The (Computational) Social Choice Take on Indivisible Participatory Budgeting

Simon Rey
s.j.rey@uva.nl

and

Jan Maly
j.f.maly@uva.nl

Institute for Logic, Language and Computation (ILLC)
University of Amsterdam
Amsterdam, the Netherlands

Abstract

In this survey, we review the literature investigating participatory budgeting as a social choice problem. Participatory Budgeting (PB) is a democratic tool aiming at making budgeting decisions in a more democratic manner. Specifically, citizens are asked to vote on who to allocate a given amount of money to a set of projects that can potentially be funded. From a social choice perspective, it corresponds then to the problem of aggregating opinions about which projects should be funded, into a budget allocation satisfying a budget constraint. This problem has received substantial attention in recent years and the literature is growing at a fast pace. In this survey, we present the most important research directions from the literature, each time presenting a large set of representative results. We only focus on the *indivisible* case, that is, PB problems in which projects can either be fully funded or not at all.

The aim of the survey is to present a comprehensive overview of the state of the research on PB. We aim at providing both a general overview of the main research questions that are being investigated, and formal and unified definitions of the most important technical concepts from the literature.

Of course a survey is never complete as the state of the research keeps changing. This document is intended to be a living document that gets updated every now and then as the literature grows. If you feel that some papers are not presented correctly, or simply missing, feel free to contact us. We will be more than happy to correct it.

Acknowledgements

We would like to thank all who helped writing and improving this survey. Throughout the process, we have received very valuable feedback from Piotr Faliszewski, Martin Lackner, Jannik Peters and Dominik Peters, we are very grateful to all of you.

Contents

Chapter 1

Introduction

Participatory budgeting (PB) is a recent democratic innovation that aims to involve citizens in budgeting decisions. It is one of the most successful democratic innovations in recent years (Wampler, McNulty and Touchton, 2021). Since its first implementation in Porto Alegre, Brazil, in 1989, it spread around the world and is now implemented in every continent and in most countries (Dias, 2018; Dias, Enríquez and Júlio, 2019). Given its success, a wide variety of processes have been implemented. Most of these processes however follow the same key steps (Wampler, 2000; Cabannes, 2004; Shah, 2007):

- 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 that proposed 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.

Note that the steps above have been phrased as if the organising entity was a municipality, the typical case. However, the scale of the process can vary significantly, from a neighbourhood of a city—as, for instance, in Amsterdam (City of Amsterdam, 2022)—to subnational entities—for example regional departments in Peru (Shah, 2007). There also are examples of PB processes implemented in schools[1], or housing communities[2].

It is also interesting to note that not all the processes include a voting stage. Indeed, sometimes the PB process is just organised as a deliberative mechanism throughout where the set of projects to implement is determined meeting after meeting. This was typically the case for the first PB process implemented in Brazil (Cabannes, 2004).

As should be clear from the typical structure outlined above, several steps of a PB process involve the citizens' participation: first when submitting proposals, and second when voting on the shortlisted projects. This perspective on PB makes it a typical social choice problem and a burgeoning literature on PB has emerged from the (computational) social choice community,

[1] See for instance participatorybudgeting.org/pb-at-ps139 for an example of PB processes within primary schools.

[2] See 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.

focusing predominantly on the voting stage. The aim of this survey is to present the main findings coming from this line of research. It complements the first survey of Aziz and Shah (2021) which was written before the sharp increase of publications on the topic and only covers the literature until 2019.

In contrast to the survey of Aziz and Shah (2021), we only focus on indivisible participatory budgeting, also called *discrete* PB (Aziz and Shah, 2021), that is, the special case of PB where projects can either be fully funded or not all (projects cannot be fractionally implemented). Within this framework, we present what we believe to be the most important concepts and results. We aim to provide a comprehensive set of definitions and to unify concepts and notations that appeared in different publications.

We will first present the basic model and our notations (Chapter 2). Then we turn to the different ballot formats that have been proposed for PB (Chapter 3). Once the design of the ballots will be clarified, we will discuss rules for aggregating said ballots (Chapter 4). We will then present how to asses the quality of these rules in terms of fairness (Chapter 5) and other axiomatic properties (Chapter 6). After that, we will look at the algorithmic aspects of PB (Chapter 7). Having discussed the standard model for PB, we will then present variations and extensions of the standard model that have been introduced (Chapter 8). We will finally provide interesting pointers to go beyond what we presented in the survey, be it related frameworks or actual implementation of PB in practice (Chapter 9). We will conclude this survey by mentioning what we consider to be the most important directions for future work (Chapter 10).

Chapter 2

Preliminaries

Almost the entire computational social choice literature focuses on the voting stage of PB. The only exception we are aware of is the work of Rey, Endriss and de Haan (2021). The voting stage will also be the main focus of this survey. In the following we introduce the standard model of the voting stage of PB processes. We then try to clarify different related concepts: preferences, utilities, satisfaction and ballots.

2.1 The Standard Model of PB

The voting stage of a PB process is represented as a tuple of three elements $I = \langle \mathcal{P}, c, b \rangle$ called an *instance* where $\mathcal{P} = \{p_1, \ldots, p_m\}$ is the *set of projects*; $c : \mathcal{P} \to \mathbb{R}_{>0}$ 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[1] that both the costs and the budget limit have to be positive. 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 $b \in \mathbb{N}_{>0}$. These instances are especially interesting because they correspond to multi-winner elections (Lackner and Skowron, 2023).

Let $\mathcal{N} = \{1, \ldots, n\}$ be the *set of voters* involved in the PB process, these are the citizens participating in the process, not the officials/organisers. When facing an instance $I = \langle \mathcal{P}, c, b \rangle$, they are asked to submit their preferences over the projects in $\mathcal{P}$. They do so by submitting a ballot whose format is determined by the rules of the process. Several *ballot formats* have been considered for PB as we shall see later. For now, let us denote by A_i the ballot that voter $i \in \mathcal{N}$ is submitting. The vector $\boldsymbol{A} = (A_1, \ldots, A_n)$ of the ballots of the voters is called a *profile*. Note that we use the terms voters and agents interchangeably, purely for stylistic reasons.[2]

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 $\textsc{Feas}(I)$ the set of all *feasible* budget allocations for instance I, defined as $\textsc{Feas}(I) = \{\pi \subseteq \mathcal{P} \mid c(\pi) \leq b\}$.

Budget allocations are determined using PB rules. A *PB rule* R is a function taking as input an instance I and a profile $\boldsymbol{A}$ and returning a set of feasible budget allocations $\mathrm{R}(I, \boldsymbol{A}) \subseteq \textsc{Feas}(I)$. 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 thus potentially return several *tied* budget allocations. Unless explicitly stated, we will assume rules to be *resolute*. At a few occasions we will discuss randomised PB rules, which

[1] Motamed, Soeteman, Rey and Endriss (2022) is the only paper considering negative costs, which only is relevant since they also consider multi-dimensional costs.

[2] The reader may get bored to always read the same terminology all the time.

are rules that return for any instance I and profile $\boldsymbol{A}$, not a budget allocation, but a probability distribution over $\textsc{Feas}(I)$.

In the coming chapters, and particularly in chapters 5 and 6, we will introduce several properties of budget allocations. To avoid unnecessary definitions, we will use the exact same properties for rules. For a given property $\mathcal{X}$ of a budget allocation, we say that a rule R satisfies $\mathcal{X}$ if for every instance I and profile $\boldsymbol{A}$, the outcome of the (resolute) rule $\mathrm{R}(I, \boldsymbol{A})$ satisfies $\mathcal{X}$. When needed, we will explicitly specify how properties of budget allocations are lifted to *irresolute* rules.

2.2 The Voters: Preferences, Utility, Satisfaction and Ballots

Going through the literature on PB, and more generally about 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 cannot be claimed to capture it 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 ballot to be truthful (Gibbard, 1973; Satterthwaite, 1975; Dietrich and List, 2007; 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 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, about their internal state and about 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 that reflect their view on the possible outcomes of the decision making scenario. Remember from the above that this information may not be accessible in full by 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.*, they are based on ranking of the outcomes.

- **Utility**: The utility of a voter is a specific type of preferences 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 is often used synonymously with their utility. In computationally 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 shall see concrete example later in this survey. 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. This information is formatted in the format required by the type of ballots that is being used. Let us emphasise once again that a ballot is the sole the information submitted by the (potentially strategically-behaving) voter and not necessarily a representation of their private information.

In the following we will adopt those definitions and try to use the terms accordingly.

Chapter 3

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 possibility for the agents to submit their preferences over all possible budget allocations. These could take the forms of orderings over $\textsc{Feas}(I)$, or utility functions associating a score to every feasible budget allocation $\pi \in \textsc{Feas}(I)$. This approach clearly cannot be implement in real life as the size of $\textsc{Feas}(I)$ is exponential in the number of projects, which in itself might already be quite large (in 2023 there were 138 projects in the municipal Warsaw PB process[1]).

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 format 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. In what follows, we distinguish between cardinal ballots (Section 3.1)—that associate a score to each projects—and ordinal ballots (Section 3.2)—that require agents to rank the projects. We will conclude this section by comparing the different formats (Section 3.3) and discussing how to define satisfaction function based on different ballots (Section 3.4).

To get an overview of the different ballot formats that have been introduced and the papers studying them, we present in Table 3.1 a classification of the papers we have reviewed, based on the ballot format they are considering.

3.1 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 sometimes 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 2.2.

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

[1]See 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) — Lackner, Maly and Rey (2021) — Los, Christoff and Grossi (2022) — Maly, Rey, Endriss and Lackner (2023) — Motamed, Soeteman, Rey and Endriss (2022) — Rey, Endriss and de Haan (2020) — Rey, Endriss and de Haan (2021) — Sreedurga, Bhardwaj and Narahari (2022) — Talmon and Faliszewski (2019)
t-Approval	Fairstein, Benadè and Gal (2023)
Knapsack	Benadè, Nath, Procaccia and Shah (2021) — Fairstein, Benadè and Gal (2023) — Goel, Krishnaswamy, Sakshuwong and Aitamurto (2019)
t-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)

*Laruelle (2021) considers that the agents submit weak and complete rankings over the projects, that are then converted into cardinal scores (via positional scoring functions) for the aggregation.

Table 3.1: Papers studying indivisible PB based on the type of ballots they consider. We categorised the papers based on the main ballot format used in their study, not necessarily based on *all* the format mentioned in the paper.

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.

3.1.1 Approval Ballots

When using approval ballots, agents are asked to submit a subset of projects they approve of. We represent approval ballots as cardinal ballots by requiring the score of each project to either be 0 or 1. For agent $i \in \mathcal{N}$, their approval ballot $A_i : \mathcal{P} \to \{0, 1\}$ 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. We will sometimes call voter $i \in \mathcal{N}$ a *supporter* of project $p \in \mathcal{P}$ whenever $A_i(p) = 1$.

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).

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 to the project. One way to circumvent this issue is by enforcing additional constraints on the ballots, that allow us to interpret them more accurately.

3.1.2 Semantically Enriched Approval Ballots

As explained above, the semantics of approval ballots is not well defined. This lead 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 $\sum_{p \in \mathcal{P}} A_i(p) \leq t$ for the ballot A_i of all agents $i \in \mathcal{N}$. This allows us to get some understanding of the not approved projects: they are not part of the top-t projects of the voter (assuming that voters can actually order the projects based on their preferences).

One important variation of approval ballots, both in theoretical terms and because of its actual usage, is the *knapsack ballot* (Goel, Krishnaswamy, Sakshuwong and Aitamurto, 2019). A knapsack ballot is an approval ballot with the additional constraint that the total cost of the approved projects cannot exceed the budget limit b. Formally speaking, it is an approval ballot A_i such that $c(\{p \in \mathcal{P} \mid A_i(p) = 1\}) \leq b$. 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 semantically enriched variation of approval ballot are *t-threshold approval ballots* (Benadè, Nath, Procaccia and Shah, 2021; Fairstein, Benadè and Gal, 2023). Here, agents are assumed to have private additive utility functions that they are aware off, and they are asked to submit an approval ballot, approving of a project if and only if it provides them with utility at least $t \in \mathbb{R}$.

3.1.3 Cumulative Ballots

When using cumulative ballots (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.

3.2 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 cost, 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 one 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 ranking of the projects based on their value-for-money. Note that this is only well defined when agents are assumed to have private utility function that they are aware off.

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 $p \sim p'$ if $p \succsim p'$ and $p' \succsim p$; and $p \succ p'$ if $p \succsim p'$ but not $p' \succsim p$, for any two projects p and p'. Of course rankings by value or value-for-money can be considered either as strict or weak rankings.

