

Fairness in Participatory Budgeting via Equality of Resources

Jan Maly

ILLC, University of Amsterdam
Amsterdam, The Netherlands
j.f.maly@uva.nl

Ulle Endriss

ILLC, University of Amsterdam
Amsterdam, The Netherlands
u.endriss@uva.nl

Simon Rey

ILLC, University of Amsterdam
Amsterdam, The Netherlands
s.j.rey@uva.nl

Martin Lackner

DBAI, TU Wien
Vienna, Austria
lackner@dbai.tuwien.ac.at

ABSTRACT

We introduce a family of normative principles to assess fairness in the context of participatory budgeting. These principles are based on the fundamental idea that budget allocations should be fair in terms of the resources invested into meeting the wishes of individual voters. This is in contrast to earlier proposals that are based on specific assumptions regarding the satisfaction of voters with a given budget allocation. We analyse these new principles in axiomatic, algorithmic, and experimental terms.

KEYWORDS

Computational Social Choice; Participatory Budgeting; Fairness

ACM Reference Format:

Jan Maly, Simon Rey, Ulle Endriss, and Martin Lackner. 2023. Fairness in Participatory Budgeting via Equality of Resources. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 9 pages.

1 INTRODUCTION

Budgeting, *i.e.*, the allocation of money or other sparse resources to specific projects, is one of the key decisions any political body or organisation has to take. Participatory budgeting (PB) was developed in the 1990s as a method for making such decisions in a more democratic way, by putting it to a vote [7, 22]. It has found rapid adoption worldwide, in particular at the municipal level [26].

The most common form of eliciting the views of the voters is to ask which projects they approve of [13], but the question of which voting rules should be used to select the projects to be funded is not yet settled. In this paper, we advocate for the use of fairness measures based on the resources spent on behalf of individual voters to guide the search for the best and most equitable voting rules. Specifically, we focus on a measure of *equality of resources* [8, 9] called the *share*, recently introduced by Lackner et al. [15]. It is computed by equally dividing the cost of each funded project amongst the supporters of that project.

Let us briefly motivate this approach. Suppose 40% of citizens of a city support funding more cycling infrastructure, while 60% are in favour of more car infrastructure. Then, under the kind of voting rule usually employed in practice, where the projects with

the most support get selected, only car-centric projects would get funded. This clearly is not desirable. Instead, one would hope to select a *proportional* outcome [3, 4, 10, 17, 19], funding a mixture of cycling and car infrastructure projects. But how should one define proportionality? So far, the literature has focused on generalisations from approval-based multiwinner voting, where we often aim for a proportional distribution of *satisfaction* amongst voters, assuming that each approved candidate provides the same satisfaction to all of their supporters [11, 16]. However, lifting this assumption to the richer framework of PB is questionable, as projects vary in cost.

So, given their approval ballot, how should one infer a voter’s satisfaction for a set of selected projects?¹ Most researchers assume that the satisfaction of a voter is either equal for all approved projects [17, 19, 25] or proportional to the cost of a project [4, 10, 15, 23]. Both assumptions are problematic. Regarding the former, for example, in the 2019 Toulouse participatory budgeting process a cycling infrastructure project costing 390,000 euros (“Tous à la Ramée à vélo! A pied, en trottinette et rollers!”) and a project about installing birdhouses costing 2,000 euros (“Nichoires pour mésanges”) were proposed (see Pabulib [24]). It seems unlikely that both projects offer the same utility to their supporters. But full proportionality of utility and cost also seems implausible, because the cost effectiveness of different projects can vary widely. Consider, for example, a scenario where two parks of equal size could be built in different neighbourhoods. Now, it might be more expensive to build the park in one neighbourhood due to higher property prices. In that case, there is no reason to assume that the more expensive park offers more utility to its supporters. Crucially, these two examples show that higher cost sometimes implies higher utility, while sometimes it does not. This makes it hard to imagine a way of estimating utilities in a principled way that works for both examples.

To circumvent these difficulties we propose to develop fairness measures that are not based on *equality of welfare* but that instead aim for *equality of resources*, an idea first proposed by Ronald Dworkin [8, 9]. In other words, we do not aim for a fair distribution of satisfaction, but instead we strive to invest the same effort into satisfying each voter. The advantage is that the amount of resources spent is a quantity we can measure objectively. This idea can be formalised through the notion of *share* [15]. Ideally, we want to

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom. © 2023 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

¹In principle, there is also another possibility, namely to directly ask voters for their satisfaction (or utility). But this would impose a significant cognitive burden on them, and it is debatable whether it is even possible to elicit utilities in a way that allows for interpersonal comparisons [6, 14].

find a budget allocation where each voter has the same share. Let us emphasise that we do not interpret the share as a measure of satisfaction, but rather of a distribution of resources. Interestingly, fairness notions based on share also provide an explanation on how each voter’s part of the budget was spent. In contrast to the related notion of *priceability* [20], here all supporters of a project “contribute” the same amount. As priceability allows for unequal contribution of voters, it does not qualify as a notion based on equality of resources.

In this paper, we investigate the viability of the share as a basis of fairness notions in PB in several complementary ways. First, we propose several axioms that formalise what it means for an outcome to be fair in terms of share. We observe that it is not always possible to guarantee everyone their *fair share*, which we define as the budget divided by the numbers of voters. For this reason, we consider several relaxations, such as the *justified share*, where we only aim to allocate to a voter the resources they deserve by virtue of being part of a cohesive group. Moreover, we identify a version of MES [19], that satisfies all share-based axioms known to be satisfiable by a tractable voting rule. Finally, using data from a large number of real-life PB exercises [24], we analyse to what extent it is possible to provide voters with their fair share in practice and how well established PB voting rules meet share-based desiderata.

Roadmap. After introducing our model in Section 2, we investigate the fair share in Section 3 and the justified share in Section 4. Relationships between the concepts are discussed in Section 5, while Section 6 reports on an experimental study. Due to space constraints, some proofs are only available in the full version [18].

2 THE MODEL

A PB problem is described by an *instance* $I = \langle \mathcal{P}, c, b \rangle$ where \mathcal{P} is the set of available *projects*, $c : \mathcal{P} \rightarrow \mathbb{N}$ is the *cost function*—mapping any project $p \in \mathcal{P}$ to its cost $c(p) \in \mathbb{N}$ —and $b \in \mathbb{N}$ is the *budget limit*. We write $c(P)$ instead of $\sum_{p \in P} c(p)$ for sets of projects $P \subseteq \mathcal{P}$. If $c(p) = 1$ for all $p \in \mathcal{P}$, then I belongs to the *unit-cost setting*.

