The Method of Equal Shares for Participatory Budgeting

Dominik Peters, Grzegorz Pierczyński, and Piotr Skowron

Abstract

Participatory budgeting involves a group of voters collectively deciding which projects should be funded using a common budget. In this paper, we propose a new voting rule for participatory budgeting, called the Method of Equal Shares. This method assigns each voter an equal portion of the budget, which can only be spent on projects they vote for. The rule then iterates through the projects and implements those that can be funded using the combined budget shares of their supporters, with carefully chosen policies about how to prioritize between projects and how to divide a project's cost between its supporters. The Method of Equal Shares is compatible with any ballot input format that can be interpreted as additive utilities, including approval and score voting. We argue that this method provides proportional representation, so that groups with similar preferences have an influence on the outcome that is proportional to the group size. We formalize this by proving that the method satisfies an axiom known as Extended Justified Representation up to 1 Project (EJR-1).

Note: This is a short version of our paper "Proportional Participatory Budgeting with Additive Utilities", published at NeurIPS'21, based on the October 2022 arXiv-version of the paper (arXiv:2008.13276) which contains some new and updated material (including a detailed example and a new proof of the main result).

1 Introduction

A growing number of cities now uses Participatory Budgeting (PB) to decide how to spend a part of their budgets [9, 21]. Through a voting system, PB allows the residents of a city to decide which projects will be funded by the city authorities [18]. This increases civic involvement in government, by increasing the number of issues that are decided by democratic vote, and by allowing residents to submit their own project proposals [8, 2].

To count the votes, most cities use a variant of a simple protocol: Each voter is allowed to vote for a certain number of project proposals. Then, the projects with the highest number of votes are funded, until the budget limit is reached. While simple and intuitive, this voting rule gives too much voting power to pluralities: cohesive groups (if they are larger than other groups) can dictate the entire outcome, leading to a *tyranny of the largest minority*.

To see this, consider Circleville, a fictional city divided into four districts. A map of the city is shown in Figure 1. The districts have similar sizes, but Northside has the largest population. Suppose \$400k have been allocated to PB (\$1 per person), and suppose that all the project proposals

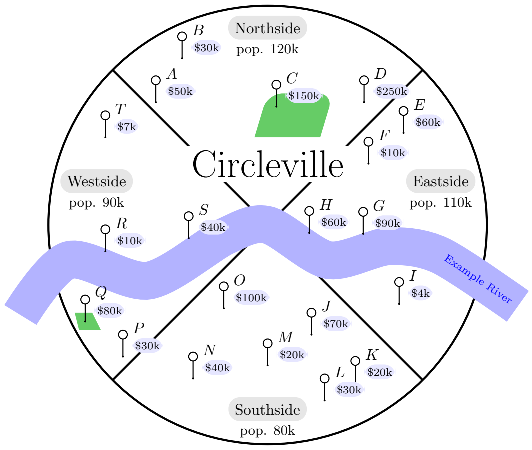

Figure 1: Map of Circleville, showing the locations and costs of the PB project proposals.

are of a local character (such as school renovations), and that residents only vote for projects that concern their own district. For example, every Northside resident will cast votes for projects A, B, C, and D, but no one else votes for these. Because Northside is the most populous district, the Northside projects will all receive the highest number of votes, and the voting rule will spend the entire budget on Northside projects (via projects C and D). The 280k residents of the other three districts are left empty-handed.

To circumvent this obvious issue, many cities have opted to hold separate elections for each district. The budget is divided in advance between the districts (e.g., in proportion to their number of residents), each project is assigned to a district, and voters only vote in their local election. While this avoids spending the entire budget in Northside, this fix introduces many other problems. For example, projects on the boundary of two districts (such as A and P) need to be assigned to one of them. Residents of the other district may be in favor of the boundary project, but cannot vote for it. Thus boundary projects are less likely to be funded, even if they would be more valuable overall. Similarly, projects without a specific location that benefit the entire city are difficult to handle. Also, interest groups that are not geographic in nature will be underserved; for instance, parents across the city might favor construction of a large playground (project C), but with separate district elections, parents cannot form a voting block. Similarly, bike riders across the city cannot express their joint interest in the construction of a bike trail along Example River (projects R, S, H, and G).

