2019. Tight FPT Approximations for k -Median and k -Means. In *Proceedings of the 46th International Colloquium on Automata, Languages, and Programming (ICALP)*. 42:1–42:14. (Cited on page 226)

- Marie-Jean-Antoine-Nicolas de Caritat, Marquis de Condorcet. 1785. *Essai sur l'Application de l'Analyse à la Probabilité des Décisions Rendues à la Pluralité des Voix*. Facsimile reprint of original published in Paris, 1972, by the Imprimerie Royale. (Cited on pages 8, 119, and 121)
- Vincent Conitzer, Rupert Freeman, and Nisarg Shah. 2017. Fair Public Decision Making. In *Proceedings of the 18th ACM Conference on Economics and Computation (ACM-EC)*. 629–646. (Cited on page 74)
- Vincent Conitzer and Tuomas Sandholm. 2005. Common Voting Rules as Maximum Likelihood Estimators. In *Proceedings of the 21st Annual Conference on Uncertainty in Artificial Intelligence (UAI)*. 145–152. (Cited on pages 121, 122, 123, and 132)
- Vincent Conitzer and Toby Walsh. 2016. Barriers to Manipulation in Voting. In *Handbook of Computational Social Choice*, Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (Eds.). Cambridge University Press, New York, Chapter 6. (Cited on page 9)
- Dirk van Dalen. 2013. *Logic and Structure* (5th ed.). Springer-Verlag. (Cited on page 143)
- Andreas Darmann, Christian Klamler, and Ulrich Pferschy. 2009. Maximizing the Minimum Voter Satisfaction on Spanning Trees. *Mathematical Social Sciences* 58, 2 (2009), 238–250. (Cited on page 74)
- Andreas Darmann, Christian Klamler, and Ulrich Pferschy. 2011. Finding Socially Best Spanning Trees. *Theory and Decision* 70, 4 (2011), 511–527. (Cited on page 74)
- Adnan Darwiche and Pierre Marquis. 2002. A Knowledge Compilation Map. *Journal of Artificial Intelligence Research* 17 (2002), 229–264. (Cited on pages 146 and 158)
- Amrita Dhillon and Susana Peralta. 2002. Economic Theories of Voter Turnout. *The Economic Journal* 112, 480 (2002), F332–F352. (Cited on page 26)
- Nelson Dias (Ed.). 2018. *Hope for democracy: 30 years of participatory budgeting*. Epopéia Records. (Cited on page 5)
- Nelson Dias, Sahsil Enríquez, and Simone Júlio (Eds.). 2019. *The Participatory Budgeting World Atlas*. Epopéia Records. (Cited on page 5)
- Franz Dietrich. 2014. Scoring Rules for Judgment Aggregation. *Social Choice and Welfare* 42, 4 (2014), 873–911. (Cited on page 143)
- Franz Dietrich and Christian List. 2007a. Arrow's Theorem in Judgment Aggregation. *Social Choice and Welfare* 29, 1 (2007), 19–33. (Cited on page 141)

- Franz Dietrich and Christian List. 2007b. Strategy-Proof Judgment Aggregation. *Economics & Philosophy* 23, 3 (2007), 269–300. (Cited on page 24)
- Franz Dietrich and Kai Spiekermann. 2019. Jury Theorems. In *The Routledge Handbook of Social Epistemology*, Miranda Fricker, Peter J. Graham, David Henderson, and Nikolaj J.L.L. Pedersen (Eds.). Routledge. (Cited on page 119)
- Rodney G. Downey and Michael R. Fellows. 2013. *Fundamentals of Parameterized Complexity*. Springer. (Cited on pages 22 and 154)
- Conal Duddy, Ashley Piggins, and William S. Zwicker. 2016. Aggregation of Binary Evaluations: a Borda-Like Approach. *Social Choice and Welfare* 46 (2016), 301–333. (Cited on page 220)
- Michael Dummett. 1984. *Voting Procedures*. Oxford University Press. (Cited on page 51)
- Ronald Dworkin. 1981a. What is Equality? Part 1: Equality of Welfare. *Philosophy & Public Affairs* 10, 3 (1981), 185–246. (Cited on page 78)
- Ronald Dworkin. 1981b. What is Equality? Part 2: Equality of Resources. *Philosophy & Public Affairs* 10, 4 (1981), 283–345. (Cited on page 79)
- Edith Elkind, Piotr Faliszewski, Ayumi Igarashi, Pasin Manurangsi, Ulrike Schmidt-Kraepelin, and Warut Suksompong. 2022. The Price of Justified Representation. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*. 4983–4990. (Cited on page 110)
- Edith Elkind, Martin Lackner, and Dominik Peters. 2022. Preference Restrictions in Computational Social Choice: A Survey. *arXiv preprint arXiv:2205.09092* (2022). (Cited on page 9)
- Edith Elkind and Arkadii Slinko. 2016. Rationalizations of Voting Rules. In *Handbook of Computational Social Choice*, Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (Eds.). Cambridge University Press, Chapter 8. (Cited on pages 65 and 119)
- 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, New York, Chapter 17, 399–426. (Cited on page 144)
- Ulle Endriss (Ed.). 2017. *Trends in Computational Social Choice*. AI Access. (Cited on page 8)