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, that we could call *t-ordinal ballots*.

3.3 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.

3.3.1 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 submitted their ballots according to the specific format.

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 3.2: 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.

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 3.2 presents their findings. Note that they 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.[2]

3.3.2 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 is 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. 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 k-approval ballots outperform all the other ballot formats, though not by a large margin.

3.4 Ballot-Based Satisfaction

Before we consider how to use the ballots to determine budget allocations through PB rules, let us discuss how to model satisfaction based on the different ballot format we have introduced. Many

[2] The 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 information alone is not enough to make a high quality decision.

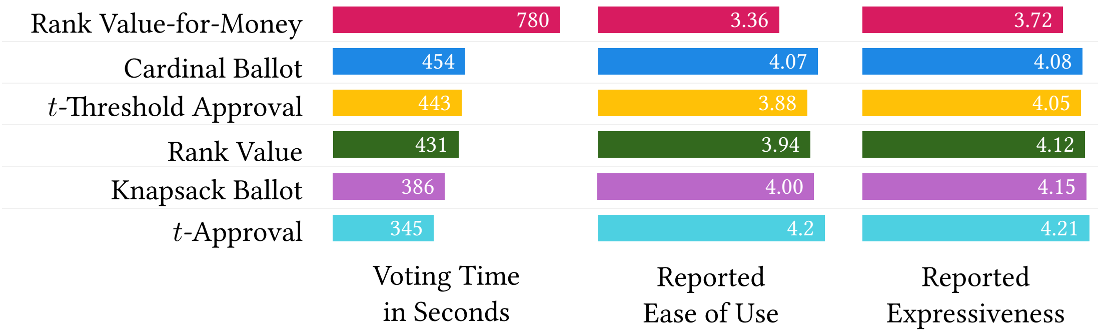

Figure 3.1: Some of the experimental findings of Fairstein, Benadè and Gal (2023) comparing different ballot formats. The voting time column indicates the time in seconds it took participants to submit their opinion for each ballot format. The reported ease of use and expressiveness columns represents 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.

of the concepts that we will introduce in the rest of this paper rely on measures of satisfaction.

3.4.1 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 under the *additivity* assumption. This means the satisfaction of a voter is the sum of the score they submitted for the projects that have been selected.

3.4.2 Approval Ballots

When it comes to 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. Let us provide their definition.

Definition 1 (Approval-Based Satisfaction Functions)**.** *Given an instance* $I = \langle \mathcal{P}, c, b \rangle$ *and a profile* $\boldsymbol{A}$*, an* (approval-based) satisfaction function *is a mapping* $sat : 2^{\mathcal{P}} \to \mathbb{R}_{\geq 0}$ *satisfying the following two conditions:*

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

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

$$sat_i(\pi) = sat(\{p \in \pi \mid A_i(p) = 1\}).$$

Note that in contrast to the case of cardinal ballots, satisfaction functions are not generally assumed to be additive. However, we will sometimes make this assumption, *i.e.*, requiring that $sat(P) = \sum_{p \in P} sat(\{p\})$ for any $P \subseteq \mathcal{P}$.

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 author's 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). See the discussion below for more details.

Several satisfaction functions have been introduced in the literature, we define them below.

- **Cardinality Satisfaction Function** (Talmon and Faliszewski, 2019): measures the satisfaction of the voters as the number of selected and approved projects:
$$sat^{card}(P) = |P|.$$
- **Cost Satisfaction Function** (Talmon and Faliszewski, 2019): measures the satisfaction of the voters as the cost of the selected and approved projects:
$$sat^{cost}(P) = c(P).$$
Note that with indivisible projects, this is equivalent to the *overlap satisfaction function* of Goel, Krishnaswamy, Sakshuwong and Aitamurto (2019).
- **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:
$$sat^{CC}(P) = \mathbb{1}_{P \neq \emptyset}.$$
- **Share** (Lackner, Maly and Rey, 2021): measures the resources the decision maker used to satisfy the voters:
$$sat^{share}(P) = \sum_{p \in P} \frac{c(p)}{|\{i \in \mathcal{N} \mid A_i(p) = 1\}|}.$$
It is important to keep in mind that the share has not been introduced as a satisfaction function but can still be interpreted as one (while being cautious as to how to use it).
- **Square Root and Log Satisfaction Functions** (Brill, Forster, Lackner, Maly and Peters, 2023): measures the satisfaction of the voters as (marginally) diminishing when the cost of a project increases:
$$sat^{\log}(P) = \log(1 + c(P)) \qquad\qquad sat^{\sqrt{}}(P) = \sqrt{c(P)}.$$

Both the cardinality and the cost satisfaction are quite standard within the literature, even though they can easily be criticised: there is no good reason to assume that the satisfaction of an agent is the same for two projects, one being very expensive while the other being particularly cheap; though it is also not sensible to assume a perfect correlation between satisfaction and cost.

In general, all the above apply seamlessly to all approval-like ballots (t-approval, knapsack, t-threshold...). Some satisfaction functions are however more meaningful with some ballots than others. In particular, additional satisfaction functions could be interesting to study when using semantically richer approval ballots.

It is also worth noting that Brill, Forster, Lackner, Maly and Peters (2023) presented results that apply to whole classes of satisfaction functions, and not just functions from the list above. Some of these results will be presented later (notably in Section 5.1 and 5.3).

Finally, note that a satisfaction function as defined by Brill, Forster, Lackner, Maly and Peters (2023) only depends on the projects in the winning bundle that the voter approved. Therefore, these functions cannot capture satisfaction functions that also depend on the non-approved projects in the winning bundle or on the approved projects that have not been funded.

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 it to be of very limited relevance when the projects are indivisible.

Lackner, Maly and Rey (2021) proposed the notion of *relative satisfaction*, which normalises the satisfaction of a voter by the maximum satisfaction achievable:

$$relsat_{sat}(P) = \frac{sat(P)}{\max\{sat(P') \mid P' \in \textsc{Feas}(I) \text{ and } A_i(p) = 1, \forall p \in P'\}},$$

where sat is any satisfaction function. Lackner, Maly and Rey (2021) only considered relative satisfaction associated with the cost satisfaction function sat^{cost}. This can also not be modelled as an approval-based satisfaction function, as it depends on the full approval ballot of the voter.

3.4.3 Ordinal Ballots

To measure satisfaction with ordinal ballots, one can associate each project in the ordering to a given satisfaction level. This is usually done through *positional scoring functions* that associate to each project 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 contest-dependent. Note that this assumption is never explicitly stated and that this reflects our understanding of their definitions.

Chapter 4

Participatory Budgeting Rules

We have seen many ways of collecting the opinion of the voters. The next natural step is thus to use that information to select "good" budget allocations. This is done through the use of PB rules. In this section we will present the main rules that have been introduced in the literature. Note that in what follows, and in almost the entirety of the paper, we will mainly focus on cardinal and approval ballots.

Our exposition will start with welfare maximising rules (Section 4.1). 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 4.2), the maximin support rule (Section 4.3), and the method of equal shares (Section 4.4). A brief overview of the other rules that have been introduced in PB will conclude this part of our survey (Section 4.5).

4.1 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 as follows: for a given instance $I = \langle \mathcal{P}, c, b \rangle$, budget allocation $\pi \in \textsc{Feas}(I)$ and a utility function $\mu_i : 2^{\mathcal{P}} \rightarrow \mathbb{R}_{\geq 0}$ for every agent $i \in \mathcal{N}$:

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

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 of an agent given their 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 $\boldsymbol{A}$, select the budget allocation that maximises Util-SW:[1]

$$\textsc{Util-SW}(I, \boldsymbol{A}, \pi) = \sum_{i \in \mathcal{N}} \sum_{p \in P} A_i(p).$$

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

[1]Note 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 by the voters.

Given an instance I and a profile $\boldsymbol{A}$, selecting the budget allocation $\pi \in \textsc{Feas}(I)$ that maximises $\textsc{Util-SW}(I, \boldsymbol{A}, \pi)$ defines a rule, the utilitarian welfare maximising rule.

We thus have seen a first example of a PB rule. In what follows we will review other utilitarian rules that have been introduced in the literature.

4.1.1 Exact Welfare Maximisation with Approval Ballots

In Section 3.4, we have introduced so-called satisfaction functions to measure the satisfaction of the voters when using approval ballots. The definition of the utilitarian social welfare can then be parametrised by a satisfaction function. Given a satisfaction function sat, the utilitarian social welfare of a budget allocation $\pi \in \textsc{Feas}(I)$ given an instance $I = \langle \mathcal{P}, c, b \rangle$, profile $\boldsymbol{A}$ of approval ballots is defined as:

$$\textsc{Util-SW}[sat](I, \boldsymbol{A}, \pi) = \sum_{i \in \mathcal{N}} sat_i(\pi).$$

Remember that $sat_i(\pi) = sat(\{p \in \pi \mid A_i(p) = 1\})$.

Among the first PB rules to have been introduced in the literature are two utilitarian welfare maximisation rules (Talmon and Faliszewski, 2019). They make use of the cardinality and cost satisfaction functions.

The *cardinality welfare maximising rule* $\textsc{MaxCard}$ is defined for any instance I and approval profile $\boldsymbol{A}$ as:

$$\begin{aligned}\textsc{MaxCard}(I, \boldsymbol{A}) &= \arg\max_{\pi \in \textsc{Feas}(I)} \textsc{Util-SW}\left[sat^{card}\right](I, \boldsymbol{A}, \pi) \\ &= \arg\max_{\pi \in \textsc{Feas}(I)} \sum_{i \in \mathcal{N}} |\{p \in \pi \mid A_i(p) = 1\}|.\end{aligned}$$

Similarly, the *cost welfare maximising rule* $\textsc{MaxCost}$ is defined for any instance I and approval profile $\boldsymbol{A}$ as:

$$\begin{aligned}\textsc{MaxCost}(I, \boldsymbol{A}) &= \arg\max_{\pi \in \textsc{Feas}(I)} \textsc{Util-SW}\left[sat^{cost}\right](I, \boldsymbol{A}, \pi) \\ &= \arg\max_{\pi \in \textsc{Feas}(I)} \sum_{i \in \mathcal{N}} c(\{p \in \pi \mid A_i(p) = 1\}).\end{aligned}$$

These definitions give rise to *irresolute rules*. Remember that we want to work with resolute rules in this survey. They can be obtained by using some fixed tie-breaking mechanism between all budget allocations maximising $\textsc{Util-SW}$.

Interestingly, these two rules can be reinterpreted in terms of *approval score*: given an instance I and a profile $\boldsymbol{A}$ of approval ballots, the approval score of a project $p \in \mathcal{P}$ in $\boldsymbol{A}$, denoted by $app(p, \boldsymbol{A})$, is defined as $app(p, \boldsymbol{A}) = |\{i \in \mathcal{N} \mid A_i(p) = 1\}|$. For any I and $\boldsymbol{A}$, we then have:

$$\begin{aligned}\textsc{MaxCard}(I, \boldsymbol{A}) &= \arg\max_{\pi \in \textsc{Feas}(I)} \sum_{p \in \pi} app(p, \boldsymbol{A}), \\ \textsc{MaxCost}(I, \boldsymbol{A}) &= \arg\max_{\pi \in \textsc{Feas}(I)} \sum_{p \in \pi} app(p, \boldsymbol{A}) \cdot c(p).\end{aligned}$$

These two formulation will prove useful when drawing parallel with the knapsack problem (Kellerer, Pferschy and Pisinger, 2004).

As we will see later (Section 7.2), it is computationally difficult to compute the outcome of these two rules, at least at a theoretical level.[2] For this reason, greedy approximations of the utilitarian social welfare have also been considered.

[2] In practice, there are rather efficient techniques for solving knapsack problems that can be used.

4.1.2 Greedy Approximation of the Welfare Maximiser with Approval Ballots

Exploiting the connection between the maximisation of Util-SW and various knapsack problems. We can use the prolific literature on the topic (Kellerer, Pferschy and Pisinger, 2004) to derive PB rules approximating the maximum utilitarian social welfare. Let us first define the general scheme of a greedy rule.

Definition 2 (Greedy Scheme)**.** *Consider an instance $I = \langle \mathcal{P}, c, b \rangle$ and a strict ordering $\rhd$ over $\mathcal{P}$. The* greedy scheme $\text{Greed}(I, \rhd)$ *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 $\rhd$. When considering project p for current budget allocation π, p is selected (added to π) if and only $c(\pi \cup \{p\}) \leq b$. If there is a next project according to $\rhd$, it is considered; otherwise π is the output of* $\text{Greed}(I, \rhd)$.

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).