To solve these problems, we could hold a single city-wide election, but use a voting system which ensures that money is spent *proportionally* in the sense that for each group of voters with similar interests, the voting systems spends an amount of money proportional to the group's size. These groups of voters should not be pre-specified, but automatically identified based on the votes [5, 10]. Proportionality should make it safe to run a single city-wide election, and enjoy the efficiency benefits of avoiding separate district elections.

We introduce a voting rule for participatory budgeting that we call the *Method of Equal Shares*, generalizing a previously-proposed committee election method [15]. The method is based on the simple idea that to give each voter roughly equal influence on the outcome, a good first step is to split the overall budget into equal parts, assign each voter one of the parts, and stipulate that the method is only allowed to spend this part of the budget on projects that the voter has voted for. There are many ways to instantiate this generic scheme. Roughly speaking, the Method of Equal Shares first tries to fund projects that have the highest total reported utility divided by cost (if those projects can be afforded using the remaining budget shares of their supporters) and divides a project's cost so that voters contribute an amount that is proportional to their reported utility for the project.

We define the method so that its input consists of a non-negative utility number for each project and each voter. We implicitly assume that voter satisfaction can be measured as the sum of the utility numbers of the projects selected by the voting rule. Thus, the method can be used together with any input format that can be converted into an additive utility presentation. This includes almost all of the input formats actually used in practice: some cities allow voters to distribute a fixed number of points between projects (for example, 10), which immediately gives a utility representation; others allow voters to rank projects which can be converted into utilities using Borda scores. Most cities use some variant of approval voting, which can be converted by taking utility 1 for approved projects and utility 0 for non-approved projects, though there are other conversion possibilities such as taking the cost of the project as the utility in case it is approved ("cost utilities") [7].

Our main result is that the Method of Equal Shares satisfies an axiom called *extended justified representation up to one project* (EJR-1). Intuitively, this axiom requires that no group of voters with common interests is underserved. It is a generalization of the well known EJR axiom from the literature on approval-based committee voting [3, 12]. We believe that this axiom is a good formalization of the idea of proportional representation. We end the paper by a discussion of subsequent literature and several open questions.

2 Preliminaries

For each $t \in \mathbb{N}$, write $[t] = \{1, 2, \dots, t\}$. An *election* is a tuple $(N, C, b, \mathrm{cost}, \{u_i\}_{i \in N})$, where:

- $N = [n]$ and $C = \{c_1, \dots, c_m\}$ are the sets of *voters* and *candidates* (or *projects*).
- $b \in \mathbb{Q}_{\geqslant 0}$ is the available *budget*.
- $\mathrm{cost}\colon C \to \mathbb{Q}_{\geqslant 0}$ specifies the *cost* that needs to be paid if a candidate is selected. For each $T \subseteq C$, we write $\mathrm{cost}(T) = \sum_{c \in T} \mathrm{cost}(c)$ for the total cost of T.
- For each voter $i \in N$, the function $u_i\colon C \to \mathbb{Q}_{\geqslant 0}$ defines i's additive *utility function*.[1] If a set $T \subseteq C$ of candidates is implemented, i's overall utility is $u_i(T) = \sum_{c \in T} u_i(c)$. For a subset $S \subseteq N$ of voters, we further write $u_S(T) = \sum_{i \in S} \sum_{c \in T} u_i(c)$ for the total utility enjoyed by S if T is implemented. We assume that $u_N(c) > 0$ for each $c \in C$, so that every candidate is assigned positive utility by at least one voter.

A subset of candidates $W \subseteq C$ is *feasible* if $\mathrm{cost}(W) \leqslant b$. Our goal is to choose a feasible subset of candidates, which we call an *outcome*, based on voters' utilities. An *aggregation rule* (or, in short, a *rule*) is a function $\mathcal{R}$ that for each election E selects a feasible outcome $W = \mathcal{R}(E)$ called the *winning outcome*.[2] We call this the *general PB model*.