- 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 page 141)
- Ulle Endriss, Umberto Grandi, Ronald de Haan, and Jérôme Lang. 2016. Succinctness of Languages for Judgment Aggregation. In *Proceedings of the 15th International Conference on Principles of Knowledge Representation and Reasoning (KR)*. 176–185. (Cited on page 142)
- 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 (2020), 687–731. (Cited on page 146)
- Ulle Endriss and Jérôme Lang (Eds.). 2006. *Proceedings of the 1st International Workshop on Computational Social Choice (COMSOC)*. ILLC, University of Amsterdam. (Cited on page 8)
- Eithan Ephrati and Jeffrey S. Rosenschein. 1993. Multi-Agent Planning as a Dynamic Search for Social Consensus. In *Proceedings of the 13th International Joint Conference on Artificial Intelligence (IJCAI)*. 423–429. (Cited on page 8)
- Bruno Escoffier, Jérôme Lang, and Meltem Öztürk. 2008. Single-Peaked Consistency and its Complexity. In *Proceedings of the 18th European Conference on Artificial Intelligence (ECAI)*. 366–370. (Cited on page 9)
- Patricia Everaere, Sébastien Konieczny, and Pierre Marquis. 2014. Counting Votes for Aggregating Judgments. In *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 1177–1184. (Cited on page 144)
- Brandon Fain, Ashish Goel, and Kamesh Munagala. 2016. The Core of the Participatory Budgeting Problem. In *Proceedings of the 12th International Workshop on Internet and Network Economics (WINE)*. 384–399. (Cited on pages 46, 47, and 74)
- Brandon Fain, Kamesh Munagala, and Nisarg Shah. 2018. Fair Allocation of Indivisible Public Goods. In *Proceedings of the 19th ACM Conference on Economics and Computation (ACM-EC)*. 575–592. (Cited on pages 47 and 74)
- Roy Fairstein, Gerdus Benadè, and Kobi Gal. 2023. Participatory Budgeting Design for the Real World. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*. (Cited on pages 25, 30, and 247)
- Roy Fairstein, Dan Vilenchik, Reshef Meir, and Kobi Gal. 2022. Welfare vs. Representation in Participatory Budgeting. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 409–417. (Cited on page 43)

- Piotr Faliszewski, Edith Hemaspaandra, and Lane A Hemaspaandra. 2009. How Hard is Bribery in Elections? *Journal of Artificial Intelligence Research* 35 (2009), 485–532. (Cited on page 9)
- Piotr Faliszewski and Jörg Rothe. 2016. Control and Bribery in Voting. In *Handbook of Computational Social Choice*, Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (Eds.). Cambridge University Press, New York, Chapter 7. (Cited on page 9)
- Piotr Faliszewski, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. 2017. Multiwinner Voting: A New Challenge for Social Choice Theory. In *Trends in Computational Social Choice*, Ulle Endriss (Ed.). Chapter 2. (Cited on pages 74 and 220)
- 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 page 221)
- Till Fluschnik, Piotr Skowron, Mervin Triphaus, and Kai Wilker. 2019. Fair Knapsack. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*. 1941–1948. (Cited on pages 25, 33, 67, 68, 69, and 74)
- Rupert Freeman, David M. Pennock, Dominik Peters, and Jennifer Wortman Vaughan. 2021. Truthful Aggregation of Budget Proposals. *Journal of Economic Theory* 193 (2021), 105234. (Cited on page 74)
- Rupert Freeman, Seyed Majid Zahedi, and Vincent Conitzer. 2017. Fair and Efficient Social Choice in Dynamic Settings. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*. 4580–4587. (Cited on page 177)
- Rupert Freeman, Seyed Majid Zahedi, Vincent Conitzer, and Benjamin C. Lee. 2018. Dynamic Proportional Sharing: A Game-Theoretic Approach. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 2, 1 (2018), 3:1–3:36. (Cited on page 177)
- Martin Fürer and Huiwen Yu. 2011. Packing-Based Approximation Algorithm for the k -Set Cover Problem. In *Proceedings of the 22nd International Symposium on Algorithms and Computation (ISAAC)*. 484–493. (Cited on pages 82 and 85)
- Michael R. Garey and David S. Johnson. 1979. *Computers and Intractability*. W.H. Freeman. (Cited on pages 93, 161, 188, and 190)
- Allan Gibbard. 1973. Manipulation of Voting Schemes: a General Result. *Econometrica* (1973), 587–601. (Cited on pages 8 and 24)
- Corrado Gini. 1912. *Variabilità e Mutabilità: Contributo allo Studio delle Distribuzioni e Delle Relazioni Statistiche*. P. Cuppini. (Cited on page 185)