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

$$\text{GreedCard}(I, \boldsymbol{A}) = \{\text{Greed}(I, \rhd) \mid \rhd \text{ is compatible with } {}^{app}/_{c}\}.$$

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

$$\text{GreedCost}(I, \boldsymbol{A}) = \{\text{Greed}(I, \rhd) \mid \rhd \text{ is compatible with } app\}.$$

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

Interestingly, we know from the knapsack literature that these two greedy rules approximate their respective welfare objective within a factor 2 (Kellerer, Pferschy, and Pisinger, 2004, Chapter 2).[3] This is particularly clear when considering the approval-score based definitions of MaxCard and MaxCost.

A final important fact to keep in mind is that the greedy cost welfare rule, GreedCost, is actually the rule that is the most widely used in practice. This makes it an important rule to consider in any analysis.

4.1.3 Other Welfare-Based Rules

On top of the four rules we defined above, Talmon and Faliszewski (2019) introduce five extra rules. They additionally consider welfare defined in terms of sat^{CC} (see Section 3.4), 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

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

social welfare in Section 7.2—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.

Our focus was mainly on approval ballots, though a similar approach has been followed for cardinal ones. Fluschnik, Skowron, Triphaus and Wilker (2019) study utilitarian and 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 7.2). Their motivation is more algorithmic, however, and they don't necessarily aim to devise PB rules.

Finally coming to ordinal ballots, Laruelle (2021) studies welfare optimising rules with weak ordinal ballots where positional scoring functions are used to measure 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 7.2 for a formal definition).

4.2 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 aims to provide *proportional representation*, which will be studied in more detail in Chapter 5.

This rule can only be applied with approval ballots. It was formally studied in the multi-winner literature by Brill, Freeman, Janson and Lackner (2017), and has then been adapted for the PB setting by Los, Christoff and Grossi (2022).

Definition 3 (Sequential Phragmén, Continuous Formulation)**.** *Given an instance I and a profile $\boldsymbol{A}$ of approval ballots, 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 t money. For any time t, let $P_t^\star$ 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, $P_t^\star$ is non-empty, if there exists a $p \in P^\star$ such that $c(\pi \cup \{p\}) > b$, the process stops; otherwise one project from $P_t^\star$ is selected, the budget of its supporters is set to 0, and the process resumes.

Breaking the ties among the projects in any $P_t^\star$ in the above definition will lead to a resolute rule. In the irresolute variants, 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[4], which however comes at the cost of exhaustiveness (see Sections 5.3 and 6.1).

The sequential Phragmén rule can also be formalised in a discrete fashion where the *loads* of the voters are to be balanced. These two formulation are equivalent. We provide below the second formulation (see, *e.g.*, Brill, Forster, Lackner, Maly, and Peters, 2023).

[4]Note that phrasing 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.

Definition 4 (Sequential Phragmén, Discrete Formulation)**.** *Given an instance I and a profile $\boldsymbol{A}$ of approval ballots, the* sequential Phragmén rule, SeqPhrag, *constructs a budget allocation π, initially empty, iteratively as follows. A load $\ell_i : 2^{\mathcal{P}} \to \mathbb{R}_{\geq 0}$, 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:*

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

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

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

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

As before, to obtain a resolute rule one needs to break the ties among the projects in any $P^\star$. The irresolute variant is obtained by considering all possible ways of breaking such ties.

4.3 The Maximin Support Rule

We can also adapt the definition of sequential Phragmén slightly and allow a redistribution of the loads in each round. This leads to the definition of the *maximin support rule*. This rule was first introduced by Aziz, Lee and Talmon (2018) in the PB setting. Note that they named it sequential Phragmén though the actual rule they define is a generalisation of the maximin support rule from multi-winner voting (Sánchez-Fernández, Fernández-García, Fisteus and Brill, 2022).[5]

The maximin support rule is defined for approval ballots as follows.

Definition 5 (Maximin Support Rule)**.** *Given an instance I and a profile $\boldsymbol{A}$ of approval ballots, the* maximin support rule, MaximinSupp, *constructs a budget allocation π, initially empty, iteratively as follows.*

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

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

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

If there exists $p \in P^\star$ such that $c(\pi \cup \{p\}) > b$, the maximin support rule terminates and outputs π. Otherwise, a project $p \in P^\star$ 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 $P^\star$. 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 MaximinSupp 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.

[5] This was observed by Brill, Forster, Lackner, Maly and Peters (2023).

4.4 The Method of Equal Shares

The next rule we introduce is called the *Method of Equal Shares* (formerly known as Rule X). It is similar to SeqPhrag (in its continuous formulation) or MaximinSupp except that agents receive all their money initially.

This rule has been introduced for PB by Peters, Pierczyński and Skowron (2021)[6] for generic cardinal ballots, based on the version for multi-winner voting introduced by Peters and Skowron (2020). We provide their definition below.[7]

Definition 6 (Method of Equal Shares for Cardinal Ballots)**.** *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile $\boldsymbol{A}$ of cardinal ballots, the* method of equal shares, MES, *constructs a budget allocation π, initially empty, iteratively as follows.*

A load $\ell_i : 2^{\mathcal{P}} \to \mathbb{R}_{\geq 0}$, 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:

$$\gamma_i(\pi, \alpha, p) = \min\left({}^{b}\!/_{n} - \ell_i(\pi), \alpha \cdot A_i(p)\right).$$

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 said to be α-affordable, for $\alpha \geq 0$, if

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

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

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 $\alpha^\star$-affordable where $\alpha^\star$ 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.

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

We observe that MES does not necessarily spend the whole budget, *i.e.*, it is not *exhaustive* (see Definition 54). 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 6.1.

The definition of MES can easily be adapted for approval ballots. Note that since approval ballots are just restricted cases of cardinal ones, MES can actually be used as is. However, since the cardinal ballot are meant to represent the voters' utility for the projects, it is more natural to parametrize MES for approval ballots by a satisfaction function (Brill, Forster, Lackner, Maly and Peters, 2023).

Definition 7 (Method of Equal Shares for Approval Ballots)**.** *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile $\boldsymbol{A}$ of approval ballots, the* method of equal shares for the satisfaction function sat, MES[sat], *constructs a budget allocation π, initially empty, iteratively as follows.*

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

[7] For a description of the Method of Equal Shares aimed at non-experts, see equalshares.net, a website maintained by Dominik Peters.

A load $\ell_i : 2^{\mathcal{P}} \to \mathbb{R}_{\geq 0}$, 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:

$$\gamma_i(\pi, \alpha, p) = A_i(p) \cdot \min\left({}^{b}/_{n} - \ell_i(\pi), \alpha \cdot sat(\{p\})\right).$$

This is the amount i would pay to buy project p for a given α. Importantly, i only contribute to p if $A_i(p) = 1$, 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

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

A project is thus α-affordable if, for the given α, all the agents can contribute enough to afford 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 $\alpha^\star$-affordable where $\alpha^\star$ 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 sat is not additive, MES$[sat]$ is still well-defined.

4.5 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 non unit-cost setting as shown 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 Votes* 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.[8]

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

Chapter 5

Fairness in Indivisible Participatory Budgeting

All the ingredients are now in place: We have introduced the basic setting (Chapter 2), discussed how the voters can submit their opinion (Chapter 3), and finally how to take these opinions into account to select budget allocations (Chapter 4). However, defining the rules is only the first step, we still need to assess their respective merits.

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 justified representation (Section 5.1) which will naturally lead us to the idea of the core (Section 5.2). We will then discuss the idea of priceability (Section 5.3). Broadening our perspective, our next focus will be fairness in ordinal PB (Section 5.4), and more generally all the other notions of fairness that have been introduced (Section 5.5). In an attempt to unify everything, we will then draw taxonomies linking the requirements to each others (Section 5.7). We will conclude by discussing fairness in extended models of PB (Section 5.6).

5.1 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, Sánchez-Fernández 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 presents the main results from the literature. Figures 5.1 and 5.2 summarise (among others) the results presented in this section.

5.1.1 Justified Representation with Cardinal Ballots

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

5.1.1.1 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 8 ((α, P)-Cohesive Groups). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile $\boldsymbol{A}$ of cardinal ballots, a non-empty group of agents $N \subseteq \mathcal{N}$ is said to be* (α, P)-cohesive, *for a function $\alpha : \mathcal{P} \to [0, 1]$ and a set of projects $P \subseteq \mathcal{P}$, if the following two conditions are satisfied:*

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

Overall, for any (α, P)-cohesive group of agents $N \subseteq \mathcal{N}$, the following holds: (i) N is 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 $\alpha^{\min}$ and defined for any profile $\boldsymbol{A}$ and subset of agents $N \subseteq \mathcal{N}$ as:

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

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 $(\alpha^{\min}_{N,\boldsymbol{A}}, P)$-cohesive.

5.1.1.2 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.*[1]

Definition 9 (Strong Extended Justified Representation). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ a profile $\boldsymbol{A}$ of cardinal ballots, a budget allocation $\pi \in$ Feas(I) is said to satisfy* strong extended justified representation *(Strong-EJR) if for all $P \subseteq \mathcal{P}$, all $(\alpha^{\min}_{N,\boldsymbol{A}}, P)$-cohesive groups N, 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.

[1]Note 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 $\alpha^{\min}$. The two definitions are however equivalent and we believe this one to be clearer since it requires one less universal quantification.

Definition 10 (Extended Justified Representation)**.** *Given an instance $I = \langle \mathcal{P}, c, b\rangle$ a profile $\boldsymbol{A}$ of cardinal ballots, a budget allocation $\pi \in \textsc{Feas}(I)$ is said to satisfy* extended justified representation *(EJR) if for all $P \subseteq \mathcal{P}$ and all $(\alpha^{\min}_{N,\boldsymbol{A}}, P)$-cohesive groups N, there exists $i \in N$ such that:*

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

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

Theorem 11 (Peters, Pierczyński and Skowron 2021)**.** *For every instance I, there exists a budget allocation $\pi \in \textsc{Feas}(I)$ that satisfies EJR.*

What Peters, Pierczyński and Skowron (2021) actually prove is that a *greedy cohesive rule*[2] always returns a feasible budget allocation that satisfies EJR (it even satisfies *full justified representation*, see Section 5.5.1). 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 12 (Peters, Pierczyński and Skowron 2021)**.** *There is no strongly polynomial time algorithm that computes, given an instance I and a profile $\boldsymbol{A}$ of cardinal ballots, a budget allocation $\pi \in \textsc{Feas}(I)$ satisfying EJR unless* $\texttt{P} = \texttt{NP}$*, 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 $\texttt{coNP}$-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 could 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 of at most one project away from the objective.

Definition 13 (Extended Justified Representation up to One Project)**.** *Given an instance $I = \langle \mathcal{P}, c, b\rangle$ a profile $\boldsymbol{A}$ of cardinal ballots, a budget allocation $\pi \in \textsc{Feas}(I)$ is said to satisfy* extended justified representation up to one project *(EJR-1) if for all $P \subseteq \mathcal{P}$ and all $(\alpha^{\min}_{N,\boldsymbol{A}}, P)$-cohesive groups N, there exists $i \in N$ such that either $\sum_{p \in \pi} A_i(p) \geq \sum_{p \in P} \alpha(p)$, or for some $p^\star \in P \setminus \pi$ we have:*

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

One might be surprised by the strict inequality in the definitions above. It is there for technical reasons: It ensures that EJR and EJR-1 coincide in the unit cost setting when used with approval

[2]The idea behind a greedy cohesive rule is to consider all the cohesive groups, and to greedily select sets of projects P for which there is a suitable (α, P)-cohesive group with "maximum" α. This is a general scheme for procedures as the notion of "suitable cohesive group" differs depending on the goal. Such procedures have notably been defined and studied by Aziz, Lee and Talmon (2018), Peters, Pierczyński and Skowron (2021) and Maly, Rey, Endriss and Lackner (2023).

ballots.[3] It also has interesting consequences in terms of the fairness properties that EJR-1 implies.[4]

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 14 (Peters, Pierczyński and Skowron 2021)**.** *For every instance I and profile $\boldsymbol{A}$ of cardinal ballots,* MES$(I, \boldsymbol{A})$ *satisfies EJR-1.*

5.1.1.3 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 15 (Proportional Justified Representation)**.** *Given an instance $I = \langle \mathcal{P}, c, b\rangle$ a profile $\boldsymbol{A}$ of cardinal ballots, a budget allocation $\pi \in$ Feas(I) is said to satisfy* proportional justified representation *(PJR) if for all $P \subseteq \mathcal{P}$ and all $(\alpha^{\min}_{N,\boldsymbol{A}}, P)$-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 11 and 12 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 $\texttt{P} = \texttt{NP}$. To see why the second point holds, observe that PJR and EJR coincide when there is only a single agent and that Theorem 12 holds for one-agent profiles.