There are two interesting special cases of our model:

Committee elections. In this case, the budget is an integer, $b \in \mathbb{N}$ (the committee size) and each candidate costs 1. Then W is an outcome if and only if $|W| \leqslant b$. In this special case we also refer to outcomes as *committees*, and we say that the election satisfies the *unit cost assumption*.

Approval utilities. In this case, for each $i \in N$ and $c \in C$ it holds that $u_i(c) \in \{0, 1\}$. The *approval set* of voter i is $A(i) := \{c \in C\colon u_i(c) = 1\}$, and we say that i *approves* candidate c if $c \in A(i)$. If $c \in A(i) \cap W$, we say that c is a *representative* of i in W.

Often we combine of these special cases, and study *approval-based committee elections*.

3 The Method of Equal Shares

Peters and Skowron [15] introduced an aggregation rule for approval-based committee elections that they called Rule X. In that setting (approval-based committee elections) the rule satisfies a combination of appealing proportionality properties. Here, we extend it to the more general model of participatory budgeting, that is, to the model with arbitrary costs and utilities. We will call this rule the *Method of Equal Shares*, because it works by dividing the available budget equally between voters, and (to the extent it is possible) sharing the cost of each project equally between the voters who approve the project. For brevity, we often refer to the rule simply as *Equal Shares*.

Definition 1 (Method of Equal Shares)**.** Each voter is initially given an equal fraction of the budget, i.e., each voter is given b/n monetary units. We start with an empty outcome $W = \emptyset$ and sequentially add candidates to W. To add a candidate c to W, we need the voters to pay for c. For each selected candidate $c \in W$ we write $p_i(c)$ for the amount that voter i pays for c; we need that $\sum_{i \in N} p_i(c) = \mathrm{cost}(c)$. We write $p_i(W) = \sum_{c \in W} p_i(c) \leqslant b/n$

[1] The rules we propose will be invariant under rescaling all utility functions by a common factor, so without loss of generality we could assume that $u_i(c) \in [0, 1]$ for all $i \in N$ and $c \in C$. The rules will not be invariant under rescaling by different factors, or by shifting utilities.

[2] Sometimes there are ties. For the results of this paper it does not matter how these ties are broken.

for the total amount i has paid so far; then write $b_i = {}^b\!/n - p_i(W)$ for the amount of money that i has left (clearly, in the first round $b_i = {}^b\!/n$ for all $i \in N$). For $\rho \geqslant 0$, we say that a candidate $c \notin W$ is *ρ-affordable* if

$$\sum_{i \in N} \min\left(b_i, u_i(c) \cdot \rho\right) = \text{cost}(c).$$

If no candidate is ρ-affordable for any ρ, Equal Shares terminates and returns W. Otherwise it selects a candidate $c \notin W$ that is ρ-affordable for a minimum ρ. Individual payments are given by

$$p_i(c) = \min\left(b_i, u_i(c) \cdot \rho\right) \qquad \square$$

Intuitively, when Equal Shares adds a candidate c, it asks voters to pay an amount proportional to their utility $u_i(c)$ for c; in particular, the cost per unit of utility is ρ. If a voter does not have enough money, the rule asks the voter to pay all the money the voter has left, which is $b_i = {}^b\!/n - p_i(W)$. Throughout the execution of Equal Shares, the value of ρ for the selected candidates increases. Thus, candidates are added in decreasing order of utility per dollar.[3]

Special Case: Approval Utilities

In the special case where voters have approval utilities ($u_i(c) \in \{0, 1\}$ for all $i \in N$ and $c \in C$), the definition of Equal Shares becomes somewhat more straightforward. At each step, a candidate $c \notin W$ is ρ-affordable if and only if the cost of c can be covered by the voters approving c in such a way that the maximum payment of any voter is ρ. Voters who have less than ρ left spend all their remaining money, and the other voters pay exactly ρ. This way, the cost is shared as equally as possible among supporters of a project, which motivates the name Method of Equal Shares. See also the following example for an illustration.