- Ashish Goel, Anilesh K. Krishnaswamy, Sukolsak Sakshuwong, and Tanja Aitamurto. 2019. Knapsack Voting for Participatory Budgeting. *ACM Transactions on Economics and Computation* 7, 2 (2019), 8:1–8:27. (Cited on pages 16, 25, 28, 31, 32, 64, 65, 74, and 121)
- Thomas klein Goldewijk. 2022. *Fairness in Perpetual Participatory Budgeting*. MSc Thesis. Universiteit van Amsterdam, ILLC. (Cited on pages 189, 207, and 214)
- Ronald L. Graham, Donald E. Knuth, and Oren Patashnik. 1994. *Concrete Mathematics: a Foundation for Computer Science* (2nd ed.). Addison–Wesley. (Cited on page 22)
- Umberto Grandi and Ulle Endriss. 2011. Binary Aggregation with Integrity Constraints. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*. 204–209. (Cited on page 142)
- 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 141, 146, and 148)
- Sven Ove Hansson. 2001. *The Structure of Values and Norms*. Cambridge University Press. (Cited on page 26)
- Edith Hemaspaandra, Lane A. Hemaspaandra, and Jörg Rothe. 1997. Exact Analysis of Dodgson Elections: Lewis Carroll’s 1876 Voting System is Complete for Parallel Access to NP. *Journal of the ACM* 44, 6 (1997), 806–825. (Cited on page 9)
- Edith Hemaspaandra, Holger Spakowski, and Jörg Vogel. 2005. The Complexity of Kemeny Elections. *Theoretical Computer Science* 349, 3 (2005), 382–391. (Cited on page 9)
- D. Ellis Hershkowitz, Anson Kahng, Dominik Peters, and Ariel D. Procaccia. 2021. District-Fair Participatory Budgeting. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*. 5464–5471. (Cited on pages 9, 25, 68, and 70)
- Olivier Hudry. 1989. *Recherche d’Ordres Médiants : Complexité, Algorithmique et Problèmes Combinatoires*. Ph.D. Dissertation. École Nationale Supérieure des Télécommunications, Paris. (Cited on page 8)
- Anil K. Jain and Richard C. Dubes. 1988. *Algorithms for Clustering Data*. Prentice-Hall. (Cited on page 221)
- Pallavi Jain, Krzysztof Sornat, and Nimrod Talmon. 2020. Participatory Budgeting with Project Interactions. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*. 386–392. (Cited on pages 25, 27, 68, 71, and 72)

- Pallavi Jain, Krzysztof Sornat, Nimrod Talmon, and Meirav Zehavi. 2021. Participatory Budgeting with Project Groups. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*. 276–282,. (Cited on pages 25, 68, 70, 71, and 159)
- Pallavi Jain, Nimrod Talmon, and Laurent Bulteau. 2021. Partition Aggregation for Participatory Budgeting. In *Proceedings of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 665–673. (Cited on page 72)
- Svante Janson. 2016. Phragmén’s and Thiele’s Election Methods. *arXiv preprint arXiv:1611.08826* (2016). (Cited on pages 20 and 225)
- Zhihao Jiang, Kamesh Munagala, and Kangning Wang. 2020. Approximately Stable Committee Selection. In *Proceedings of the 52nd Annual ACM Symposium on Theory of Computing (STOC)*. 463–472. (Cited on pages 25 and 48)
- Tapas Kanungo, David M. Mount, Nathan S. Netanyahu, Christine D. Piatko, Ruth Silverman, and Angela Y. Wu. 2004. A Local Search Approximation Algorithm for k -Means Clustering. *Computational Geometry* 28, 2-3 (2004), 89–112. (Cited on page 226)
- Oded Kariv and S. Louis Hakimi. 1979. An Algorithmic Approach to Network Location Problems. II: The p -Medians. *SIAM Journal on Applied Mathematics* 37, 3 (1979), 539–560. (Cited on page 226)
- Richard M. Karp. 1972. Reducibility among Combinatorial Problems. In *Complexity of Computer Computations*, Raymond E. Miller, James W. Thatcher, and Jean D. Bohlinger (Eds.). Springer-Verlag, 85–103. (Cited on pages 93, 151, 157, and 190)
- Richard M. Karp and Richard J. Lipton. 1980. Some Connections Between Nonuniform and Uniform Complexity Classes. In *Proceedings of the 12th Annual ACM Symposium on Theory of Computing (STOC)*. 302–309. (Cited on pages 157 and 159)
- Hans Kellerer, Ulrich Pferschy, and David Pisinger. 2004. *Knapsack Problems*. Springer-Verlag. (Cited on pages 19, 20, and 69)
- D. Marc Kilgour. 2016. Approval Elections with a Variable Number of Winners. *Theory and Decision* 81, 2 (2016), 199–211. (Cited on page 220)
- Angelika Kimmig, Guy van den Broeck, and Luc de Raedt. 2017. Algebraic Model Counting. *Journal of Applied Logic* 22 (2017), 46–62. (Cited on pages 147 and 148)
- Boas Kluiving, Adriaan de Vries, Pepijn Vrijbergen, Arthur Boixel, and Ulle Endriss. 2020. Analysing Irresolute Multiwinner Voting Rules with Approval Ballots via SAT Solving. In *Proceedings of the 24th European Conference on Artificial Intelligence (ECAI)*. 131–138. (Cited on page 64)