Interestingly, the problem of checking whether a budget allocation satisfies PJR or not is still a 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, Sánchez-Fernández and Skowron, 2018).

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

This last axiom concludes our section on cardinal ballots, we will now focus on approval ballots.

5.1.2 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.

[3] EJR 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.

[4] Notably, 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 $|N|/n \cdot p \geq c(P)$ 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.

5.1.2.1 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 16 (P-cohesive groups)**.** *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile $\boldsymbol{A}$ 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: a group of agents that deserves some satisfaction in the final outcome. When using approval ballots, we will use *(approval-based) satisfaction functions* as introduced in Section 3.4 as measures of satisfaction.

5.1.2.2 Extended Justified Representation for Approval Ballots

Having defined cohesive groups, 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 17 (Strong-EJR for Approval Ballots)**.** *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ a profile $\boldsymbol{A}$ of approval ballots, and a satisfaction function sat, a budget allocation $\pi \in \textsc{Feas}(I)$ 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 quite appealing, but not satisfiable in general.

Proposition 18. *For any satisfaction function sat, there exists an instance I such that no budget allocation $\pi \in \textsc{Feas}(I)$ satisfies Strong-EJR[sat].*[5]

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

Definition 19 (Extended Justified Representation for Approval Ballots)**.** *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ a profile $\boldsymbol{A}$ of approval ballots, and a satisfaction function sat, a budget allocation $\pi \in \textsc{Feas}(I)$ 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 20 (Brill, Forster, Lackner, Maly and Peters 2023)**.** *For every instance I and every satisfaction function sat, there exists a budget allocation $\pi \in \textsc{Feas}(I)$ that satisfies EJR[sat].*

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

[5] We are not aware of this result existing in the literature. The proof is rather simple, it relies on a counter example 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. 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 $b = 2$.

Theorem 21 (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 $\boldsymbol{A}$ of approval ballots, a budget allocation $\pi \in$ Feas(I) satisfying EJR-[sat] unless $\mathtt{P} = \mathtt{NP}$, 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 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 $\mathtt{coNP}$ 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 22 (Extended Justified Representation up to Any Project for Approval Ballots)**.** *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ a profile $\boldsymbol{A}$ of approval ballots, and a satisfaction function sat, a budget allocation $\pi \in$ 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 $p^\star \in P \setminus \pi$ we have $sat_i(\pi \cup \{p^\star\}) > sat_i(P)$.*

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, the outcome of MES$[sat]$ always satisfies EJR-X$[sat]$.

Definition 23 (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:*

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

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

Before moving on to PJR for approval ballots, let us touch on another topic. Fairstein, Vilenchik, Meir and Gal (2022) study the consequences of satisfying EJR on measures of social welfare and representation.[6] Specifically, they provide bounds on the social welfare and representation guarantees of rules satisfying EJR$[sat^{card}]$. In other words, they compare the maxi-

[6] 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.

mally achievable social welfare with respect to sat^{card} and sat^{CC} to the social welfare achieved by rules satisfying EJR.[7]

Theorem 25 (Fairstein, Vilenchik, Meir and Gal 2022)**.** *Let* R *be a PB rule that satisfies EJR[sat^{card}]. Then for any instance $I = \langle \mathcal{P}, c, b \rangle$ and profile $\boldsymbol{A}$ of approval ballots, we have:*

$$\frac{c_{\min}}{n \cdot b} \left\lfloor \frac{b}{c_{\max}} \right\rfloor \leq \frac{\displaystyle\sum_{i \in \mathcal{N}} sat_i^{card}(\mathrm{R}(I, \boldsymbol{A}))}{\displaystyle\max_{\pi \in \textsc{Feas}(I)} \sum_{i \in \mathcal{N}} sat_i^{card}(\pi)} \leq \frac{4}{\sqrt{n}} - \frac{1}{n},$$

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

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

$$\frac{1}{n} \leq \frac{\displaystyle\sum_{i \in \mathcal{N}} sat^{CC}(\mathrm{R}(I, \boldsymbol{A}))}{\displaystyle\max_{\pi \in \textsc{Feas}(I)} \sum_{i \in \mathcal{N}} sat^{CC}(\pi)} \leq \frac{1}{n-1},$$

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

5.1.2.3 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 26 (Proportional Justified Representation for Approval Ballots)**.** *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ a profile $\boldsymbol{A}$ of approval ballots, and a satisfaction function sat, a budget allocation $\pi \in \textsc{Feas}(I)$ is said to satisfy* proportional justified representation for *sat (PJR[sat]) if for all $P \subseteq \mathcal{P}$ and all P-cohesive groups N, we have:*

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

Similar adaptions of PJR to the PB setting have also been studied. PJR[sat^{cost}] is equivalent to the BPJR-L property introduced by Aziz, Lee and Talmon (2018).[8] Aziz, Lee and Talmon (2018) also defined variants of (B)PJR based on the relative budget, which will be discussed in Section 5.5.2. Finally, PJR[sat^{card}] has been introduced by Los, Christoff and Grossi (2022).

For now, let us focus on PJR[sat]. It should be clear that for any satisfaction function sat, EJR[sat] implies PJR[sat]. Thus, for any instance I and profile $\boldsymbol{A}$ of approval ballots, there exists a budget allocation satisfying PJR[sat], however for a large class of satisfaction function, it cannot be computed in polynomial time (see Theorem 21 for the exact condition on the satisfaction function). Finally checking that PJR[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.

[7] Let us also mention that Lackner and Skowron (2020) studied the same questions in the multiwinner voting setting.

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

Definition 27 (Proportional Justified Representation up to Any Project for Approval Ballots)**.** *Given an instance $I = \langle \mathcal{P}, c, b\rangle$ a profile $\boldsymbol{A}$ of approval ballots, and a satisfaction function sat, a budget allocation $\pi \in \textsc{Feas}(I)$ is said to satisfy* proportional justified representation up to any project for sat *(PJR-X[sat]) if for all $P \subseteq \mathcal{P}$ and all P-cohesive groups N and any $p^\star \in P \setminus \pi$, we have:*

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

Remember that we know for any DNS function sat (Definition 23), that EJR-X[sat] can be satisfied (Theorem 24). Since PJR-X[sat] is implied by EJR-X[sat], this result also applies to PJR-X[sat]. Brill, Forster, Lackner, Maly and Peters (2023) actually prove something stronger: PJR-X[sat] can be satisfied *simultaneously* for every DNS function sat.

Theorem 28 (Brill, Forster, Lackner, Maly and Peters 2023)**.** *Let I be an instance I and $\boldsymbol{A}$ a profile. Then,* $\textsc{SeqPhrag}(I, \boldsymbol{A})$*,* $\textsc{MaximinSupp}(I, \boldsymbol{A})$ *and* MES$[sat^{card}](I, \boldsymbol{A})$ *satisfy PJR-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 5.3).

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-1[sat^{card}]. It also generalises the result of Aziz, Lee and Talmon (2018) that MaximinSupp satisfies a property called Local-BPJR-L[sat^{cost}] as explained below.

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

Before Brill, Forster, Lackner, Maly and Peters (2023) introduced their definition of PJR parameterised by a satisfacion function, Aziz, Lee and Talmon (2018) defined PJR for PB with approval ballots. As we have mentioned before, they introduced an axiom called BPJR-L—that is equivalent to PJR[sat^{cost}]—and proved that budget allocations satisfying it could not be found in polynomial time (unless $\mathtt{P} = \mathtt{NP}$). Due to this observation, they introduced Local-BPJR-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-BPJR-L[sat^{cost}].

Definition 29 (Local Budget Proportional Justified Representation for the Budget Limit)**.** *Given an instance $I = \langle \mathcal{P}, c, b\rangle$ a profile $\boldsymbol{A}$ of approval ballots, and a satisfaction function sat, a budget allocation $\pi \in \textsc{Feas}(I)$ is said to satisfy* Local-BPJR-L[sat] *if for all $P \subseteq \mathcal{P}$ and all P-cohesive groups N, it is the case that for every $P^\star \subseteq P$ such that $\{p \in \pi \mid \exists i \in N, A_i(p) = 1\} \subsetneq P^\star$ we have:*

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

One of the main result 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}] 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 $c\left(\bigcup_{i \in N}\{p \in \pi \mid A_i(p) = 1\}\right) \geq \ell$. 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.

5.2 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.

5.2.1 The Core with Cardinal Ballots

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

Definition 30 (The Core of PB with Cardinal Ballots)**.** *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile $\boldsymbol{A}$ of cardinal ballots, a budget allocation $\pi \in$ 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 $i^\star \in N$ with:*

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

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 (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 no feasible budget allocation is in the core. They strengthen a first counter-example provided by Fain, Munagala and Shah (2018) without the unit-cost assumption.[9]

Proposition 31 (Peters, Pierczyński and Skowron 2021)**.** *There exists an instance I with unit costs and profile $\boldsymbol{A}$ of cardinal ballots such that no budget allocation $\pi \in$ Feas(I) is in the core, even if for every agent $i \in \mathcal{N}$ and project $p \in \mathcal{P}$ we have $A_i(p) \in \{0, 1, 2\}$.*

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

5.2.2 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 32 (The α-sat Approximate Core of PB with Cardinal Ballots)**.** *Given an instance $I = \langle \mathcal{P}, c, b \rangle$, a profile $\boldsymbol{A}$ of cardinal ballots, and a scalar $\alpha \geq 1$, a budget allocation $\pi \in \textsc{Feas}(I)$ 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 $i^\star \in N$ and a project $p^\star \in \mathcal{P}$ with:*

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

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.

Theorem 33 (Peters, Pierczyński and Skowron 2021)**.** *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile $\boldsymbol{A}$ of cardinal ballots, let $u_{\max}$ and $u_{\min}$ be the highest and lowest possible satisfaction of a voter, defined as:*

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

Then, $\mathrm{MES}(I, \boldsymbol{A})$ *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 $\mathcal{O}(\log(|\textsc{Feas}(I)|))$-sat approximate core is always non empty for any instance I. Moreover, it also implies 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, however it is unknown if the corresponding budget allocation can be computed in polynomial-time.

Theorem 34 (Munagala, Shen, Wang and Wang 2022)**.** *For every instance I and profile $\boldsymbol{A}$ 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 it also allows Munagala, Shen, Wang and Wang (2022) to 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 35 (The α-entitlement approximate core of PB with Cardinal Ballots)**.** *Given an instance $I = \langle \mathcal{P}, c, b \rangle$, a profile $\boldsymbol{A}$ of cardinal ballots, and a scalar $\alpha \geq 1$, a budget allocation $\pi \in \textsc{Feas}(I)$ 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 $i^\star \in N$ with:*

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

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

Theorem 36 (Jiang, Munagala and Wang 2020)**.** *For every instance I and profile $\boldsymbol{A}$ 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 $\boldsymbol{A}$ of cardinal ballots and a budget allocation $\pi \in \textsc{Feas}(I)$, 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 a logarithmic approximation algorithm.

5.2.3 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 37 (The Core of PB with Approval Ballots)**.** *Given an instance $I = \langle \mathcal{P}, c, b \rangle$, a profile $\boldsymbol{A}$ of approval ballots and a satisfaction function sat, a budget allocation $\pi \in \textsc{Feas}(I)$ 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 $i^\star \in N$ with:*

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

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

5.3 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.[10]

[10] Note that we changed the terminology to avoid using the terms "budget" and "price" that can be confused with the basic elements of an instance. This typically avoids sentences such as "π is priceable for a budget $B \geq b$".

Definition 38 (Priceability for Cardinal Ballots)**.** *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile $\boldsymbol{A}$ 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, $\gamma_i : \mathcal{P} \to [0, \alpha]$ 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 project they derive satisfaction from.*
- ***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 contribution.*
- ***C3**: $\sum_{p \in \mathcal{P}} \gamma_i(p) \leq \alpha$ for all $i \in \mathcal{N}$: No agent contributes more than 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 contribution to be funded.*
- ***C5**: $\sum_{i \in \mathcal{N} \mid A_i(p) > 0} \left(\alpha - \sum_{p \in \mathcal{P}} \gamma_i(p) \right) \leq c(p)$ for all $p \in \mathcal{P} \setminus \pi$: No group of agents with positive utility for a project p has more than $c(p)$ money left.*

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 do some tie-breaking 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 the latter always returns priceable budget allocations. Maybe more surprisingly, it also is the case of sequential Phragmén and the maximin support rules.