Example

Let us go through a small example with 5 projects, 10 voters $N = \{i_1, \ldots, i_{10}\}$ with approval utilities, and a budget of $b = \$100$. Costs and utilities are shown in the table below.

	cost	i_1	i_2	i_3	i_4	i_5	i_6	i_7	i_8	i_9	i_{10}
Project 1	\$36	1	1	1	1	1	0	0	0	0	0
Project 2	\$36	1	1	1	1	1	1	0	0	0	0
Project 3	\$25	1	1	1	1	1	0	0	0	0	0
Project 4	\$24	0	0	0	1	1	1	1	0	0	0
Project 5	\$24	0	0	0	0	0	1	1	1	1	0

Sharing the available budget equally among the voters, everyone starts with \$10. We can now calculate, for each project, the value ρ such that the project is ρ-affordable. As discussed above, because we have approval utilities, this just entails spreading the cost as equally as possible among project supporters.

[3] An exception are voters who do not have enough money to contribute $u_i(c) \cdot \rho$ and instead contribute all of their remaining money (which could be nothing); those voters obviously get higher utility per dollar.

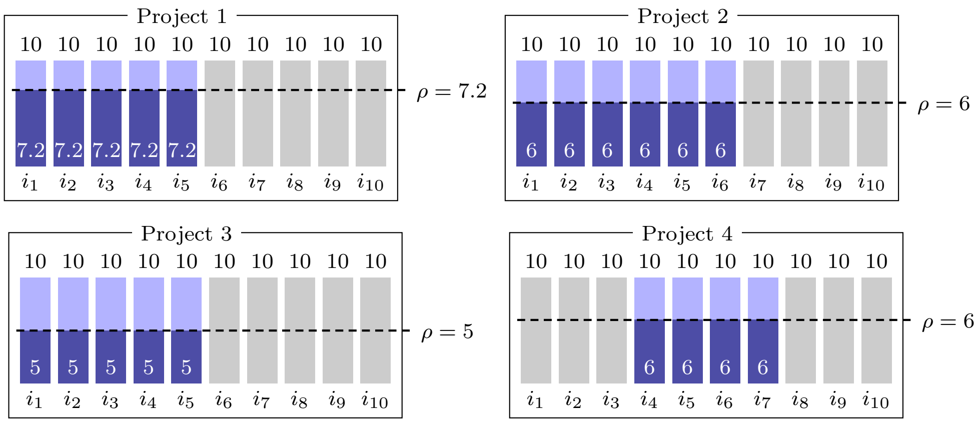

We do not include a picture for Project 5 because it is similar to Project 4 (same cost, same number of supporters, and thus also $\rho = 6$). Equal Shares will choose to implement Project 3, since it has the lowest value of ρ. Comparing the projects at this stage, we can identify a few principles:

- *More supporters is better*: Project 2 costs the same as Project 1, but it has more supporters. Thus its cost can be spread more thinly (i.e., it is affordable for a lower ρ), which will lead Equal Shares to prefer Project 2 over Project 1.
- *Cheaper is better*: Project 3 has the same supporters as Project 1, but it costs less. Again this leads to a lower ρ, and Equal Shares will prefer Project 3 over Project 1.

In fact, the first project selected by Equal Shares will always be one that (i) can be afforded by its supporters and (ii) subject to this, it has the highest number of approvers divided by cost.

After Equal Shares selects Project 3, voters i_1 through i_5 each pay \$5 for it. We can note that Projects 1 and 2 (both costing \$36) are now not affordable anymore: the supporters of Project 1 have \$25 left, and the supporters of Project 2 have \$35 left. However, Projects 4 and 5 are still affordable.