- Martin Lackner. 2020. Perpetual Voting: Fairness in Long-Term Decision Making. In *Proceedings of the 24th AAI Conference on Artificial Intelligence (AAAI)*. 2103–2110. (Cited on page 177)
- 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 page 221)
- Martin Lackner and Jan Maly. 2023. Proportional Decisions in Perpetual Voting. In *Proceedings of the 37th AAI Conference on Artificial Intelligence (AAAI)*. (Cited on page 177)
- 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 12)
- Martin Lackner and Piotr Skowron. 2020. Utilitarian Welfare and Representation Guarantees of Approval-Based Multiwinner Rules. *Artificial Intelligence* 288 (2020), 103366. (Cited on pages 43 and 110)
- Martin Lackner and Piotr Skowron. 2021. Consistent Approval-Based Multi-Winner Rules. *Journal of Economic Theory* 192 (2021), 105173. (Cited on page 248)
- Martin Lackner and Piotr Skowron. 2023. *Multi-Winner Voting with Approval Preferences*. Springer-Verlag. (Cited on pages 4, 15, 36, 37, 49, 61, 62, 74, 108, 220, and 224)
- Julian Lamont and Christi Favor. 2017. Distributive Justice. In *The Stanford Encyclopedia of Philosophy* (Winter 2017 ed.), Edward N. Zalta (Ed.). Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/win2017/entries/justice-distributive> (Cited on page 77)
- Jérôme Lang, Gabriella Pigozzi, Marija Slavkovik, and Leendert van der Torre. 2011. Judgment aggregation rules based on minimization. In *Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*. 238–246. (Cited on pages 144 and 168)
- Jérôme Lang, Gabriella Pigozzi, Marija Slavkovik, Leendert van der Torre, and Srdjan Vesic. 2017. A Partial Taxonomy of Judgment Aggregation Rules and their Properties. *Social Choice and Welfare* 48, 2 (2017), 327–356. (Cited on page 144)
- 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 141)

- 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. (Cited on page 223)
- Annick Laruelle. 2021. Voting to Select Projects in Participatory Budgeting. *European Journal of Operational Research* 288, 2 (2021), 598–604. (Cited on pages 25, 32, 33, 34, and 67)
- Jean-François Laslier, Hervé Moulin, Remzi Sanver, and William S. Zwicker (Eds.). 2018. *The Future of Economic Design*. Springer-Verlag. (Cited on page 8)
- Shira B. Lewin. 1996. Economics and Psychology: Lessons for Our Own Day From the Early Twentieth Century. *Journal of Economic Literature* 34, 3 (1996), 1293–1323. (Cited on page 26)
- Richard J. Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. 2004. On Approximately Fair Allocations of Indivisible Goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce (ACM-EC)*. 125–131. (Cited on pages 9 and 83)
- Christian List and Philip Pettit. 2002. Aggregating Sets of Judgments: An Impossibility Result. *Economics & Philosophy* 18, 1 (2002), 89–110. (Cited on pages 141 and 142)
- Maaïke 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 pages 20, 25, 36, 40, 42, 44, 45, 50, 54, 56, and 57)
- 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 25, 68, and 72)
- Jan Maly, Simon Rey, Ulle Endriss, and Martin Lackner. 2023. Fairness in Participatory Budgeting via Equality of Resources. In *Proceedings of the 22th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. (Cited on page 10)
- Pierre Marquis. 2015. Compile! In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*. 4112–4118. (Cited on page 157)
- Kenneth O. May. 1952. A Set of Independent Necessary and Sufficient Conditions for Simple Majority Decision. *Econometrica* (1952), 680–684. (Cited on page 8)
- Iain McLean and Arnold Urken (Eds.). 1995. *Classics of Social Choice*. University of Michigan Press. (Cited on page 8)

- Reshef Meir. 2018. *Strategic Voting*. Morgan & Claypool Publishers. (Cited on pages 24 and 220)
- 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 74)
- Michael K. Miller and Daniel Osherson. 2009. Methods for Distance-based Judgment Aggregation. *Social Choice and Welfare* 32, 4 (2009), 575–601. (Cited on page 144)
- Nima Motamed, Arie Soeteman, Simon Rey, and Ulle Endriss. 2022. Participatory Budgeting with Multiple Resources. In *Proceedings of the 19th European Conference on Multi-Agent Systems (EUMAS)*. 330–347. (Cited on pages 10, 25, 55, 64, 68, 71, and 72)
- Hervé Moulin. 1980. On Strategy-Proofness and Single Peakedness. *Public Choice* (1980), 437–455. (Cited on page 8)
- Kamesh Munagala, Yiheng Shen, and Kangning Wang. 2022. Auditing for Core Stability in Participatory Budgeting. In *Proceedings of the 18th International Workshop on Internet and Network Economics (WINE)*. 292–310. (Cited on pages 25 and 48)
- Kamesh Munagala, Yiheng Shen, Kangning Wang, and Zhiyi Wang. 2022. Approximate Core for Committee Selection via Multilinear Extension and Market Clearing. In *Proceedings of the SODA22 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 2229–2252. (Cited on pages 25, 47, and 48)
- Chris Murray. 2019. *Toronto’s Participatory Budgeting Pilot Evaluation*. City Manager’s Report. City of Toronto. Available at www.toronto.ca/legdocs/mmis/2019/bu/bgrd/backgroundfile-124370.pdf (Last accessed on 16 May 2023). (Cited on page 218)
- Klaus Nehring and Marcus Pivato. 2019. Majority Rule in the Absence of a Majority. *Journal of Economic Theory* (2019), 213–257. (Cited on pages 143 and 144)
- Fanny Pascual, Krzysztof Rzdca, and Piotr Skowron. 2018. Collective Schedules: Scheduling Meets Computational Social Choice. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 667–675. (Cited on page 74)
- Deval Patel, Arindam Khan, and Anand Louis. 2021. Group Fairness for Knapsack Problems. In *Proceedings of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 1001–1009. (Cited on pages 25, 68, and 71)
- Dominik Peters. 2018. Proportionality and Strategyproofness in Multiwinner Elections. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 1549–1557. (Cited on pages 24, 63, and 64)