Let $\mathcal{N} = \{1, \dots, n\}$ be a set of *agents*. When facing a PB instance, each agent is asked to submit a (not necessarily feasible) *approval ballot* representing the subset of projects they approve of. The approval ballot of agent $i \in \mathcal{N}$ is denoted by $A_i \subseteq \mathcal{P}$, and the resulting vector $\mathbf{A} = (A_1, \dots, A_n)$ is called a *profile*. We assume without loss of generality that every project is approved by at least one agent.

Given an instance $I = \langle \mathcal{P}, c, b \rangle$, we need to select a set of projects $\pi \subseteq \mathcal{P}$ to implement. Such a *budget allocation* π has to be *feasible*, i.e., we require $c(\pi) \leq b$. Let $\mathcal{A}(I) = \{\pi \subseteq \mathcal{P} \mid c(\pi) \leq b\}$ be the set of feasible budget allocations for I .

Choosing allocations is done by means of (resolute) *PB rules*. Such a rule F is a function that maps any instance I and profile \mathbf{A} over I to a single feasible budget allocation $F(I, \mathbf{A}) \in \mathcal{A}(I)$. Whenever ties occur between several outcomes, we assume that they are broken in a fixed and consistent manner (e.g., lexicographically).

We are going to propose several fairness properties we might want a rule to satisfy. All of these properties will be defined in terms of the fundamental notion of an agent’s *share*.

Definition 1 (Share). Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile \mathbf{A} , the *share of agent i for a subset of projects $P \subseteq \mathcal{P}$* is defined as follows:

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

When clear from context, we shall omit the arguments of I and \mathbf{A} . We interpret an agent’s share as the amount of resources spent by the decision maker on satisfying the needs of that agent. It is important to note that the share cannot be captured via independent cardinal utility functions as the share of an agent depends on the ballots submitted by the other agents. Let us illustrate the concept of share on an example.

Example 1. Consider a PB instance with three projects such that $c(p_1) = 8$ and $c(p_2) = c(p_3) = 2$, and a budget limit $b = 8$. The profile \mathbf{A} is composed of four ballots such that $A_1 = A_2 = \{p_1, p_2\}$, $A_3 = \{p_1\}$ and $A_4 = \{p_3\}$. The most commonly used PB rule, which greedily picks the most approved projects, would select the bundle $\{p_1, p_2\}$. This gives agents 1 and 2 a share of $6/3 + 2/2 = 3$, agent 3 a share of $6/3 = 2$ and agent 4 a share of 0. We claim that $\{p_1, p_3\}$ would be a fairer bundle as it gives agent 1, 2 and 3 a share of $6/3 = 2$ and agent 4 a share of $2/1 = 2$. Hence, all agents have the same share.

In the sequel, we will introduce several properties of budget allocations. We shall extend them to properties of rules so that rule F is said to satisfy fairness property \mathcal{F} defined for budget allocations if, for every I and \mathbf{A} , the outcome $F(I, \mathbf{A})$ satisfies \mathcal{F} .

3 FAIR SHARE

The first fairness property we study is based on the idea that each voter deserves $1/n$ of the budget—a fundamental idea familiar, for instance, from the classical fair division (“cake cutting”) literature [21]. So a perfect allocation would give each voter a share of b/n (unless they do not approve of enough projects for this to be possible).

Definition 2 (Fair Share). Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile \mathbf{A} , the *fair share of agent $i \in \mathcal{N}$* is defined as:

$$\text{fairshare}(i) = \min\{b/n, \text{share}(A_i, i)\}.$$

A budget allocation $\pi \in \mathcal{A}(I)$ is said to satisfy *fair share (FS)* if for every agent we have $\text{share}(\pi, i) \geq \text{fairshare}(i)$.

It is easy to see that for some instances, no budget allocation would provide a fair share, and thus no rule can possibly satisfy FS. Take for instance two projects of cost 1, a budget limit of 1 and two agents each approving of a different project. Then, both agents have a fair share of $\min\{1/2, 1\} = 1/2$. However, whichever project is selected (at most one can be selected), the share of one agent would be 0.

Even more, we can show that no polynomial-time computable rule can return an FS allocation whenever one exists. Indeed, checking whether an FS allocation exists is NP-complete.

Proposition 1. Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile \mathbf{A} , checking whether there exists a feasible budget allocation $\pi \in \mathcal{A}(I)$ that satisfies FS is NP-complete, even in the unit-cost setting.

The proof involves a reduction from 3-SET-COVER [12]. Due to these shortcomings of FS, we introduce two relaxations that are inspired by relaxations of important, satisfaction based fairness axioms, Extended Justified Representation up to one project (EJR-1) [19] and Local-Budget Proportional Justified Representation (Local BPJR) [4].

Definition 3 (FS up to one project). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile A , a budget allocation $\pi \in \mathcal{A}(I)$ is said to satisfy fair share up to one project (FS-1) if, for every agent i , there is a project $p \in \mathcal{P}$ such that:*

$$\text{share}(\pi \cup \{p\}, i) \geq \text{fairshare}(i).$$

Thus, FS-1 requires that every agent is only one project away from their fair share. Unfortunately, FS-1 is not always satisfiable.

Proposition 2. *There exist instances I for which no budget allocation $\pi \in \mathcal{A}(I)$ provides FS-1.*

PROOF. Consider an instance with three projects of cost 3 and a budget limit $b = 5$. Consider three agents, with approval ballots $\{p_1, p_2\}$, $\{p_1, p_3\}$ and $\{p_2, p_3\}$, respectively.

Here the fair share of each agent is $5/3 \approx 1.67$. As a single project only yields a share of 1.5 to an agent who approves of it, for any agent to reach their fair share threshold, two projects must be selected. However, a feasible budget allocation can select at most one project, meaning that for one agent none of the projects they approve of will be selected. \square

As for FS, we can show that deciding whether an FS-1 budget allocation exists is NP-complete.

Proposition 3. *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile A , checking whether there exists a feasible budget allocation $\pi \in \mathcal{A}(I)$ that satisfies FS-1 is NP-complete, even in the unit-cost setting.*

Alternatively, we can require that every project that is not part of the winning budget allocation should give some voter at least their fair share when that project is added.²

Definition 4 (Local-FS). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile A , a budget allocation $\pi \in \mathcal{A}(I)$ is said to satisfy local fair share (Local-FS) if there is no project $p \in \mathcal{P} \setminus \pi$ such that, for all agents $i \in \mathcal{N}$ with $p \in A_i$, we have:*

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

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

- p can be selected without exceeding the budget limit b ;
- some voter i^* receives more than their fair share.

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

In contrast to FS-1, we can always find an allocation that satisfies Local-FS. Indeed, an adaption of the *Method of Equal Shares* (MES)³ [19] satisfies Local-FS. Our definition closely resembles the definition of MES for PB with additive utilities [19]. We adapt it by plugging in the share. Note that this rule can be executed in polynomial time.