Proposition 39 (Peters, Pierczyński and Skowron 2021)**.** *For every instance I and profile $\boldsymbol{A}$ of cardinal ballots,* MES$(I, \boldsymbol{A})$ *is priceable.*

Proposition 40 (Los, Christoff and Grossi 2022)**.** *For every instance I and profile $\boldsymbol{A}$ of approval ballots,* SeqPhrag$(I, \boldsymbol{A})$ *is priceable.*

Proposition 41 (Brill, Forster, Lackner, Maly and Peters 2023)**.** *For every instance I and profile $\boldsymbol{A}$ of approval ballots,* MaximinSupp$(I, \boldsymbol{A})$ *is priceable.*

Priceability and PJR. In the context of multi-winner voting, links have been drawn between PJR and pricebility (Peters and Skowron, 2020). Brill, Forster, Lackner, Maly and Peters (2023) extend this result for PB with approval ballots. They show that priceability implies PJR-X$[sat^{cost}]$. More importantly, they show that a stronger notion of priceability implies PJR-X$[sat]$ for all DNS functions sat (see Definition 23).

Theorem 42 (Brill, Forster, Lackner, Maly and Peters 2023)**.** *For every instance $I = \langle \mathcal{P}, c, b \rangle$ and profile $\boldsymbol{A}$ of approval ballots, consider a budget allocation $\pi \in$* 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 extra 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 moving 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 MES[sat^{card}], SeqPhrag and MaximinSupp provide budget allocations that are priceable for their extended notion of priceability.

5.4 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 agents submit ordinal ballots.

Definition 43 (Solid Coalition)**.** *Let $I = \langle \mathcal{P}, c, b \rangle$ be an instance and $\boldsymbol{A} = (\succsim_i)_{i \in \mathcal{N}}$ a profile of weak ordinal ballots. Given a subset of projects $P \subseteq \mathcal{P}$, a group of voters $N \subseteq \mathcal{P}$ is a P-*solid coalition *if for all voter $i \in N$ and project $p \in P$, 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* (Dummett, 1984). Before defining them, we introduce a new notation. Interpret a weak order $\succsim$ over $\mathcal{P}$ as a vector of indifference classes $\succsim = (P_1, P_2, \ldots)$ such that all projects in P_j are preferred to the ones in $P_{j+1} \cup P_{j+2} \cup \cdots$. Then, let $top(\succsim, k)$, for $k \in \mathbb{N}$ be defined as $top(\succsim, k) = P_1 \cup \cdots \cup P_{j^\star} \cup P_{j^\star+1}$ where $j^\star \in \mathbb{N}_{\geq 0}$ is the largest number such that $|\bigcup_{j=1}^{j^\star} P_j| < k$.

Definition 44 (Comparative Proportionality for Solid Coalitions)**.** *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and profile $\boldsymbol{A} = (\succsim_i)_{i \in \mathcal{N}}$ of weak ordinal ballots, a budget allocation $\pi \in \textsc{Feas}(I)$ 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:*

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

Definition 45 (Inclusion Proportionality for Solid Coalitions)**.** *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile $\boldsymbol{A} = (\succsim_i)_{i \in \mathcal{N}}$ of weak ordinal ballots, a budget allocation $\pi \in \textsc{Feas}(I)$ 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 $p^\star \in P \setminus \{p \in \pi \mid \exists i \in N \text{ such that } p \in top(\succsim_i, |P|)\}$ such that:*

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

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 46 (Aziz and Lee 2021)**.** *There exist an instance I and a profile $\boldsymbol{A}$ of weak ordinal ballots such that no $\pi \in \textsc{Feas}(I)$ satisfies CPSC.*

For every instance I and a profile $\boldsymbol{A}$ of weak ordinal ballots there exists $\pi \in \textsc{Feas}(I)$ 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.

5.5 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 receive less attention, we will go a bit faster on them.

5.5.1 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 47 (Full Justified Representation for Cardinal Ballots). *Let $I = \langle \mathcal{P}, c, b \rangle$ be an instance and $\boldsymbol{A}$ a profile of cardinal ballots. 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 $\boldsymbol{A}$, a budget allocation $\pi \in \textsc{Feas}(I)$ satisfies full justified representation *(FJR) if for all $P \subseteq \mathcal{P}$, all $\beta \in \mathbb{R}$ and all weakly (β, P)-cohesive group N, there exists $i \in N$ 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 48 (Peters, Pierczyński and Skowron 2021). *For any instance I and profile $\boldsymbol{A}$ of cardinal ballots, there exists a budget allocation $\pi \in \textsc{Feas}(I)$ 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 49 (Full Justified Representation for Approval Ballots). *Let $I = \langle \mathcal{P}, c, b \rangle$ be an instance, $\boldsymbol{A}$ a profile of approval ballots and sat a satisfaction function. A budget allocation $\pi \in \textsc{Feas}(I)$ satisfies* full justified representation for sat *(FJR[sat]) if for every groups of voters $N \subseteq \mathcal{N}$ and subset of project $P \subseteq \mathcal{P}$ such that $|N|/n \cdot b \geq c(P)$, there exists $i \in N$ for whom:*

$$sat_i(\pi) \geq sat_i(P).$$

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.

5.5.2 Variants with Relative Budget

Most of the proportionality requirements we introduce 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 adaptions 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 8 and 16—this does not lead to any division by 0).

5.5.3 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 are laminar proportional in the multi-winner setting (namely MES and SeqPhrag) cease to be on PB instances.

5.5.4 Proportionality for Solid Coalitions

In section 5.4 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 give these definitions below. Observe that they are both closely related to PJR.

Definition 50 (CPSC with Approval Ballots)**.** *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile $\boldsymbol{A}$ of approval ballots, a budget allocation $\pi \in \textsc{Feas}(I)$ is said to satisfy* CPSC *if the following two conditions hold:*

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

Definition 51 (IPSC with Approval Ballots)**.** *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile $\boldsymbol{A}$ of approval ballots, a budget allocation $\pi \in \textsc{Feas}(I)$ 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 PJR-X[sat^{cost}]. One can actually prove that IPSC implies PJR-X[sat^{cost}]. Indeed, if a budget allocation π fails PJR-X[sat^{cost}], then the P-cohesive N witnessing this violation would also be a witness of the violation 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.

5.5.5 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 proportional representation axioms for this setting. We present here what they call *proportional representation*.

Definition 52 (Proportional Representation with Cumulative Ballots)**.** *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile $\boldsymbol{A}$ of cumulative ballots, a budget allocation $\pi \in$ Feas(I) is said to satisfy* proportional representation *if for every $\ell \in \{1, \ldots, b\}$, 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 $p \in \mathcal{P} \setminus P$, 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.

5.5.6 Equality of Resources

In the context of PB with approval ballots, one needs to go through the concept of satisfaction functions to define proportionality requirements. Based on the observation that no satisfaction function can ever be satisfactory, Maly, Rey, Endriss and Lackner (2023) suggest to define fairness criteria not in terms of satisfaction but in terms of the resources spent on an agent, the so called *equality of resources.* They use the concept of *share* (Lackner, Maly and Rey, 2021) to measure the amount of resources spent on an agent and aim at providing every agents with their *fair share.*

Definition 53 (Fair Share)**.** *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile $\boldsymbol{A}$ of approval, $\pi \in$ Feas(I) is said to satisfy* fair share *if for every agent $i \in \mathcal{N}$ we have:*

$$\sum_{p \in \pi} A_i(p) \cdot \frac{c(p)}{|\{i' \in \mathcal{N} \mid A_{i'}(p) = 1\}|} \geq \min \left\{ \frac{b}{n}, \sum_{\substack{p \in \mathcal{P} \\ A_i(p)=1}} \frac{c(p)}{|\{i' \in \mathcal{N} \mid A_{i'}(p) = 1\}|} \right\}.$$

This requirement is not satisfiable in general, which motivated Maly, Rey, Endriss and Lackner (2023) to introduce several weakenings, either based on direct relaxations ("up-to-one project" or "local" variants), or on cohesive groups. Most notably, they prove that a variant of MES provides good fair share guarantees both theoretically and empirically.

5.6 Fairness in Extended Settings

We now mention some papers that have studied fairness in PB beyond the standard model. This section overlaps in some way with Chapter 8 though we only focus on fairness requirements here and only discuss the work briefly.

- 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 PJR[sat^{cost}] (or BPJR-L) and EJR.

- Lackner, Maly and Rey (2021) introduce a fairness theory for long-term PB where several instances are considered. They introduce several fairness concepts for their setting and study under which conditions they can be satisfied.

- 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.

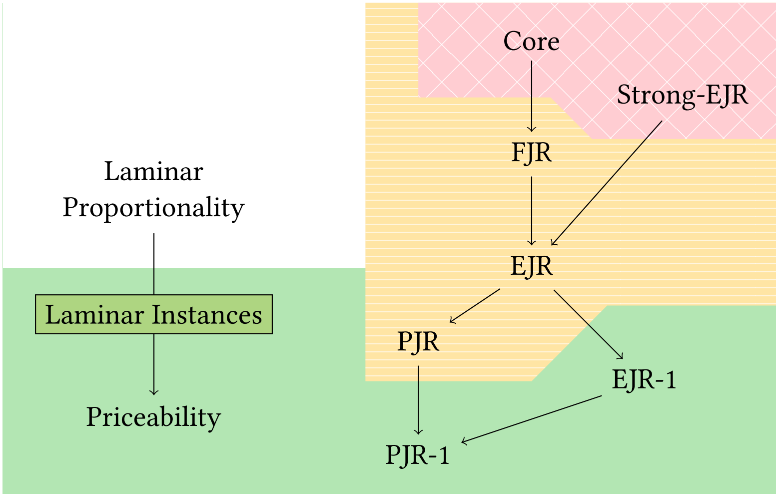

These concepts cannot always be satisfied.
These concepts can always be satisfied, however finding a suitable budget allocation cannot be done in polynomial time unless $\mathtt{P} = \mathtt{NP}$.
These concepts can always be satisfied, and a suitable budget allocation can be found in polynomial time.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

Figure 5.1: 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 arrow between the 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 5.1.

5.7 Taxonomies of Proportionality in PB

Throughout this chapter 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 5.1. Figure 5.2 presents the taxonomy for approval ballots. All the details are available on the figures. We also summarise which rule satisfy which axioms, in Table 5.1 for cardinal ballots, and in Table 5.2 for approval ones.

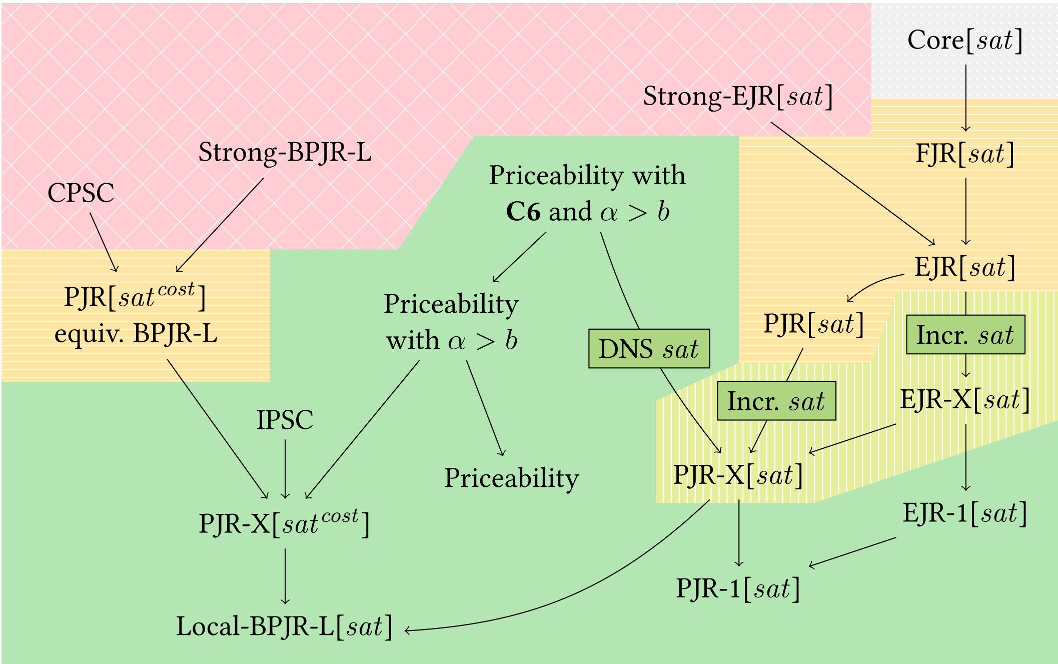

These concepts cannot always be satisfied.
These concepts can always be satisfied, however finding a suitable budget allocation cannot be done in polynomial time unless $\mathtt{P} = \mathtt{NP}$.
These concepts can always be satisfied, and a suitable budget allocation can be found in polynomial time when sat is a DNS function.
These concepts can always be satisfied, and a suitable budget allocation can be found in polynomial time when sat is additive (for the concepts depending on sat).
It is unknown whether the core can always be satisfied or not.