- Domink Peters. 2019. *Fair Division of the Commons*. DPhil Thesis. University of Oxford. (Cited on page 63)
- Dominik Peters, Grzegorz Pierczyński, and Piotr Skowron. 2021. Proportional Participatory Budgeting with Additive Utilities. In *Proceedings of the 35th Annual Conference on Neural Information Processing Systems (NeurIPS)*. 12726–12737. (Cited on pages 21, 22, 25, 27, 36, 37, 38, 39, 40, 42, 47, 48, 49, 50, 52, 53, 56, 57, 58, 59, 66, and 95)
- Dominik Peters and Piotr Skowron. 2020. Proportionality and the Limits of Welfarism. In *Proceedings of the 21st ACM Conference on Economics and Computation (ACM-EC)*. 793–794. (Cited on pages 21, 37, 49, 50, and 96)
- Gabriella Pigozzi. 2006. Belief Merging and the Discursive Dilemma: An Argument-Based Account to Paradoxes of Judgment Aggregation. *Synthese* 152, 2 (2006), 285–298. (Cited on page 144)
- Marcus Pivato. 2019. Realizing Epistemic Democracy. In *The Future of Economic Design*, Jean-François Laslier, Hervé Moulin, M. Remzi Sanver, and William S. Zwicker (Eds.). Springer-Verlag, 103–112. (Cited on pages 65 and 119)
- Osmany Porto de Oliveira. 2017. *International Policy Diffusion and Participatory Budgeting: Ambassadors of Participation, International Institutions and Transnational Networks*. Springer-Verlag. (Cited on page 5)
- Ariel D. Procaccia, Sashank J. Reddi, and Nisarg Shah. 2012. A Maximum Likelihood Approach for Selecting Sets of Alternatives. In *Proceedings of the 28th Annual Conference on Uncertainty in Artificial Intelligence (UAI)*. 695–704. (Cited on pages 120 and 121)
- Ariel D. Procaccia and Jeffrey S. Rosenschein. 2006. The Distortion of Cardinal Preferences in Voting. In *Proceedings of the International Workshop on Cooperative Information Agents X (CIA)*. 317–331. (Cited on page 29)
- Simon Rey. 2022. Designing Mechanisms for Participatory Budgeting: Doctoral Consortium. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 1860–1862. (Cited on page 10)
- Simon Rey and Ulle Endriss. 2023. Epistemic Selection of Costly Alternatives: The Case of Participatory Budgeting. *arXiv preprint arXiv:2304.10940* (2023). (Cited on page 11)
- 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 11 and 159)

- Simon Rey, Ulle Endriss, and Ronald de Haan. 2021. Shortlisting Rules and Incentives in an End-to-End Model for Participatory Budgeting. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*. 370–376. (Cited on page 12)
- Simon Rey and Jan Maly. 2023. The (Computational) Social Choice Take on Indivisible Participatory Budgeting. *arXiv preprint arXiv:2303.00621* (2023). (Cited on page 10)
- Jack Robertson and William Webb. 1998. *Cake-Cutting Algorithms*. A. K. Peters. (Cited on page 81)
- Neil Robertson and Paul D Seymour. 1983. Graph Minors. I. Excluding a Forest. *Journal of Combinatorial Theory, Series B* 35, 1 (1983), 39–61. (Cited on page 154)
- Jörg Rothe (Ed.). 2015. *Economics and Computation*. Springer-Verlag. (Cited on page 74)
- Luis Sánchez-Fernández, Edith Elkind, Martin Lackner, Norberto Fernández, Jesús 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 40)
- Luis Sánchez-Fernández, Norberto Fernández-García, Jesús A Fisteus, and Markus Brill. 2022. The Maximin Support Method: An Extension of the D’Hondt Method to Approval-Based Multiwinner Elections. *Mathematical Programming* (2022). (Cited on page 34)
- Mark Allen Satterthwaite. 1975. Strategy-Proofness and Arrow’s Conditions: Existence and Correspondence Theorems for Voting Procedures and Social Welfare Functions. *Journal of Economic Theory* 10, 2 (1975), 187–217. (Cited on pages 8 and 24)
- Herbert E. Scarf. 1967. The Core of an N Person Game. *Econometrica* (1967), 50–69. (Cited on page 47)
- Thomas J. Schaefer. 1978. The Complexity of Satisfiability Problems. In *Proceedings of the 10th Annual ACM Symposium on Theory of Computing (STOC)*. 216–226. (Cited on page 188)
- Anwar Shah (Ed.). 2007. *Participatory Budgeting*. The World Bank. (Cited on pages 6, 7, and 218)
- Yves Sintomer, Carsten Herzberg, and Anja Röcke. 2008. Participatory Budgeting in Europe: Potentials and Challenges. *International Journal of Urban and Regional Research* 32, 1 (2008), 164–178. (Cited on page 5)