²We stress that this formulation of Local-FS relies on our assumption that every project p is approved by at least one agent.

³The rule used to be named Rule X until recently. The new name—*method of equal share*—is not related to the definition of share, first introduced by Lackner et al. [15].

Definition 5 ($\text{MES}_{\text{share}}$). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile A , $\text{MES}_{\text{share}}$ constructs a budget allocation π , initially empty, iteratively as follows. A load $\ell_i : 2^{\mathcal{P}} \rightarrow \mathbb{R}_{\geq 0}$, is associated with every agent $i \in \mathcal{N}$, initialised as $\ell_i(\emptyset) = 0$ for all $i \in \mathcal{N}$. It represents the total contribution of the agents for a given budget allocation. Given π and a scalar $\alpha \geq 0$, the contribution of agent $i \in \mathcal{N}$ for project $p \in \mathcal{P} \setminus \pi$ is defined by:*

$$\gamma_i(\pi, \alpha, p) = \min(b/n - \ell_i(\pi), \alpha \cdot \text{share}(\{p\}, i)).$$

For a specific budget allocation π , a project $p \in \mathcal{P} \setminus \pi$ is said to be α -affordable, for $\alpha \geq 0$, if $\sum_{i \in \mathcal{N}} \gamma_i(\pi, \alpha, p) \cdot \mathbb{1}_{p \in A_i} = c(p)$, i.e., if p can be afforded with none of its supporters contributing more than α .

At a given round with current budget allocation π , if no project is α -affordable for any α , $\text{MES}_{\text{share}}$ terminates. Otherwise, it selects a project $p \in \mathcal{P} \setminus \pi$ that is α^* -affordable where α^* is the smallest α such that one project is α -affordable (π is updated to $\pi \cup \{p\}$). The agents’ loads are then updated: If $p \notin A_i$, then $\ell_i(\pi \cup \{p\}) = \ell_i(\pi)$, and otherwise $\ell_i(\pi \cup \{p\}) = \ell_i(\pi) + \gamma_i(\pi, \alpha, p)$.

Theorem 4. *$\text{MES}_{\text{share}}$ satisfies Local-FS.*

PROOF. Given a budget allocation π and a scalar $\alpha > 0$, we say that agent $i \in \mathcal{N}$ contributes in full to project $p \in A_i$ if we have: $\gamma_i(\pi, \alpha, p) = \alpha \cdot \text{share}(\{p\}, i)$.

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

In fact, the proof of Theorem 4 establishes a slightly stronger statement: there is no project $p \in \mathcal{P} \setminus \pi$ such that for all agents $i \in \mathcal{N}$ with $p \in A_i$ we have $\text{share}(\pi \cup \{p\}, i) \leq b/n$. In other words, any project added to π gives at least one voter *more* than their fair share.

4 JUSTIFIED SHARE

Local-FS and FS-1 require the outcome to be, in some sense, close to satisfying FS. Another idea for weakening FS is to spend on a voter only the resources they can claim to deserve by virtue of being part of a cohesive group. This idea is inspired by the well-known axioms of justified representation extensively studied in the literature on approval-based committee elections [1, 2, 16, 20]. Before exploring this idea further, let us define what we mean by cohesive groups.

Definition 6 (P -cohesive groups). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile A , for a set of projects $P \subseteq \mathcal{P}$ we say that a non-empty group of agents $N \subseteq \mathcal{N}$ is P -cohesive, if $P \subseteq \bigcap_{i \in N} A_i$ and $\frac{|N|}{n} \geq \frac{c(P)}{b}$.*

So a group N is cohesive relative to a set P of projects if, first, everyone in N approves of all the projects in P and, second, N is large enough—relative to the size n of the society and the budget b —so as to “deserve” the resources needed for the projects in P .

In the unit-cost setting, one of the strongest proportionality properties known to be satisfiable by a polynomial-time-computable rule is called Extended Justified Representation (EJR) [2, 20]. It ensures that one member of every cohesive group receives the satisfaction deserved due to their cohesiveness. Peters et al. [19] generalised EJR to the setting of PB with additive utilities. This generalisation will be our blue-print for modifying EJR to deal with share. Ideally, we would want to satisfy the following property.

Definition 7 (Strong Extended Justified Share). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile A , a budget allocation $\pi \in \mathcal{A}(I)$ is said to satisfy strong extended justified share (Strong-EJS) if for all $P \subseteq \mathcal{P}$ and all P -cohesive groups N , we have $\text{share}(\pi, i) \geq \text{share}(P, i)$ for all agents $i \in N$.*

The idea behind Strong-EJS is the following: since every P -cohesive group S controls enough budget to fund P , every agent in S deserves to enjoy at least as much share as what they would have gotten if P had been the outcome. Intuitively, this is very similar to Strong-EJR, a property that is known not to be always satisfiable [1]. The same holds for Strong-EJS: there exist instances for which no budget allocation can satisfy this axiom.

Example 2. *Consider the following instance and profile with three projects p_1, p_2 and p_3 of cost 1, a budget limit $b = 2$, and four agents $1, \dots, 4$ such that 1 approves project p_1 , 2 approves project p_1 and p_2 , 3 approves p_1 and p_3 and 4 approves p_2 and p_3 . Note that $\{1, 2, 3\}$ is $\{p_1\}$ -cohesive, $\{2, 4\}$ is $\{p_2\}$ -cohesive and $\{3, 4\}$ is $\{p_3\}$ -cohesive. Hence, to satisfy Strong-EJS, one needs to select all three projects which is not possible within the given budget limit.*

Observe that in this scenario it is not even possible to guarantee each P -cohesive group the same average share as they receive from P . We thus weaken Strong-EJS and introduce (simple) EJS.

Definition 8 (Extended Justified Share). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile A , a budget allocation $\pi \in \mathcal{A}(I)$ is said to satisfy extended justified share (EJS), if for all $P \subseteq \mathcal{P}$ and all P -cohesive groups N , there is an agent $i \in N$ such that $\text{share}(\pi, i) \geq \text{share}(P, i)$.*

The difference between Strong-EJS and EJS is the switch from a universal to an existential quantifier: for the former, we impose a lower bound on the share of every agent in a cohesive group, while for the latter this lower bound only applies to one agent of each cohesive group. Therefore, in Example 2 both $\{p_1, p_3\}$ and $\{p_2, p_3\}$ satisfy EJS, as either agent 3 or agent 4 satisfies the share requirement for every cohesive group.

We observe that EJR and EJS, while similar in spirit, do not coincide, not even in the unit-cost case.