Incr. sat**:** the link only applies for satisfaction functions that are strictly increasing, *i.e.*, such that for all $P \subseteq \mathcal{P}$ and $P' \subsetneq P$, we have $sat(P') < sat(P)$.
DNS sat**:** the link only applies for DNS function, see Definition 23.
PJR[sat^{cost}**] equiv. BPJR-L:** these two concepts are equivalent.
Priceability with $\alpha > b$**:** priceable for a price system $\langle \alpha, (\gamma_i)_{i \in \mathcal{N}} \rangle$ where $\alpha > b$.
Priceability with C6 and $\alpha > b$**:** see Theorem 42.

Figure 5.2: Taxonomy of the proportionality requirements for PB with approval ballots where sat is an arbitrary satisfaction function. An arrow between two concepts means that any budget allocation satisfying one, also satisfies the other. Some arrows are only valid for some satisfaction functions, the conditions are indicate on the arrow. All missing arrows are known to be missing.

The links between EJR, PJR, Local-BPJR-L and priceability concepts are due to Brill, Forster, Lackner, Maly and Peters (2023). The link from Strong-BPJR-L and BPJR-L is due to Aziz, Lee and Talmon (2018). The link between CPSC and PJR[sat^{cost}] is due to Aziz and Lee (2021). The one between IPSC and PJR-X[sat^{cost}] has never been published. The absence of arrows between the core, EJR and priceability is due to Los, Christoff and Grossi (2022). The link between FJR and EJR is due to Peters, Pierczyński and Skowron (2021). For the satisfiability of the concepts, see Table 5.2.

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 5.1: Rules satisfying each of the fairness property we introduced for generic 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 5.2: Rules satisfying each of the fairness property we introduced for approval ballots. sat is an arbitrary satisfaction function.

Chapter 6

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 that have been introduced.

Our analysis will start with a discussion around exhaustiveness (Section 6.1) and a presentation of the monotonicity axioms that have been introduced for PB (Section 6.2). From there, we will move on to axioms relating to strategic behaviour of the agents (Section 6.3). We will conclude this section by our usual discussion of the concepts that exist in the literature but which do not fit in the previous sections (Section 6.4).

6.1 Exhaustiveness

Let us start with *exhaustiveness*, an efficiency requirement that states that the budget should not be underused. 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 an axiom called *budget monotonicity* 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).

Let us first introduce the idea of exhaustiveness.

Definition 54 (Exhaustiveness)**.** *Given an instance $I = \langle \mathcal{P}, c, b \rangle$, a feasible budget allocation $\pi \in \textsc{Feas}(I)$ is said to be* exhaustive *if there are no project $p \in \mathcal{P} \setminus \pi$ such that $c(\pi \cup \{p\}) \leq b$.*

Table 6.1 summarizes which of the usual rules satisfy exhaustiveness. The results for the welfare maximizing and greedy rules are straightforward. Interestingly, the fact that SeqPhrag,

	Exhaustiveness	
	General Instances	Unit-Cost Instances
MaxCard	✓	✓
GreedCard	✓	✓
MaxCost	✓	✓
GreedCost	✓	✓
SeqPhrag	✗	✓
MaximinSupp	✗	✓
MES	✗	✗

Table 6.1: Satisfaction of exhaustiveness for different rules.

MaximinSupp and MES fail exhaustiveness is due the fact that they are priceable. Indeed, the two requirements are incompatible.

Proposition 55 (Peters, Pierczyński and Skowron 2021)**.** *There exists an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile $\boldsymbol{A}$, 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 again 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 6.2), 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 know 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.

6.2 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 variation of the instance. We will define these axioms in the following and present what is known about them.

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

Definition 56 (Discount monotonicity)**.** *A PB rule* R *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 $p^\star \in \mathcal{P}$, we have $c(p^\star) > c'(p^\star)$, and $c(p) = c'(p)$ for all $p \in \mathcal{P} \setminus \{p^\star\}$, it is the case that $p^\star \in \mathrm{R}(I, \boldsymbol{A})$ implies $p^\star \in \mathrm{R}(I', \boldsymbol{A})$ for all profiles $\boldsymbol{A}$.*

Thus, a rule is discount monotonic if whenever the price of a selected project p decreases, the rule would still 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 57 (Limit monotonicity)**.** *A PB rule* R *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* $\mathrm{R}(I, \boldsymbol{A}) \subseteq \mathrm{R}(I', \boldsymbol{A})$ *for all profiles* $\boldsymbol{A}$*.*

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

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

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

Definition 58 (Splitting monotonicity)**.** *A PB rule* R *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* $\boldsymbol{A}$ *and* $\boldsymbol{A}'$ *and any project* $p \in \mathrm{R}(I, \boldsymbol{A})$ *such that* I' *and* $\boldsymbol{A}'$ *are the result of splitting project* p *into a subset of projects* P *given* I *and* $\boldsymbol{A}$*, it is the case that* $\mathrm{R}(I', \boldsymbol{A}') \cap P \neq \emptyset$*.*

Definition 59 (Merging monotonicity)**.** *A PB rule* R *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‘* $\boldsymbol{A}$ *and* $\boldsymbol{A}'$*, and any subset of projects* $P \subseteq \mathrm{R}(I, \boldsymbol{A})$ *such that* I' *and* $\boldsymbol{A}'$ *are the result of merging project set* P *into project* p *given* I *and* $\boldsymbol{A}$*, it is the case that* $p \in \mathrm{R}(I', \boldsymbol{A}')$*.*

These two axioms thus impose 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 6.2 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).

	Monotonicity			
	Limit	Discount	Splitting	Merging
MaxCard	✗	✓	✓	✗
GreedCard	✗	✓	✓	✗
MaxCost	✗	✗	✓	✓
GreedCost	✗	✗	✗	✓
MES	✗			

Table 6.2: 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 montonicity as it already did on 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—that can return more than one budget allocation—in two different ways general ways. For a given irresolute rule R:

- 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 $\boldsymbol{A}$, and for every budget allocation $\pi \in \mathrm{R}(I, \boldsymbol{A})$ that satisfy a specific pre-condition, it must be the case that for every suitable I' and $\boldsymbol{A}'$, there exist a budget allocation $\pi' \in \mathrm{R}(I', \boldsymbol{A}')$ satisfying the specific post-condition;
- Rey, Endriss and de Haan (2020) extend them in a *universal* fashion: for a given instance I and profile $\boldsymbol{A}$ if every budget allocation $\pi \in \mathrm{R}(I, \boldsymbol{A})$ satisfies a specific pre-condition, it must be the case that for every suitable I' and $\boldsymbol{A}'$ and every $\pi' \in \mathrm{R}(I', \boldsymbol{A}')$, the required post-condition is satisfied.

6.3 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.[1]

Definition 60 (Strategy-Proofness for Cardinal Ballots). *A PB rule* R *is said to be* strategy-proof *if for every instance I and profile $\boldsymbol{A}$ of cardinal ballots, and for every agent $i \in \mathcal{N}$, there exists no cardinal ballot A_i' such that for the profile $\boldsymbol{A}' = (A_1, \ldots, A_{i-1}, A_i', A_{i+1}, \ldots, A_n)$ we have:*

$$\sum_{p \in \mathrm{R}(I, \boldsymbol{A}')} A_i(p) > \sum_{p \in \mathrm{R}(I, \boldsymbol{A})} A_i(p).$$

[1]Note 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.

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

Definition 61 (Strategy-Proofness for Approval Ballots)**.** *Given a satisfaction function sat, a PB rule* R *is said to be* strategy-proof for sat *if for every instance I and profile $\boldsymbol{A}$ of approval ballots, for every agent $i \in \mathcal{N}$, there exists no approval ballot A'_i such that for the profile $\boldsymbol{A}' = (A_1, \ldots, A_{i-1}, A'_i, A_{i+1}, \ldots, A_n)$ we have:*

$$sat(\mathrm{R}(I, \boldsymbol{A}') \cap A_i) > sat(\mathrm{R}(I, \boldsymbol{A}) \cap A_i).$$

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 62 (Peters 2018)**.** *A PB rule* R *is said to be* weakly-proportional on unit-cost instances *if for every unit-cost instance I and profile $\boldsymbol{A}$ 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 ($\boldsymbol{A}$ 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 $p \in \mathrm{R}(I, \boldsymbol{A})$.*

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

Note 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 additional axiom.

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 5.1. 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 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 63 (Approximate Strategy-Proofness for Approval Ballots)**.** *Given a satisfaction function sat, a PB rule* R *is said to be* approximately strategy-proof for sat *if for every instance I and profile $\boldsymbol{A}$ of approval ballots, for every agent $i \in \mathcal{N}$, there exists no approval ballot A'_i such that for the profile $\boldsymbol{A}' = (A_1, \ldots, A_{i-1}, A'_i, A_{i+1}, \ldots, A_n)$, for all $p \in \mathcal{P}$ we have:*

$$sat(\mathrm{R}(I, \boldsymbol{A}') \cap A_i) > sat((\mathrm{R}(I, \boldsymbol{A}) \cap A_i) \cup \{p\}).$$

Proposition 64 (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.[2] 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 ballot are knapsack ballots.

Example 65. *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 monotonic[3] and for every project $p \in \mathcal{P}$, agent 3 prefers π' over $\pi \cup \{p\}$.*

6.4 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 for now.

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).

[2] Let 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 $\boldsymbol{A}$ of approval ballots in the same manner to obtain a profile $\boldsymbol{A}'$ of approval ballots. Now, it is clear that the approval scores of the projects in $\boldsymbol{A}'$ are equal to those of the projects in $\mathcal{P}$ they come from in $\boldsymbol{A}$. 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 $\textsc{GreedCost}(I', \boldsymbol{A}')$ 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 $\textsc{GreedCost}(I', \boldsymbol{A}')$. We thus have $\textsc{GreedCost}(I, \boldsymbol{A}) \cup \{p\} = \pi'$. 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}.

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

Chapter 7

Algorithmic Approach

Another large part of the literature focuses on the algorithmic aspects of PB. This usually concerns computing the outcome 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 7.1), then on the complexity of welfare maximisation (Section 7.2), and finally on the other algorithmic problems that have been studied (Section 7.3).

7.1 Outcome Determination of Standard PB Rules

The main focus of the computational 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 R.

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

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 formulate. 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 actual definition of GreedCard, GreedCost and SeqPhrag should make it somewhat obvious that computing their outcome 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.

7.2 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 $\textsc{Util}$-SW so that the reader does not need to get back to Section 4.1.

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

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

  This is the most standard definition of social welfare simply considering the sum of the satisfactions of the individuals. A budget allocation maximising $\textsc{Util}$-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 $\boldsymbol{A}$ of cardinal ballots, the *Chamberlin-Courant social welfare* achieved by a budget allocation π is defined as:

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

  The Chamberlin-Courant social welfare assumes that agents only consider one projects from each budget allocation, the one that leads to the highest satisfaction. Maximising CC-SW corresponds thus to aim 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 $\boldsymbol{A}$ of cardinal ballots, the *egalitarian social welfare* achieved by a budget allocation π is defined as:

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

  The egalitarian social welfare assumes that the welfare of a society is the satisfaction of its most dissatisfied member. Maximising $\textsc{Egal}$-SW hence means maximising the satisfaction of the worst off voter.

  $\textsc{Egal}$-SW is studied by Sreedurga, Bhardwaj and Narahari (2022) under the name *maximin PB.*

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

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

The Nash social welfare measures 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 7.1.

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 Chapter 8 for more details. Hershkowitz, Kahng, Peters and Procaccia (2021) studied welfare maximisation in a model in which projects are grouped into district. 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-resources 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).

7.3 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 projects, 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.

Util-SW	Weakly NP-complete	▶ Even with one voter
	Pseudopoly. solvable	
	Poly. solvable	▶ With approval ballots and sat^{card}
Nash-SW	Strongly NP-complete	▶ 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 $p \in \mathcal{P}$
	W[1]-hard	▶ 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 $\max_{i \in \mathcal{N}} \lvert\{\sum_{p \in \pi} A_i(p) \mid \pi \in \textsc{Feas}(I)\}\rvert$
CC-SW	Pseudopoly. solvable	▶ For single-peaked and single-crossing profiles
	Strongly NP-complete	▶ Even for binary valuations, i.e., 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
Egal-SW	Strongly NP-complete	▶ Even if $A_i(p) \in \{0, c(p)\}$ for all $i \in \mathcal{N}$ and $p \in \mathcal{P}$
	Pseudopoly. 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 $\frac{\max_{p \in \mathcal{P}} c(p)}{GCD\{c(p) \mid p \in \mathcal{P}\}}$ is constant.