- Yves Sintomer, Carsten Herzberg, Anja Röcke, and Giovanni Allegretti. 2012. Transnational Models of Citizen Participation: The Case of Participatory Budgeting. *Journal of Deliberative Democracy* 8, 2 (2012). (Cited on page 175)
- Piotr Skowron, Piotr Faliszewski, and Arkadii Slinko. 2019. Axiomatic Characterization of Committee Scoring Rules. *Journal of Economic Theory* 180 (2019), 244–273. (Cited on page 248)
- Piotr Skowron, Arkadii Slinko, Stanisław Szufa, and Nimrod Talmon. 2020. Participatory Budgeting with Cumulative Votes. *arXiv preprint arXiv:2009.02690* (2020). (Cited on pages 25, 28, 36, and 55)
- Gogulapati Sreedurga, Mayank Ratan Bhardwaj, and Y. Narahari. 2022. Maxmin Participatory Budgeting. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*. 489–495. (Cited on pages 25, 33, 62, 65, 67, 68, and 69)
- Dariusz Stolicki, Stanisław Szufa, and Nimrod Talmon. 2020. Pabulib: A Participatory Budgeting Library. *arXiv preprint arXiv:2012.06539* (2020). (Cited on pages 24, 111, and 120)
- Nimrod Talmon and Piotr Faliszewski. 2019. A Framework for Approval-Based Budgeting Methods. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*. 2181–2188. (Cited on pages 17, 19, 25, 31, 33, 59, 60, 62, 68, and 69)
- Zoi Terzopoulou and Ulle Endriss. 2019. Optimal Truth-Tracking Rules for the Aggregation of Incomplete Judgments. In *Proceedings of the 12th International Symposium on Algorithmic Game Theory (SAGT)*. 298–311. (Cited on page 121)
- William Thomson. 2001. On the Axiomatic Method and its Recent Applications to Game Theory and Resource Allocation. *Social Choice and Welfare* 18, 2 (2001), 327–386. (Cited on page 8)
- Brian Wampler. 2000. A Guide to Participatory Budgeting. *Third conference of the International Budget Project* (2000). (Cited on page 6)
- Brian Wampler. 2012. Participatory Budgeting: Core Principles and Key Impacts. *Journal of Public Deliberation* (2012). (Cited on page 6)
- Brian Wampler, Stephanie McNulty, and Michael Touchton. 2021. *Participatory Budgeting in Global Perspective*. Oxford University Press. (Cited on pages 5, 6, 218, and 248)
- H. Peyton Young. 1974. An Axiomatization of Borda’s Rule. *Journal of Economic Theory* 9, 1 (1974), 43–52. (Cited on page 8)
- H. Peyton Young. 1995. Optimal Voting Rules. *Journal of Economic Perspectives* 9, 1 (1995), 51–64. (Cited on page 121)

- H. Peyton Young and Arthur Levenglick. 1978. A Consistent Extension of Condorcet's Election Principle. *SIAM Journal on Applied Mathematics* 35, 2 (1978), 285–300. (Cited on page 8)
- 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. (Cited on pages 220 and 230)

Index

A

Additive JA rule	143
Asymmetric	167
Adrian	219
Alina	180
Allocation stage	15
Argmax rule	125
Monotonic	125
Arianna	139
Awareness set	222

B

Ballot	26
Approval	16, 27
Cardinal	27
Cumulative	28
Enriched approval	27
Knapsack	16
Ordinal	28
Budget allocation	16
Exhaustive	16
Feasible	16