Example 3. *Consider an instance with four voters and six projects with unit cost and $b = 4$, where the approvals are as follows: $A_1 = \{p_1, p_2, p_3\}$, $A_2 = \{p_1, p_2, p_4\}$, $A_3 = A_4 = \{p_4, p_5, p_6\}$. It is now easy to check that $\{p_3, p_4, p_5, p_6\}$ satisfies EJS but not EJR, while $\{p_1, p_4, p_5, p_6\}$ satisfies EJR but not EJS.*

The first question that presents itself is whether EJS is always achievable. This is indeed the case. To see this, one just needs to adapt the well-known greedy cohesive procedure for satisfying EJR, which was first introduced by Aziz et al. [1] and extended to PB by Peters et al. [19], to the share setting.

Proposition 5. *For every instance $I = \langle \mathcal{P}, c, b \rangle$ and every profile A , there exists a budget allocation $\pi \in \mathcal{A}(I)$ that satisfies EJS.*

However, the greedy approach in general needs exponential time. This turns out to be unavoidable, unless $P = NP$, as can be shown by a standard reduction from SUBSET SUM.

Theorem 6. *There is no polynomial-time algorithm that, given an instance I and a profile A as input, always computes a budget allocation satisfying EJS, unless $P = NP$.*

On the other hand, we recall that the greedy approach generally runs in FPT-time, when parameterized by the number of projects [1]. This is also the case in the share setting.

Proposition 7. *For every instance $I = \langle \mathcal{P}, c, b \rangle$ and every profile A , we can compute a budget allocation $\pi \in \mathcal{A}(I)$ that satisfied EJS in time $O(n \cdot 2^{|\mathcal{P}|})$.*

We have seen that EJS can always be satisfied. However, this is not entirely satisfactory, given that no tractable rule can satisfy it. Unfortunately, in many PB applications, the use of intractable rules is not practical due to the large instance sizes. Therefore, we try to find fairness notions that can be satisfied in polynomial time by relaxing EJS. First, similar to the property of EJR up to one project (EJR-1) proposed by Peters et al. [19], we can define EJS up to one project, requiring that at least one agent in every cohesive group is at most one project away from being satisfied.⁴

Definition 9 (EJS-1). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile A , a budget allocation $\pi \in \mathcal{A}(I)$ is said to satisfy extended justified share up to one project (EJS-1) if for all $P \subseteq \mathcal{P}$ and all P -cohesive groups N there is an agent $i \in N$ for which there exists a project $p \in \mathcal{P}$ such that $\text{share}(\pi \cup \{p\}, i) \geq \text{share}(P, i)$.*

It is straightforward to adapt the proof of Peters et al. [19] that MES satisfies EJR up to one project to our setting to prove that $\text{MES}_{\text{share}}$ satisfies EJS up to one project.

Proposition 8. *$\text{MES}_{\text{share}}$ satisfies EJS-1.*

In particular, this implies, together with Example 3, that already in the unit-cost case MES and $\text{MES}_{\text{share}}$ are indeed different rules.

Finally, note that we can define a local variant of EJS, based on a similar motivation as Local-FS.

Definition 10 (Local-EJS). *Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile A , a budget allocation $\pi \in \mathcal{A}(I)$ is said to satisfy local extended justified share (Local-EJS), if there is no P -cohesive group N , where $P \subseteq \mathcal{P}$, for which there exists a project $p \in P \setminus \pi$ for which it holds for all agents $i \in N$ that $\text{share}(\pi \cup \{p\}, i) < \text{share}(P, i)$.*

The idea behind Local-EJS is that there is no P -cohesive group N that can claim that they could “afford” another project p without

⁴We note that in Definition 9 we require that $\text{share}(\pi \cup \{p\}, i) \geq \text{share}(P, i)$ instead of a strict inequality as used in the definition of EJR-1. Our rationale is that adding one project guarantees to satisfy the EJS condition (but not more than that).

a single voter in N receiving more share than they deserve due to their P -cohesiveness. In this sense, any allocation that satisfies Local-EJS is a local optimum for any P -cohesive group. Now, in our setting we observe that Local-EJS is equivalent to a notion that could be called “EJS up to any project”.

Proposition 9. *Let $I = \langle \mathcal{P}, c, b \rangle$ be an instance and A a profile. An allocation π satisfies Local-EJS if and only if for every $P \subseteq \mathcal{P}$ and P -cohesive group N there exists an agent i such that for all projects $p \in P \setminus \pi$ we have $share(\pi \cup \{p\}, i) \geq share(P, i)$.*

PROOF. It is clear that the statement above implies Local-EJS. Now, let π be an allocation that satisfies Local-EJS, let $P \subseteq \mathcal{P}$ be a set of projects and N a P -cohesive group. Let $i^* \in N$ be an agent with maximal share from π in N . Consider $p \in P \setminus \pi$. By Local-EJS there is an agent i_p such that $share(\pi \cup \{p\}, i_p) > share(P, i_p)$. By the choice of i^* we have $share(\pi, i^*) \geq share(\pi, i_p)$. By the definition of share, it follows that $share(\pi \cup \{p\}, i^*) > share(P, i^*)$. \square

From this equivalence, it is easy to see that Local-EJS implies EJS-1. Unfortunately, MES_{share} fails Local-EJS, as the next example shows.

Example 4. *Consider an instance with five projects, a budget limit $b = 20$, and four agents where the costs are as follows:*

$$c(p_1) = 8, c(p_2) = 5, c(p_3) = c(p_4) = 2, c(p_5) = 10.$$

Moreover, voters 1 and 2 approve projects p_1, p_2, p_3 and p_4 and voters 3 and 4 approve p_3, p_4 and p_5 . With a suitable tie-breaking rule, MES_{share} will return the budget allocation $\pi = \{p_2, p_3, p_5\}$. Note that voters 1 and 2 are $\{p_1, p_4\}$ -cohesive and would thus deserve to enjoy a share of 4.5. However, if we add p_4 to π , voters 1 and 2 would only have a share of 3.5, showing that π fails Local-EJS.

Whether Local-EJS can always be satisfied in polynomial time remains an important open question.

Finally, we observe a crucial difference between EJR and EJS: MES_{share} does not satisfy EJS in the unit cost setting, while MES satisfies EJR in the unit-cost setting [20].

Example 5. *Assume that there are two voters 1 and 2, and three projects p_1, p_2 and p_3 , all of cost 1. The budget limit is $b = 2$. Voter 1 approves of p_1 and p_3 and voter 2 of p_2 and p_3 . Then voter 1 is $\{p_1\}$ -cohesive and hence deserves a share of 1, the same applies to voter 2 and $\{p_2\}$. Nevertheless, with a suitable tie-breaking rule, MES_{share} would first select p_3 . In that case, neither $\{p_1, p_2\}$, nor $\{p_2, p_3\}$ would satisfy EJS, as at least one voter will have only a share of $1/2$.*

However, it does satisfy Local-EJS in the unit-cost setting.