Table 7.1: Computational complexity of the decision problem corresponding to the maximisation of different types of social welfare. 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 $\boldsymbol{A}$ of cardinal ballots, and $x \in \mathbb{Q}_{\geq 0}$, is there a budget allocation $\pi \in \textsc{Feas}(I)$ such that $\mathrm{SW}(I, \boldsymbol{A}, \pi) \geq x$?

Statements for Util-SW follow immediately from the literature on the knapsack problem (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.

Note that $A_i(p) \in \{0, 1\}$ for all $i \in \mathcal{N}$ and $p \in \mathcal{P}$ simulates approval ballots with the satisfaction function sat^{card}, and $A_i(p) \in \{0, c(p)\}$ for all $i \in \mathcal{N}$ and $p \in \mathcal{P}$ simulates approval ballots used with sat^{cost}.

Chapter 8

Variations and Extensions of the Standard Model

The literature we reviewed so far studied what could be called the standard model of PB. Beyond that, a myriad of variations of the model have been introduced. In the following we delve into these variations, in an order based on their (assumed) end-goal. We will first look into variations of the standard model that aim at capturing real-life PB processes more accurately (Section 8.1). Our focus will then shift to models that propose extensions of the standard model (Section 8.2).

8.1 Towards More Accurate Models of PB

A large chunk of the variations of the standard model have been introduced with the aim of better capturing real-life PB processes. Among others, the repetitive aspect of PB, its multistage implementation, and its geographical constraints have been studied.

8.1.1 End-to-End Model for PB

We start with the integration of the several stages of a PB process in the analysis. As detailed in the introduction already, a typical PB process has several stages including two during which the citizens are consulted: for the exploration of the projects to consider, and for the vote on which of the selected proposals should actually be implemented. This two stage model has been formalised and studied by Rey, Endriss and de Haan (2021). They focus on two specific aspects: shortlisting rules, and strategic behaviours in this integrated model.

Rey, Endriss and de Haan (2021) approach the first stage of the process—the *shortlisting stage*, during which agents submit proposals for projects that could be considered throughout the process—as a multi-winner election for which there is no specific requirements regarding the size, or cost, of the outcome. They define and analyse several *shortlisting rules* that could be used to create the shortlist based on different ideas: diversity in the voters represented, diversity in the essence of the proposals, and limited size of the shortlist to reduce the cognitive burden induced by the voting stage.

In addition to the first stage, Rey, Endriss and de Haan (2021) also study their end-to-end model in its entirety. They focus on its strategic aspects. They try to answer questions of the type: can an agent improve their satisfaction with the final outcome by not proposing a project in the first stage? Their findings indicate that, unsurprisingly, it is difficult to prevent strategic behaviours.

8.1.2 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, and to some extent in Paris (City of Paris, 2020)[1]. 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 maximises 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.

8.1.3 Temporal Aspects of PB

PB processes are rarely single-shot instances, they often span several years, one PB process being organised each year. Based on that insight, Lackner, Maly and Rey (2021) introduced a model for long-term PB based on the *perpetual voting* framework (Lackner, 2020).

In their work, Lackner, Maly and Rey (2021) introduce what they call a fairness theory for long-term PB. They assume that agents are partitioned into types and they try to achieve fairness for the types over time. They study three fairness requirements based on sat^{cost}, $relsat$ and $share$: Enforcing that all types enjoy the same welfare, that all types converge towards equal welfare if the instance would be infinite or that the welfare across types is distributed optimally (according to the Gini-coefficient). Each of these fairness concepts are analysed in terms of whether they are satisfiable or not. Their findings suggest that it is difficult to provide such fairness guarantees.

8.2 Enriching the Standard Model

The works we have presented above aimed at capturing real-life PB processes more accurately. In the following, we will review works that aim at enriching the standard model by proposing extension of the model that could improve the PB process, but are, to the best of our knowledge, currently not widely implemented in practice.

8.2.1 Additional Distributional Constraints

We first focus on a strand of the literature that deals with incorporating additional constraints to 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. We will see several examples in the following.

[1] In Paris, the PB process combines local and global aspects: voters can vote on the projects for their district together with some Paris-wide projects.

Developing a very general framework for this task, Rey, Endriss and de Haan (2020) demonstrate how to encode PB problems into *judgment aggregation*, a very expressive framework for constrained aggregation (Endriss, 2016). Their framework allows for the addition of any additional constraint that can be expressed in propositional logic. They study the computational cost of such an approach, and show that as long as the constraints can be efficiently encoded in some compact logical representations, the computational overhead is not too large. They also provide an axiomatic analysis (following Talmon and Faliszewski, 2019) of some rules that can be used in this context.

A similar general approach was also considered by Fain, Munagala and Shah (2018) though in a context more general than PB. They provide a framework of public decision making with matroid, matching, and packing constraints, allowing for great flexibility on what can be modelled. Note that packing constraints correspond to what we call budget constraint.

In addition to this, several papers focus on specific constraints that can be implemented in PB.

- **Dependency constraints:** Rey, Endriss and de Haan (2020) study how to include dependency constraints in their framework described above. By dependencies, they mean that the implementation of some projects is dependent on the status of some others.

- **Categorical constraints.** These constraints model the idea that projects are grouped into categories and that additional constraints apply as to which of the projects can be selected within each category.

  Still within their framework, Rey, Endriss and de Haan (2020) introduce *quota constraints* that enforce some lower and upper quota to be satisfied for each category. They provide two example of such quotas: on the number of selected projects from a category, or on the total cost.

  Jain, Sornat, Talmon and Zehavi (2021) also study what Rey, Endriss and de Haan (2020) called cost quota constraint, and what they refer to as *PB with project groups*. 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 approximate 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 a 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 both Rey, Endriss and de Haan (2020) and 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 finally 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).

8.2.2 Interaction Between Projects

An 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.

In addition to their conceptual contribution, Jain, Sornat and Talmon (2020) present a computational analysis of welfare maximising problems in this setting. They provide a mixture of hardness results and (fixed parameter) efficient algorithms. 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.

8.2.3 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 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 modeling 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. In their framework developed to include additional constraints in PB (see above), Rey, Endriss and de Haan (2020) assume that the costs are expressed over several dimensions. More interestingly, 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 the social welfare in this setting.

8.2.4 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 provide some fairness guarantees here.

8.2.5 PB with Endogenous Funding

The standard PB model assumes that the budget is provided by the organising entity (a municipality for instance). Several authors have proposed different 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 contribution to specific projects, thus reducing the amount of public money needed to select the 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 winner 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 the same model except that agents submit cardinal ballots instead of approval ones. 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 of the problem.

8.2.6 Weighted PB

The fact that projects have different costs in PB can be interpreted as them having different weight. In their study about PB with ordinal ballots, Aziz and Lee (2021) make a symmetrical 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.

Chapter 9

Beyond the Social Choice Take on Participatory Budgeting

The focus of this survey, as its title suggests, is the (computational) social choice literature on PB. Nevertheless, some related topics are worth presenting. First, we briefly discuss some frameworks from the social choice literature that are related to PB (Section 9.1), and then we take a more general look at PB and how it is implemented in practice (Section 9.2).

9.1 Related Frameworks and Fields

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

9.1.1 Multi-Winner Voting

The most obvious 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 the budget allocation is 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.

9.1.2 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 in 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; Pascual, Rzadca and Skowron, 2018).

9.1.3 Divisible Participatory Budgeting

Throughout this paper, 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 includes 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 behaviours (Aziz, Bogomolnaia and Moulin, 2019; Freeman, Pennock, Peters and Vaughan, 2021). This setting is also closely related to that of *probabilistic social choice* (Brandt, 2018).

9.1.4 Fair Allocation

PB also relates to the literature on fair allocation (Rothe, 2015; Brandt, Conitzer, Endriss, Lang and Procaccia, 2016) 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 needs not be a budget constraint. Note that some work consider the same model but with constraints on the outcome, though not necessarily budget constraints (Fain, Munagala and Shah, 2018).

9.2 PB in Practice

So far we have taken a very theoretical look at PB. However, theoretical analysis should be grounded in some observable facts. In this section we provide pointers to real-life PB processes for the interested researcher.

For a more empirical analysis of PB, we can look to political scientists. The seminal paper on the topic is probably that of Cabannes (2004) who describes the first PB processes in Brazil. Subsequently, Sintomer, Herzberg and Röcke (2008) analysed how PB was adapted from Brazil to Europe; Wampler (2012) analysed the core principles of PB; and several books have been published, presenting the relevant literature and the recent developments regarding PB (Shah, 2007; De Oliveira, 2017; De Vries, Nemec and Špaček, 2021; Wampler and Goldfrank, 2022). Finally, several books giving an overview over the different forms of PB processes around the world have been published in recent years, such as the books from Dias (2018), Dias, Enríquez and Júlio (2019) or Wampler, McNulty and Touchton (2021).

A lot of papers we have gone through also contain experimental studies they performed. Getting access to data about real-life PB processes is of critical importance here. Thankfully, the website PaBuLib.org (Stolicki, Szufa and Talmon, 2020) provides a lot of such data.

Finally, Dominik Peters compiled a list[1] of PB instances that includes many interesting parameters for the social choice scientist: number of votes and projects, budget limit, ballot format, etc… Let us provide some interesting facts from that list below. Plenty more are to be found out by going through the list.

- Sometimes the budget is increased after the vote to afford more projects, sometimes by up to 250% (Montreal 2021, Toulouse 2019, Gdynia 2021, Cambridge 2015-2021).
- Sometimes participation is incentivised by giving a bonus to districts with high turnout or to individual voters (Rome 2019, Gdynia 2016-2018, Kraków 2019).

[1] In case the clickable link does not work: wikipedia.org/wiki/List_of_participatory_budgeting_votes.

- Sometimes there is a minimum requirement for projects to be selected, for instance, a project has to receive at least 200 ‘points” to be selected (Gdańsk 2021).
- Sometimes unused funds are transferred to the next year (Gdańsk 2014, 2018).
- Sometimes projects are partially funded by individual donors (Gdynia 2021).

Chapter 10

Conclusion

We have now presented almost the entire research that has been conducted by social choice scientist on the topic of PB. This line of research is still quite young and there are many things that can be explored further. We conclude this survey by presenting several directions we believe are worth exploring for PB.

- **Investigating more expressive ballots.** As we have seen, PB with approval ballots is the most studied framework for PB. It seems to us that these ballots are unfortunately not expressive enough for a framework in which alternatives have different cost. In particular, it is problematic that the meaning of not approving a project is unclear. One ballot format that has not received much attention but that we find appealing is the cumulative ballot format. It would also be interesting to initiate a study of PB processes where "negative" opinions can be submitted (with explicit disapproval for instance).

- **Extending the literature on fairness.** Even though the literature on fairness in PB is already quite extensive, there are still several interesting directions to pursue.

  - Knowing whether the core of PB with approval ballots is always non-empty or not is an obvious open problem. Note that this is even open for unit-cost instances.
  - Among the proportionality axioms that we know can always be satisfied, FJR is one of the strongest. However, we don't know any *natural* rule that satisfies it. Devising such a rule, therefore, is an important open problem.
  - The existing analysis of the *price of fairness* in PB (Fairstein, Vilenchik, Meir and Gal, 2022) is still at a preliminary stage and there is a lot of room for improvement.
  - Linking to our first point, when considering approval ballots, it is particularly interesting to find results that apply to whole classes of satisfaction functions rather than a single one. This would mitigate the criticism that no satisfaction function on its own is fully convincing.
  - It is known that cohesive groups for large sets of projects do not really occur in real-life instances, finding strong requirements that would not involve the concept of cohesive groups would thus be more useful in practice.

- **Deepening the axiomatic analysis.** The corpus of axioms that have been introduced in the literature about PB is still rather slim. We believe that there is a crucial need to develop that side of the literature to have other means to compare rules than fairness guarantees. Investigating how to adapt the characterisation results from the multi-winner voting literature (Skowron, Faliszewski and Slinko, 2019; Lackner and Skowron, 2021) could be an interesting starting point (though potentially rather technical).