C

Camille	140
Capped fair share ratio	111
Cohesive group	
With approval ballots	41

With cardinal ballots	37
-----------------------------	----

Completion principle	223
----------------------------	-----

Constraint monotonicity	170
-------------------------------	-----

Core

With approval ballots	49
-----------------------------	----

With cardinal ballots	46
-----------------------------	----

α -entitlement approximation,	48
--------------------------------------	----

α -sat approximation,	47
------------------------------	----

D

Daira	180
-------------	-----

Dean	176
------------	-----

Determined rule	233
-----------------------	-----

DNNF Circuit	146
--------------------	-----

E

Embedding	145
-----------------	-----

Exhaustive tractable	166
----------------------------	-----

Tractable	149
-----------------	-----

Tractable with dependencies ..	154
--------------------------------	-----

Tractable with quotas	162
-----------------------------	-----

Emma	140
------------	-----

Evaluation function	180
---------------------------	-----

Cost	181
------------	-----

Relative cost	182
---------------------	-----

Share	183
-------------	-----

Exhaustiveness	16, 59
----------------------	--------

- F**
- Fair share 81
 - Local 87
 - Up to one project 83
 - Fairness criteria for PPB
 - Convergence to equal evaluation (EQUAL- Φ) 186
 - Equal evaluation (EQUAL- Φ) .. 184
 - Gini (Φ -GINI) 185
 - First-stage strategy-proofness 232
- G**
- Greedy scheme 19
 - Greedy welfare rule
 - Cardinality 19
 - Cost 19
 - Ground truth 122
- I**
- Ideal set 223
 - Instance 15
 - For perpetual participatory budgeting 178
 - Shortlisting 222
 - With multiple resources 142
 - With project dependencies ... 152
 - With quotas on projects 159
 - Integrity constraint 143
 - Irresoluteness 17
- J**
- Jan 219
 - Judgement aggregation 142
 - Julian 180
 - Justified representation
 - With approval ballots
 - Extended, 42, 108, 109
 - Extended up to any project, 42
 - Full, 53
 - Proportional, 44
 - Proportional up to any project, 44
 - Strong extended, 41
 - With cardinal ballots
 - Extended, 38
 - Extended up to one project, 39
 - Full, 52
 - Proportional, 40
 - Strong extended, 38
 - Justified share
 - Extended 89
 - Extended up to one project 95
 - Local extended 95
 - Strong extended 88
- K**
- Kemeny rule 144
- L**
- L_1 distance to FS 112
 - Leximax rule 144
 - Lwenn 180
- M**
- Maximin support rule 34
 - Maximum likelihood estimator ... 122
 - Method of equal shares 21, 35
 - Monotonicity axiom
 - Discount 60, 170
 - Limit 61, 169
 - Merging 61, 172
 - Splitting 61, 171
- N**
- Noise model 122
 - Non-wastefulness 227
 - Normalised satisfaction 129
- O**
- Outcome determination problem ... 66
- P**
- Pathwidth 154
 - Perpetual participatory budgeting . 177

- Preferences 26
 Preprocessing
 Cohesiveness 113
 Threshold 113
 Priceability 49, 105
 Profile
 For judgment aggregation 143
 For participatory budgeting 16
 Shortlisting 222
 Proportionality for solid coalitions
 Comparative 51
 Inclusion 52
 Propositional logic 142
- R**
- Representation efficiency 228
 Resoluteness 17
 Rule 17
- S**
- Satisfaction 26
 Satisfaction function 17
 Cardinality 17
 Cost 17
 DNS 43
 Sequential Phragmén rule 20, 34
 Share 79
 Shortlisting rule 222
 Equal representation 225
 Median 226
 Nomination 224
 Sirin 139
 Slater rule 144
 Social welfare
- Chamberlin-Courant 67
 Egalitarian 67
 Nash 67, 126
 Utilitarian 18, 66, 132
 Solid coalition 51
 Solution 178
 Sophie 219
 Strategy-proofness
 With approval ballots 63
 Approximate, 64
 With cardinal ballots 63
- T**
- Tie-breaking rule 221
 Canonical 221
 Types of agents 178
 Types of projects 159
- U**
- Unanimity 235
 Project-wise 110
 Unit costs 15
 Utility 26
- V**
- Voting stage 15
- W**
- Weak reinforcement 123
- Z**
- Zoi 139
 Zoiville 139

List of Symbols

Building Blocks

\mathcal{P}	Set of projects	15
c	Cost function	15
b	Budget limit	15
I	PB instance	15
π	Budget allocation	16
$\text{FEAS}(I)$	Set of feasible budget allocations for instance I	16
$\text{FEAS}_{\text{EX}}(I)$	Set of feasible and exhaustive budget allocations for instance I	16

Voters

\mathcal{N}	Set of voters/agents	16
A_i	Ballot of agent i	16
\mathbf{A}	Profile of ballots	16
app	Approval score	16
$+$	Profile concatenator	16
\succ_i	Strict ordinal ballot of i	28

\succsim_i	Weak ordinal ballot of i	28
sat	Approval-based satisfaction function	17
sat_i	Approval-based satisfaction function of agent i	17
sat^{card}	Cardinality satisfaction function	17
sat^{cost}	Cost satisfaction function	17
sat^{CC}	Chamberlin-Courant satisfaction function	31
sat^{\log}	Log satisfaction function	31
$sat^{\sqrt{\cdot}}$	Square root satisfaction function	31
$\overline{sat^{card}}$	Normalised cardinality satisfaction	129
$\overline{sat^{cost}}$	Normalised cost satisfaction	129
$relsat_{sat}$	Relative satisfaction of sat	182