Theorem 10. *MES_{share} satisfies Local-EJS in the unit-cost case.*

5 RELATIONSHIPS BETWEEN CRITERIA

We now analyse the relationships between different fairness criteria. We start with links within the space of criteria we introduced earlier and then compare them to the notion of priceability [20].

5.1 Share-Based Fairness Criteria

The following theorem establishes the relations between share-based fairness concepts. These relations are visualised in Figure 1.




Figure 1: Taxonomy of criteria. An arrow from one criterion to another indicates that any budget allocation satisfying the former also satisfies the latter. MES_{share} satisfies the criteria boxed in green solid lines. For the criterion boxed in orange dashed lines, no efficient algorithms computing them exist (unless $P = NP$). Criteria boxed in red dotted lines are not always satisfiable. The status of Local-EJS is unknown.

Theorem 11. *Given an instance I and a profile A , for every budget allocation $\pi \in \mathcal{A}(I)$ the following statements hold:*

- (1) *If π satisfies FS, it also satisfies FS-1, Local-FS, and Strong-EJS.*
- (2) *If π satisfies FS-1, it also satisfies EJS-1.*
- (3) *If π satisfies Strong-EJS, it also satisfies EJS.*
- (4) *If π satisfies EJS, it also satisfies Local-EJS.*
- (5) *If π satisfies Local-EJS, it also satisfies EJS-1.*

This list of implications is exhaustive when closed under transitivity.

PROOF OF (1). It is easy to verify that every budget allocation satisfying FS also satisfy FS-1 and Local-FS. So let us show that FS also implies Strong-EJS. Let $i \in N$. We distinguish two cases.

First, assume $share(A_i, i) < b/n$. For FS to be satisfied, we must have $share(\pi, i) \geq share(A_i, i)$. This entails that $A_i \subseteq \pi$. Hence, the conditions for Strong-EJS are trivially satisfied for agent i .

Second, assume $share(A_i, i) \geq b/n$. Since π satisfies FS, we know that $share(\pi, i) \geq b/n$. Let $N \subseteq N$ be a P -cohesive group, for some $P \subseteq \mathcal{P}$, such that $i \in N$. By definition of a cohesive group, we know that $c(P) \leq b/n \cdot |N|$. Hence, $share(P, i) \leq b/n$. Overall, we have $share(\pi, i) \geq b/n \geq share(P, i)$ and thus π satisfies Strong-EJS. \square

The proof of claim (2) is similar to the proof of claim (1) and can be found in the full version on arXiv [18]. The proofs of claims (3) and (4) are immediately derived from the respective definitions. The proof of claim (5) is a direct consequences of Proposition 9.

The absence of any further implications between fairness criteria can be established by counterexamples. We include here a representative sample of such counterexamples. The remaining counterexamples can be found in the full version of the paper [18].

Example 6 (FS-1 not implying Local-FS). *Consider the following instance with four projects and a budget limit of $b = 6$.*

	p_1	p_2	p_3	p_4
Cost	3	3	6	1
A_1	✓	✓	✗	✗
A_2	✗	✗	✓	✓
A_3	✗	✗	✓	✓

Then $\{p_1, p_2\}$ satisfies FS-1 as agent 1 already receives (more than) their fair share, while 2 and 3 receive their fair share from $\{p_1, p_2\} \cup \{p_3\}$. However, no supporter of p_4 receives their fair share from $\{p_1, p_2\} \cup \{p_4\}$. Therefore, Local-FS is violated.

The other direction (Local-FS does not imply FS-1) follows from the fact that a Local-FS allocation always exists (as a consequence of Theorem 4) while FS-1 is not always satisfiable (Proposition 2).

Example 7 (EJS-1 not implying Local-EJS, even in the unit-cost setting). Consider an instance with two voters, 1 and 2, and six projects p_1, \dots, p_6 all of cost 1. Voter 1 approves of $\{p_1, p_2, p_3, p_4, p_5\}$ and voter 2 approves of $\{p_4, p_5, p_6\}$. The budget limit is $b = 4$. We claim that $\pi = \{p_1, p_2, p_3, p_4\}$ satisfies EJS-1 but not Local-EJS.

The share of 1 in π is 3.5 so every cohesive group containing them will satisfy the conditions for EJS-1 and Local-EJS. Consider now voter 2. Their share in π is $1/2$. Note that they are $\{p_5, p_6\}$ -cohesive and deserve thus a share of $3/2$. Since $\pi \cup \{p_6\}$ would provide them a share of $3/2$, π satisfies EJS-1. However, $\pi \cup \{p_5\}$ would only provide 2 a share of 1, showing that π fails Local-EJS.

We now turn to Local-FS and show that it does not imply EJS-1. Due to the implications shown in Theorem 11, Local-FS also does not imply any of Local-EJS, EJS, Strong-EJS, FS-1 and FS.

Example 8 (Local-FS not implying EJS-1). Consider the following instance with three projects, a budget limit of $b = 6$, and two agents. Agent 1 approves of $\{p_1, p_2, p_3\}$ and agent 2 of $\{p_2, p_3\}$. Allocation $\pi = \{p_1\}$ satisfies Local-FS: for both p_2 and p_3 , if we were to add them to π , agent 1 would have a fair share. However, it does not satisfy EJS-1: $\{2\}$ is a $\{p_2, p_3\}$ -cohesive group but neither project is selected.

5.2 Comparison with Priceability

Priceability is a fairness criterion requiring that the budget allocation can be obtained through a market-based approach [20]. It is similar in spirit to share-based criteria as it also measures the amount of money spent on each agent. However, priceability does not require the cost of a project to be equally distributed between its supporters. Instead it requires there to be some distribution of the costs of the selected projects to their supporters that satisfies certain conditions.

Definition 11 (Priceability). Given an instance $I = \langle \mathcal{P}, c, b \rangle$ and a profile A , a budget allocation π satisfies priceability if there exists an allowance $\alpha \in \mathbb{R}_{\geq 0}$ and a collection $(\gamma_i)_{i \in \mathcal{N}}$ of contribution functions, $\gamma_i : \mathcal{P} \rightarrow [0, \alpha]$ such that all of the following conditions are satisfied:

- C1: If $\gamma_i(p) > 0$ then $p \in A_i$ for all $p \in \mathcal{P}$ and $i \in \mathcal{N}$.
- C2: If $\gamma_i(p) > 0$ then $p \in \pi$ for all $p \in \mathcal{P}$ and $i \in \mathcal{N}$.
- C3: $\sum_{p \in \mathcal{P}} \gamma_i(p) \leq \alpha$ for all $i \in \mathcal{N}$.
- C4: $\sum_{i \in \mathcal{N}} \gamma_i(p) = c(p)$ for all $p \in \pi$.
- C5: $\sum_{i \in \mathcal{N} | p \in A_i} \alpha_i^* \leq c(p)$ for all $p \in \mathcal{P} \setminus \pi$, where for any $i \in \mathcal{N}$, $\alpha_i^* = \alpha - \sum_{p \in \mathcal{P}} \gamma_i(p)$ is their unspent allowance.