- **Devising a theory of explainable PB.** Taking the risk to be called trend-followers, we believe that there is room for explaining outcomes of voting procedures to the citizen. PB makes no exception here. The outcome of PB rules could be explained in a counterfactual fashion: "How much cheaper should the project have been to be selected? How many more supporters?..." A theory of explainable PB could also take the form of more principled, axiomatically-guided, approaches (see, *e.g.*, Procaccia, 2019; Boixel and Endriss, 2020).

- **Developing a Python library.** Several papers presents simulation results using similar approaches, rules, etc, meaning that a large set of authors must have somewhere on their laptop the exact same code. Some of this code is available, see e.g., Pabutools. It would be nice to gather that in a unified Python package (inspired by the `abcvoting` package Lackner, Regner, Krenn, Cela, Kompauer, Lackner, Szufa and Forster, 2021; Lackner, Regner and Krenn, 2023). That would make the results based on simulations more reliable as they would be less susceptible to having errors in the code. It would also help with reproducibility of said results. Finally, it would be a great tool for the adoption of newly developed PB rules: verified and approved code would be available.

- **Looking beyond the voting stage.** As we mentioned in the introduction, PB is usually a longer process that has several steps. Despite this, the social choice literature on PB has almost exclusively focused on the voting stage, with the end-to-end model of Rey, Endriss and de Haan (2021) being the only exception. We believe that there are many interesting questions which can be answered with social choice methods that arise from a more holistic view of the PB process. This includes, for example, which incentives voters have when planning a project that they want to propose, *i.e.*, is it beneficial to make a project as cheap as possible or to merge two similar projects.

- **Stepping outside the Western world.** It is worth pointing out that the (computational) social choice literature so far has almost exclusively used PB processes in the Western world as examples. However, there is a large diversity in the actual implementation of PB around the world. For example in Western countries, generally only a small percentage of a municipalities budget is allocated to PB and most projects funded through PB 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 the cities budget was spend through PB and many projects funded through PB addressed crucial parts of life like access to basic health care (Cabannes, 2004). For a systematic analysis of the differences between PB processes in different parts of the world, we refer to the book by Wampler, McNulty and Touchton (2021). Taking a more global perspective on PB could open interesting new research directions.

Bibliography

Stéphane Airiau, Haris Aziz, Ioannis Caragiannis, Justin Kruger, Jérôme Lang, and Dominik Peters. 2023. Portioning Using Ordinal Preferences: Fairness and Efficiency. *Artificial Intelligence* 314 (2023), 103809. (Cited on page 61)

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 61)

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 23, 24, 25, and 26)

Haris Aziz, Edith Elkind, Shenwei Huang, Martin Lackner, Luis Sánchez-Fernández, and Piotr Skowron. 2018. On the Complexity of Extended and Proportional Justified Representation. In *Proceedings of the 32rd AAAI Conference on Artificial Intelligence (AAAI)*. (Cited on pages 24 and 27)

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)*. (Cited on pages 10, 39, 52, and 58)

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 10, 52, and 58)

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)*. (Cited on pages 10, 12, 16, 23, 36, 38, 41, 43, and 59)

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)*. (Cited on pages 10, 21, 26, 30, 31, 32, 37, 41, 43, 44, and 50)

Haris Aziz and Nisarg Shah. 2021. Participatory Budgeting: Models and Approaches. In *Pathways between Social Science and Computational Social Science: Theories, Methods and Interpretations*. Springer-Verlag. (Cited on page 5)

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 10 and 52)

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).* (Cited on pages 10, 39, and 58)

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 10, 19, and 47)

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 10, 11, 12, and 13)

Jonathan Bendor, Daniel Diermeier, David A Siegel, and Michael Ting. 2011. *A Behavioral Theory of Elections.* Princeton University Press. (Cited on page 7)

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 60)

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 61)

Arthur Boixel and Ulle Endriss. 2020. Automated Justification of Collective Decisions via Constraint Solving. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS).* 168–176. (Cited on page 64)

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 61)

Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (Eds.). 2016. *Handbook of Computational Social Choice.* Cambridge University Press. (Cited on page 61)

Robert Bredereck, 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 24)

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 10, 14, 15, 16, 20, 21, 22, 28, 29, 30, 31, 35, 36, 41, and 43)

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)

Yves Cabannes. 2004. Participatory budgeting: A significant contribution to participatory democracy. *Environment and Urbanization* 16, 1 (2004), 27–46. (Cited on pages 4, 61, and 64)

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 61)

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 49)

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 21)

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 10, 52, 56, and 58)

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 33)

City of Amsterdam. 2022. Oost Begroot. https://www.amsterdam.nl/stadsdelen/oost/oost-begroot. Last accessed on March the 1st 2023. (Cited on pages 4 and 55)

City of Paris. 2020. Paris Budget Participatif. https://decider.paris.fr/decider/jsp/site/Portal.jsp. Last accessed on March the 1st 2023. (Cited on page 55)

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 61)

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 60)

Andreas Darmann, Christian Klamler, and Ulrich Pferschy. 2011. Finding Socially Best Spanning Trees. *Theory and Decision* 70, 4 (2011), 511–527. (Cited on page 60)

Osmany Porto De Oliveira. 2017. *International Policy Diffusion and Participatory Budgeting: Ambassadors of Participation, International Institutions and Transnational Networks.* Springer-Verlag. (Cited on page 61)

Michiel S. De Vries, Juraj Nemec, and David Špaček. 2021. *International Trends in Participatory Budgeting: Between Trivial Pursuits and Best Practices.* Springer-Verlag. (Cited on page 61)

Amrita Dhillon and Susana Peralta. 2002. Economic Theories of Voter Turnout. *The Economic Journal* 112, 480 (2002), F332–F352. (Cited on page 7)

Nelson Dias (Ed.). 2018. *Hope for democracy: 30 years of participatory budgeting*. Epopeia Records and Oficina, Vila Ruiva and Faro. (Cited on pages 4 and 61)

Nelson Dias, Sahsil Enríquez, and Simone Júlio (Eds.). 2019. *The Participatory Budgeting World Atlas.* Epopee Records: Officinal Coordination. (Cited on pages 4 and 61)

Franz Dietrich and Christian List. 2007. Strategy-Proof Judgment Aggregation. *Economics & Philosophy* 23, 3 (2007), 269–300. (Cited on page 7)

Michael Dummett. 1984. *Voting Procedures.* Oxford University Press. (Cited on page 36)

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 page 49)

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 56)

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)*. (Cited on pages 32 and 61)

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 32, 33, 56, and 61)

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 10, 11, 13, and 14)

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)*. (Cited on pages 29, 30, and 63)

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.). AI Access. (Cited on page 60)

Till Fluschnik, Piotr Skowron, Mervin Triphaus, and Kai Wilker. 2019. Fair Knapsack. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*. (Cited on pages 10, 20, 52, 53, and 60)

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 61)

Allan Gibbard. 1973. Manipulation of Voting Schemes: a General Result. *Econometrica* (1973), 587–601. (Cited on page 7)

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 10, 11, 12, 15, 16, 48, 49, and 61)

Sven Ove Hansson. 2001. *The Structure of Values and Norms*. Cambridge University Press. (Cited on page 7)

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)*. (Cited on pages 10, 52, and 55)

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 9, 10, 52, and 57)

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)*. (Cited on pages 10, 52, and 56)

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 57)

Svante Janson. 2016. Phragmén's and Thiele's election methods. *arXiv preprint arXiv:1611.08826* (2016). (Cited on page 20)

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 10, 33, and 34)

Hans Kellerer, Ulrich Pferschy, and David Pisinger. 2004. *Knapsack Problems*. Springer-Verlag. (Cited on pages 18, 19, and 53)

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 Multi-Agent Systems (EUMAS)*. (Cited on page 48)

Martin Lackner. 2020. Perpetual Voting: Fairness in Long-Term Decision Making. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI)*. 2103–2110. (Cited on page 55)

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)*. (Cited on pages 10, 15, 16, 39, and 55)

Martin Lackner, Peter Regner, and Benjamin Krenn. 2023. `abcvoting`: A Python Package for Approval-Based Multi-Winner Voting Rules. *Journal of Open Source Software* 8, 81 (2023), 4880. (Cited on page 64)

Martin Lackner, Peter Regner, Benjamin Krenn, Elvi Cela, Jonas Kompauer, Florian Lackner, Stanisław Szufa, and Stefan Schlomo Forster. 2021. abcvoting: A Python library of approval-based committee voting rules. https://doi.org/10.5281/zenodo.3904466 Current version: https://github.com/martinlackner/abcvoting. (Cited on page 64)

Martin Lackner and Piotr Skowron. 2020. Utilitarian Welfare and Representation Guarantees of Approval-Based Multiwinner Rules. *Artificial Intelligence* 288 (2020), 103366. (Cited on page 30)

Martin Lackner and Piotr Skowron. 2021. Consistent Approval-Based Multi-Winner Rules. *Journal of Economic Theory* 192 (2021), 105173. (Cited on page 63)

Martin Lackner and Piotr Skowron. 2023. *Multi-Winner Voting with Approval Preferences*. Springer-Verlag. (Cited on pages 6, 23, 24, 34, 46, 47, and 60)

Annick Laruelle. 2021. Voting to Select Projects in Participatory Budgeting. *European Journal of Operational Research* 288, 2 (2021), 598–604. (Cited on pages 10, 16, 20, and 51)

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 7)

Maaike Los, Zoé Christoff, and Davide Grossi. 2022. Proportional Budget Allocations: Towards a Systematization. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*. (Cited on pages 10, 20, 23, 27, 29, 30, 31, 35, 38, 40, 41, and 42)

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 10, 52, and 57)

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 pages 10, 26, and 39)

Reshef Meir. 2018. *Strategic Voting*. Morgan & Claypool Publishers. (Cited on page 7)

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 61)

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)*. (Cited on pages 6, 10, 39, 48, 52, 56, and 57)

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 10 and 34)

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 10 and 33)

Fanny Pascual, Krzysztof Rzadca, 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 60)

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)*. (Cited on pages 10, 52, and 56)

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 7, 48, and 49)

Domink Peters. 2019. *Fair Division of the Commons*. DPhil Thesis. University of Oxford. (Cited on page 48)

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)*. (Cited on pages 10, 11, 22, 23, 25, 26, 27, 29, 32, 33, 34, 35, 36, 37, 40, 41, 42, 43, 45, and 50)

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)*. (Cited on pages 22, 24, 34, and 35)

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 page 49)

Ariel D. Procaccia. 2019. Axioms Should Explain Solutions. In *The Future of Economic Design*, Jean-François Laslier, Hervé Moulin, Remzi Sanver, and William S. Zwicker (Eds.). Springer-Verlag. (Cited on page 64)

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 12)

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)*. (Cited on pages 10, 47, 56, and 57)

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)*. (Cited on pages 6, 10, 54, and 64)

Jörg Rothe (Ed.). 2015. *Economics and Computation*. Springer-Verlag. (Cited on page 61)

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 27)

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 21)

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 page 7)

Anwar Shah (Ed.). 2007. *Participatory budgeting*. The World Bank. (Cited on pages 4 and 61)

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 61)

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 63)

Piotr Skowron, Arkadii Slinko, Stanisław Szufa, and Nimrod Talmon. 2020. Participatory Budgeting with Cumulative Votes. *arXiv preprint arXiv:2009.02690* (2020). (Cited on pages 10, 12, 23, 38, and 39)

Gogulapati Sreedurga, Mayank Ratan Bhardwaj, and Y. Narahari. 2022. Maxmin Participatory Budgeting. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*. (Cited on pages 10, 19, 20, 47, 49, 51, 52, and 53)

Dariusz Stolicki, Stanisław Szufa, and Nimrod Talmon. 2020. Pabulib: A Participatory Budgeting Library. *arXiv preprint arXiv:2012.06539* (2020). (Cited on pages 9 and 61)

Nimrod Talmon and Piotr Faliszewski. 2019. A Framework for Approval-Based Budgeting Methods. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI).* (Cited on pages 10, 15, 18, 19, 44, 45, 47, 52, 53, and 56)

Brian Wampler. 2000. A Guide to Participatory Budgeting. *Third conference of the International Budget Project* (2000). (Cited on page 4)

Brian Wampler. 2012. Participatory Budgeting: Core Principles and Key Impacts. *Journal of Public Deliberation* (2012). (Cited on page 61)

Brian Wampler and Benjamin Goldfrank. 2022. *The Rise, Spread, and Decline of Brazil's Participatory Budgeting: the Arc of a Democratic Innovation.* Springer-Verlag. (Cited on page 61)

Brian Wampler, Stephanie McNulty, and Michael Touchton. 2021. *Participatory Budgeting in Global Perspective.* Oxford University Press. (Cited on pages 4, 61, and 64)