Rules

F	PB rule	17
---	-------------------	----

MAXCARD	Cardinality welfare maximising rule	18
MAXCOST	Cost welfare maximising rule	18
GREEDCARD	Greedy cardinality welfare rule	19
GREEDCOST	Greedy cost welfare rule	19
SEQPHRAG	Sequential Phragmén rule	20
MES[<i>sat</i>]	Method of equal shares for <i>sat</i>	21
MES	Method of equal shares for cardinal ballots	35
MAXIMINSUPP	Maximin support rule	34
UNTIE	Arbitrary tie-breaking rule	221
CANONUNTIE	Canonical tie-breaking rule	221
Other PB Symbols		
UTIL-SW	Utilitarian social welfare	33
UTIL-SW[<i>sat</i>]	Utilitarian social welfare for <i>sat</i>	18
CC-SW	Chamberlin-Courant social welfare	67
EGAL-SW	Egalitarian social welfare	67
NASH-SW	Nash social welfare	67
GREED	Greedy scheme	19
Equality of Resources		
<i>share</i>	Share	79

<i>fairshare</i>	Fair share	81
------------------	------------	----

Epistemic Approach

π^*	Ground truth	122
\mathcal{M}	Noise model	122
$\mathbb{P}_{\mathcal{M}}$	Probability distribution under \mathcal{M}	122
$L_{\mathcal{M}}$	Likelihood under \mathcal{M}	122
$Z_{\pi^*}^{Ncost}$	Normalisation factor of \mathcal{M}_{Ncost} with respect to π^*	127
$Z_{\pi^*}^{app}$	Normalisation factor of \mathcal{M}_{app} with respect to π^*	132

Multi-Constraint PB

\mathcal{R}	Set of resources	142
d	Number of resources	142
\mathcal{I}	Set of all instances of PB with multiple resources	142
\mathbf{b}	Budget vector	142
\mathfrak{X}	Set of propositional atoms	142
$\mathcal{L}_{\mathfrak{X}}$	Set of propositional formulas over \mathfrak{X}	142
$Lit(X)$	Set of literals in X	142
ℓ_{x_i}	Literal corresponding to x_i	142
$aug(J)$	Augmented judgment set	143
Γ	Integrity constraint	143
$\mathfrak{J}(\Gamma)$	Set of admissible judgments with respect to Γ	143
J_i	Judgment of agent i	143
\mathbf{J}	Judgment profile	143
$n_{\mathbf{J}}^{\ell}$	Number of supporters of ℓ in \mathbf{J}	143

$m(\mathbf{J})$	Majoritarian outcome	143	Φ_{marg}^{cost}	Marginal cost evaluation function	181
F	Judgement aggregation rule	143	$\Phi^{relcost}$	Relative cost evaluation function	182
F_{asy}	Asymmetric judgement aggregation rule	167	$\Phi_{marg}^{relcost}$	Marginal relative cost evaluation function	182
τ	Outcome translation	145	Φ^{share}	Share evaluation function	183
E	Embedding	145	Φ_{marg}^{share}	Marginal share evaluation function	183
TE	Tractable embedding	149	EQUAL- Φ	Equal evaluation	184
TE_{dep}	Tractable embedding with dependencies	154	GINI	Gini coefficient	185
TE_{quo}	Tractable embedding with quotas	163	GINI $_{\Phi}$	Gini coefficient with respect to Φ	185
ETE	Exhaustive tractable embedding	166	Φ -GINI	Gini optimality with respect to Φ	185
Imp	Implication set	152	EQUAL- Φ -CONV	Convergence to equal evaluation	186
\mathfrak{T}	Type system	159	End-to-End Model		
\mathcal{T}	Set of project types	159	\mathbb{P}	Set of conceivable projects	221
Q	Quota group	159	C_i	Awareness set of agent i	222
q	Quota function	160	\mathcal{C}	Awareness profile	222
f^{card}	Cardinality-type aggregator	160	P_i	Shortlisting ballot of i	222
f^{cost}	Cost-type aggregator	160	\mathbf{P}	Shortlisting profile	222
Long-Term PB			SHORT	Shortlisting rule	222
I	PPB instance	178	NOMSHORT	Nomination shortlisting rule	224
π	Solution	178	REPRSHORT	Equal representation shortlisting rule	225
\mathcal{T}	Set of agent types	178	MEDIANSHORT	Median shortlisting rule	226
Φ	Evaluation function	180			
Φ_{marg}	Marginal evaluation function	180			
Φ^{cost}	Cost evaluation function	181			

\triangleright_i	Agent i 's ordinal preferences over the projects 222	$undom(\succsim, \mathfrak{P})$	Undominated subsets of \mathfrak{P} according to \succsim 224
$top_i(P)$	Ideal subset of projects of agent i for P 223	$A_i^*(I, \mathbf{A}, F)$	Best response of i given I , \mathbf{A} and F 230
$top(P)$	Ideal profile 223	Other Symbols	
\succsim_P	Extended preferences given P 223	$\mathbb{1}_\varphi$	Indicator function for φ 21
\sim_P	Indifference part of \succsim_P 223	GCD	Greatest common divisor ... 69
\succ_P	Strict part of \succsim_P 223	Δ	Symmetric difference 222
		$\mathbf{0}_d$	Zero-vector of length d 149

Titles in the ILLC Dissertation Series

ILLC DS-2018-07: **Julian Schlöder**

Assertion and Rejection

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: **Iliaria 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