Due to the similar motivation of share based fairness concepts and priceability, it is interesting to understand the relationships between them. Let us first consider FS.

Proposition 12. There exists instances $I = \langle \mathcal{P}, c, b \rangle$ and profiles A such that there exists $\pi \in \mathcal{A}(I)$ satisfying FS, but such that no FS budget allocation π is priceable.

PROOF. Consider the following instance with four projects, a budget limit of $b = 9$, and three agents.

	p_1	p_2	p_3	p_4
Cost	1	5	3	1
A_1	✓	✗	✗	✗
A_2	✗	✓	✗	✗
A_3	✗	✗	✓	✓

In this instance, only $\pi = \{p_1, p_2, p_3\}$ satisfies FS. For the sake of contradiction, suppose that π is priceable with allowance $\alpha \in \mathbb{R}_{\geq 0}$ and contribution functions γ_1, γ_2 and γ_3 . Since only agent 2 approves of p_2 , from conditions C1 and C4 we must have $\gamma_2(p_2) = 5$. Condition C3 then implies that $\alpha \geq 5$. For similar reasons we should have $\gamma_3(p_3) = 3$ and $\gamma_3(p_1) = \gamma_3(p_2) = 0$. Condition C2 also imposes $\gamma_3(p_4) = 0$. Overall this means that $\alpha_3^* = \alpha - \gamma_3(p_3) \geq 2$. This is a violation of condition C5 for agent 3 and project p_4 . \square

Interestingly, the intuitive connection between fair share and priceability does hold when ballots are large enough.

Proposition 13. For every instance $I = \langle \mathcal{P}, c, b \rangle$ and profile A such that for every agent $i \in \mathcal{N}$, $\text{fairshare}(i) = b/n$, every budget allocation $\pi \in \mathcal{A}(I)$ that satisfies FS is also priceable.

PROOF. Consider a suitable instance $I = \langle \mathcal{P}, c, b \rangle$ and profile A . Let $\pi \in \mathcal{A}(I)$ be a budget allocation that satisfies FS. We claim that π is priceable for the allowance $\alpha = b/n$ and the contribution functions $(\gamma_i)_{i \in \mathcal{N}}$ defined for every $i \in \mathcal{N}$ and $p \in \mathcal{P}$ as:

$$\gamma_i(p) = \begin{cases} \text{share}(\{p\}, i) & \text{if } p \in A_i \cap \pi, \\ 0 & \text{otherwise.} \end{cases}$$

First note that conditions C1 and C2 of priceability are trivially satisfied for all $i \in \mathcal{N}$ and $p \in \mathcal{P}$. Now, we know that for every agent, we have $\text{share}(\pi, i) \geq \text{fairshare}(i) = b/n$. Since $\sum_{i \in \mathcal{N}} \text{share}(\pi, i) = c(\pi)$ and π is feasible, we must have $\text{share}(\pi, i) = b/n$ for all $i \in \mathcal{N}$. Overall, we have $\sum_{p \in \mathcal{P}} \gamma_i(p) = \text{share}(\pi, i) = b/n \leq \alpha$, so condition C3 also is satisfied. In addition, we have $\sum_{i \in \mathcal{N}} \gamma_i(p) = \sum_{i \in \mathcal{N}} \text{share}(\{p\}, i) = c(p)$. Condition C4 is thus immediately satisfied. Finally, as we have for every agent $\sum_{p \in \mathcal{P}} \gamma_i(p) = \text{share}(\pi, i) = b/n = \alpha$, condition C5 is vacuously satisfied. \square

Next, we consider the relation between the weaker share-based notions and priceability. The following shows that there are no other implications between priceability and share-based fairness, even if we assume that agents approve of enough projects.

Proposition 14. Local-FS, FS-1, and EJS do not imply priceability, even if $\text{fairshare}(i) = b/n$ for every agent $i \in \mathcal{N}$. Vice versa, priceability does not imply Local-FS or EJS-1, even if $\text{fairshare}(i) = b/n$ for every agent $i \in \mathcal{N}$ and the agents have an allowance of at least b/n .

PROOF. Consider an instance with two projects with $c(p_1) = 3$ and $c(p_2) = 2$, $b = 3$ and two agents such that $A_1 = \{p_1\}$ and $A_2 = \{p_2\}$. Then $\{p_1\}$ satisfies FS-1 and Local-FS as we have $\text{share}(\{p_1, p_2\}, i) > \text{fairshare}(i)$ for both agents i . Moreover, EJS is trivially satisfied, as there are no cohesive groups. On the other hand, for $\{p_1\}$ to be priceable, each agent must receive an allowance of 3. In this case, the fact that p_2 is not selected is a contradiction to C5 as $p_2 \in A_2$ and 2 has more than $c(p_2)$ unspent allowance.

Now consider the instance with four projects such that $c(p_1) = c(p_2) = 8$ and $c(p_3) = c(p_4) = 5$ and $b = 20$. There are two

agents with ballots $A_1 = \{p_1, p_2\}$ and $A_2 = \{p_2, p_3, p_4\}$. The bundle $\{p_1, p_2\}$ is priceable: consider the following contributions with an allowance of 10 per agent: $\gamma_1(p_1) = 8$ and $\gamma_1(p) = 0$ for $p \in \{p_2, p_3, p_4\}$; $\gamma_2(p_2) = 8$ and $\gamma_2(p) = 0$ for $p \in \{p_1, p_3, p_4\}$. However, $\{p_1, p_2\}$ does not satisfy Local-FS as $share(\{p_1, p_2\} \cup \{p_3\}, 2) = 9 < 10 = fairshare(2)$. Moreover, 2 is $\{p_3, p_4\}$ -cohesive but we have $share(\{p_1, p_2\} \cup \{p\}, 2) = 9 < 10 = share(\{p_3, p_4\}, 2)$ for any $p \in \{p_3, p_4\}$. Hence, $\{p_1, p_2\}$ also does not satisfy EJS-1. \square

However, Local-FS, EJS-1, and priceability are compatible in the sense that there always exists a bundle satisfies all three, namely the output of MES_{share} . This follows directly from Theorem 4, Proposition 8, and the fact that MES is priceable for every utility function, as was shown by Peters et al. [19]. It remains open whether FS-1, EJS, and Local-EJS are compatible with priceability in this sense.

6 APPROACHING FAIR SHARE IN PRACTICE

As we saw in Section 3, there exist PB instances for which it is impossible to give every agent their fair share. In this section we report on an experimental study aimed at understanding how serious a problem this is. Our study is twofold. We first investigate how close to fair share we can get. In a second experiment, we quantify how close to this optimal value certain known PB rules get.

For these experiments we use data from Pabulib [24], an online collection of real-world PB datasets. To be more precise, we used all instances from Pabulib with up to 65 projects, except for trivial instances, where either no project or the set of all projects are affordable. Three instances have been additionally omitted for the first experiment due to very high compute time. A total of 353 PB instances are covered by our analysis.

6.1 Optimal Distance to Fair Share

We propose two ways to measure how close to FS a given budget allocation is. The first one is the *average capped fair share ratio*: For every agent i with approval ballot A_i we divide their actual share by their fair share, capped at 1 in case they get more than their fair share, and take the average of this ratio over all agents:

$$\frac{1}{n} \cdot \sum_{i \in N} \min\left(\frac{share(\pi, i)}{fairshare(i)}, 1\right).$$

Our second measure is the *average L_1 distance to FS*, measuring, for every agent i , the absolute value of the difference between their actual share and their fair share:

$$\frac{1}{n} \cdot \sum_{i \in N} |share(\pi, i) - fairshare(i)|.$$

For each PB instance we computed via *integer linear programs* budget allocations yielding the optimal average capped fair share ratio and L_1 distance to FS. Moreover, to better understand what might cause an instance not to admit a good solution, we also considered different ways of preprocessing the instances by removing “problematic” projects:

- **Threshold:** We remove any project that is not approved by at least $x\%$ of agents. We considered $x = 1\%$, 5% , and 10% .
- **Cohesiveness:** We remove any project p such that its supporters do not deserve enough money to buy the project, *i.e.*, such that $\frac{|i \in N | p \in A_i|}{n} b < c(p)$.

Threshold preprocessing removes under 10% of projects for a threshold of 1%, around 10–20% for a threshold of 5%, and around 20–30% for a threshold of 10%. Cohesiveness preprocessing removes between 30% (for the largest instances) and 70% of projects (for the smallest instances)

Let us now turn to our results, presented in Figure 2. We draw the following conclusions. Without preprocessing, we can provide agents on average between 45% (for small instances) and 75% (for larger instances) of their fair share, albeit with a lot of variation. Furthermore, we can typically guarantee an L_1 distance to FS of 50% of the worst case distance. Interestingly, preprocessing helps when using the cohesiveness condition, but not with the threshold condition. Note that we do not wish to advocate preprocessing as a method to make budget decisions in practice. Rather, we use it as a way of checking whether the failure to guarantee fair share is due to the specific structure of real-life PB instances and whether similar instances ‘nearby’ might be significantly better behaved. Our experimental findings suggest that this is not the case, and that guaranteeing fair share simply is very hard across a wide range of instances. Note that, across all instances, for only one instance—with 3 projects and 198 voters—we were able to satisfy FS.

We also investigated approximations of the average capped fair share ratio. Specifically, for a number of different given approximation ratios $\alpha \in (0, 1]$, we replaced the fair share by $\alpha \cdot fairshare(i)$ in the definition. Results indicates that moving from $\alpha = 1$ to $\alpha = 0.2$, has a very small effect on the optimum value (around 10% better for $\alpha = 0.2$). We also interpret this result as stating that FS is structurally hard to satisfy.

6.2 Distance to Fair Share of Common PB Rules

We now turn to our second experiment: how close to fair share are the outputs of known PB rules in practice. We will consider the following rules: MES_{share} , MES_{card} [19], MES_{cost} [19],⁵ sequential Phragmén [17], and greedy approval [5]. Due to space constraints, we omit the definitions of the rules here and instead refer the reader to relevant references and the full version on arXiv [18].

For every PB instance and every PB rule, we compute the outcome returned by the rule and assess how close to the optimal value it is in terms of both the average capped fair share ratio and average L_1 distance to FS. Results are presented in Figure 3.

The first striking observation is that greedy approval is performing extremely well under the capped fair share ratio measure. This is particularly surprising given how oblivious to the structure of the profile greedy approval is. We postulate that this result is due to the high difference in the percentage of the budget used by the different rules: MES rules use around 40% of the budget on average, while greedy approval and sequential Phragmén use around 90% of the budget. Since using more budget can only improve the average capped fair share ratio, this is the most likely explanation for the good performance of greedy approval compared to MES. There are no standard ways to extend MES budget allocation in the literature (for PB). It is thus hard to compare rules based on the average capped fair share ratio the achieve.

⁵We write MES_{card} and MES_{cost} for the rule MES [19] used with utility functions $u_i(p) := 1$ and $u_i(p) := c(p)$ for all $i \in N$ and $p \in \mathcal{P}$, respectively.

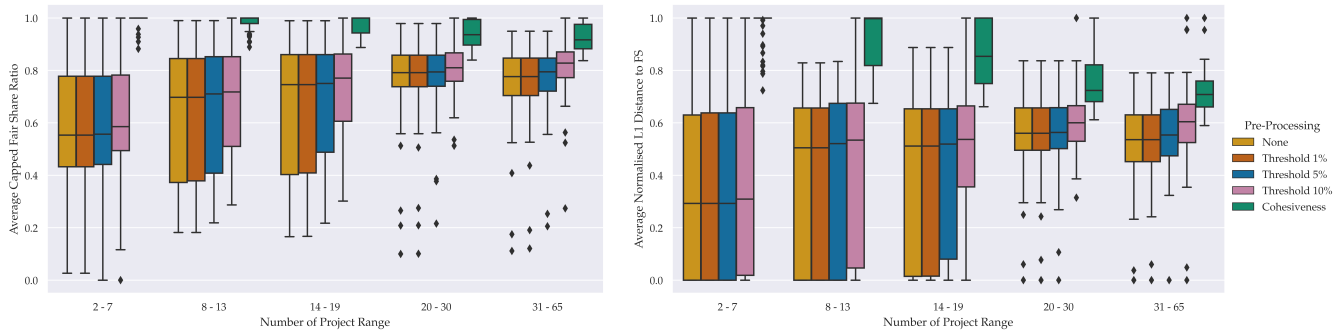


Figure 2: Average capped fair share ratio (left) and L_1 distance to FS (right) for Pabulib instances. For the latter we actually plot $1 - 1/n \cdot \sum_{i \in N} \frac{|\text{share}(\pi, i) - \text{fairshare}(i)|}{\text{fairshare}(i)}$ to obtain a normalised value for which 1 is the best.* Each range (for a number of projects) shown on the x -axis contains between 60 and 80 instances.

*Note that the empty budget allocation provides an L_1 distance to FS of $\text{fairshare}(i)$ for all $i \in N$. Normalising the L_1 distance with $\text{fairshare}(i)$, thus ensures that we display the optimal L_1 distance to FS achieved with respect to the worst case.




Figure 3: Average capped fair share ratio (left) and average L_1 distance to FS (right) for different rules on Pabulib instances. Results are normalised by the optimum value achievable in each instance, giving a score between 0 and 1 where 1 is the best.

Interestingly, the average L_1 distance to FS does not suffer this drawback. Indeed, since it also penalises rules that provide agents more than their fair share, spending more is not always better. Interpreting the results of Figure 3 in this light, we conclude that MES rules perform better than sequential Phragmén in terms of equality of resources. Interestingly, MES_{cost} performs slightly better than $\text{MES}_{\text{share}}$. It thus provides both good experimental results in terms of equality of resources and strong representation guarantees [19].

7 CONCLUSION

In this paper, we have proposed to evaluate the fairness of a participatory budgeting decision by using the share of a voter as a measure of the *resources* spent on satisfying the needs of voters rather than using the (assumed) *satisfaction* each voter might derive from an allocation. Our results suggest that this is an interesting measure of fairness that deserves further attention.

In summary, we have seen that perfect fairness in the sense of fair share is not always achievable, as is usually the case in PB, due to the discrete nature of the process. More surprisingly, our experiments show that, in practice, it is often even impossible to

achieve outcomes that are close to fair share. Nevertheless, we were able to relax the requirements of fair share to define several desiderata that can always be satisfied. Using these criteria, we are able to identify $\text{MES}_{\text{share}}$ as the polynomial-time-computable PB rule that is most equitable in terms of resources, as it satisfies both Local-FS as well as EJS-1. This result is strengthened by our experimental evaluation that shows that $\text{MES}_{\text{share}}$ selects bundles that are close to optimal with regards to distance to FS.

It is worth noting, however, that MES_{cost} performs slightly better than $\text{MES}_{\text{share}}$ in our experiments, hinting at an interesting connection between share- and satisfaction-based fairness notions. Exploring whether a meaningful compromise between these two types of fairness can be achieved is important future work, even though, such a compromise would have to be significantly weaker than EJS and EJR (in view of, e.g., Example 3).

Another important open question is, whether a (natural) PB rule exists that satisfies EJS as well as FS and FS-1 whenever they are satisfiable. Such a rule would necessarily be intractable, but it would provide strong axiomatic fairness guarantees for PB instances which are small enough to allow for the computation of rules with exponential runtime.

ACKNOWLEDGMENTS

We would like to thank the anonymous reviewers, both at this conference and at MPREF-2022, for their valuable and insightful comments. This research was funded in part by the Austrian Science Fund (FWF) under grant numbers P3189 and J4581 as well as the Dutch Research Council (NWO) under grant number 639.023.811.

REFERENCES

- [1] 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.
- [2] 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)*.
- [3] 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)*.
- [4] 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)*.
- [5] Haris Aziz and Nisarg Shah. 2020. Participatory Budgeting: Models and Approaches. In *Pathways between Social Science and Computational Social Science: Theories, Methods and Interpretations*, Tamás Rudas and Gábor Péli (Eds.). Springer-Verlag.
- [6] Charles Blackorby, Walter Bossert, and David Donaldson. 2002. Utilitarianism and the Theory of Justice. In *Handbook of Social Choice and Welfare*, Kenneth J. Arrow, Amartya K. Sen, and Kotaro Suzumura (Eds.). Elsevier.
- [7] Yves Cabannes. 2004. Participatory Budgeting: A Significant Contribution to Participatory Democracy. *Environment and Urbanization* 16, 1 (2004), 27–46.
- [8] Ronald Dworkin. 1981. What is Equality? Part 1: Equality of Welfare. *Philosophy & Public Affairs* 10, 3 (1981), 185–246.
- [9] Ronald Dworkin. 1981. What is Equality? Part 2: Equality of Resources. *Philosophy & Public Affairs* 10, 4 (1981), 283–345.
- [10] 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)*.
- [11] 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.
- [12] Martin Fürer and Huiwen Yu. 2011. Packing-Based Approximation Algorithm for the k -Set Cover Problem. In *Proceedings of the 22nd International Symposium on Algorithms and Computation (ISAAC)*.
- [13] 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.
- [14] John R. Hicks and Roy G. D. Allen. 1934. A Reconsideration of the Theory of Value. Part I. *Economica* 1, 1 (1934), 52–76.
- [15] 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)*.
- [16] Martin Lackner and Piotr Skowron. 2023. *Multi-Winner Voting with Approval Preferences*. Springer-Verlag.
- [17] Maaïke Los, Zoé Christoff, and Davide Grossi. 2022. Proportional Budget Allocations: Towards a Systematization. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*.
- [18] Jan Maly, Simon Rey, Ulle Endriss, and Martin Lackner. 2023. Fairness in Participatory Budgeting via Equality of Resources. *arXiv preprint arXiv:2205.07517* (2023). Full version of this paper.
- [19] Dominik Peters, Grzegorz Pierczynski, and Piotr Skowron. 2021. Proportional Participatory Budgeting with Additive Utilities. In *Proceedings of the 35th Annual Conference on Neural Information Processing Systems (NeurIPS)*.
- [20] Dominik Peters and Piotr Skowron. 2020. Proportionality and the Limits of Welfareism. In *Proceedings of the 21st ACM Conference on Economics and Computation (ACM-EC)*.
- [21] Jack Robertson and William Webb. 1998. *Cake-Cutting Algorithms*. A. K. Peters.
- [22] Anwar Shah (Ed.). 2007. *Participatory Budgeting*. The World Bank.
- [23] Gogulapati Sreedurga, Mayank Ratan Bhardwaj, and Y. Narahari. 2022. Maxmin Participatory Budgeting. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*.
- [24] Dariusz Stolicki, Stanisław Szufa, and Nimrod Talmon. 2020. Pabulib: A Participatory Budgeting Library. *arXiv preprint arXiv:2012.06539* (2020).
- [25] Nimrod Talmon and Piotr Faliszewski. 2019. A Framework for Approval-Based Budgeting Methods. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*.
- [26] Brian Wampler, Stephanie McNulty, and Michael Touchton. 2021. *Participatory Budgeting in Global Perspective*. Oxford University Press.