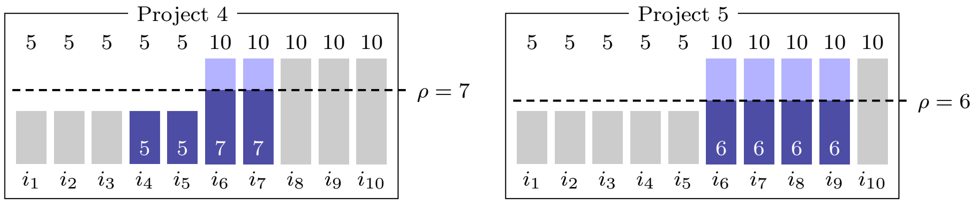

Even though Projects 4 and 5 have the same number of supporters, we can see that they now induce different ρ-values. This points to another principle:

- *Richer supporters is better*: Among otherwise identical projects, those whose supporters have more money left are preferred by Equal Shares. This makes sense, because voters with more money have not yet had their preferences satisfied as much in prior rounds.

Thus, Equal Shares next selects Project 5, and voters i_6 through i_9 each pay \$6 for it. We can check that at this point, no projects are affordable (same argument as before for Projects 1 and 2, and the supporters of Project 4 only have \$18 left). Thus, Equal Shares terminates and selects the winning outcome $W = \{\text{Project 3}, \text{Project 5}\}$. Note that $\text{cost}(W) = \$49$ while $b = \$100$, and thus a large part of the budget was not spent by Equal Shares. Indeed, we could add any one of the remaining three projects to the outcome and still satisfy the budget constraints. For strategies that "complete" the result of Equal Shares, see the next section.

Exhaustiveness

A basic and very desirable efficiency notion is *exhaustiveness*, which requires that a voting rule spends its entire budget. Of course, due to the discrete model, we cannot guarantee that the rule will spend exactly b; however, we can require that no additional project is affordable.

Definition 2 (Exhaustiveness [5])**.** An election rule $\mathcal{R}$ is *exhaustive* if for each election instance E and each non-selected candidate $c \notin \mathcal{R}(E)$ it holds that $\mathrm{cost}(\mathcal{R}(E) \cup \{c\}) > b$.

Notably, Equal Shares fails to be exhaustive. For example, suppose that $b = 1$, that we have two voters and two candidates, such that v_1 approves $\{c_1\}$ and v_2 approves $\{c_2\}$. Suppose that both candidates have cost 1. Then Equal Shares returns $W = \emptyset$. Equal Shares refuses to fund either candidate, because neither has enough support to cover its cost.

In some contexts, it may actually be a desirable feature of Equal Shares that it is not exhaustive, especially if unspent budget can be used in other productive ways (such as in next year's PB election). In other situations, unspent budget may not be reusable, such as when the budget comes from a grant where unspent money needs to be returned (and the relevant decision makers do not obtain value from the grant-maker's alternative activities), or when the 'budget' is time (for example, when we use PB to plan activities for a day-long company retreat). In such situation, one might prefer an exhaustive rule, or at least one that spends a large fraction of the available budget.

Completion by varying the budget Peters et al. [16] suggested a method of completing outcomes returned by Equal Shares, which we call "*completion by varying the budget*". In this variant, we evaluate the method with a budget $b' \geqslant b$ that is higher than the actually available budget b. We increase the value of b' gradually (each time by some fixed amount), and after each increase we recompute the outcome from scratch using Equal Shares. We stop when we reach an exhaustive outcome or when the next increase of b' would cause us to exceed the original budget b. The selected outcome is typically exhaustive, but formally there is no guarantee of this. For a more detailed discussion, see also Peters et al. [16].

An advantage of completion by varying the budget is that the voting power (in the form of the amount of budget initially distributed) continues to be equally shared, and in particular the outcome is priceable. It also mirrors the way in which the D'Hondt apportionment divisor method "completes" lower quotas. However, note that this method is not necessarily a "completion", because the final outcome need not be a superset of the Equal Shares outcome computed with the real budget b.

Completion by perturbation Since we have generalized Equal Shares to work for general additive valuations, there is another way for us to make it exhaustive. Recall that Equal Shares fails to be exhaustive in situations where the remaining projects' supporters do not have sufficient funds left. However, in elections where $u_i(c) > 0$ for all $i \in N$ and $C \in C$, every voter supports every candidate, and thus this problem never occurs. One can prove that Equal Shares is exhaustive when run on profiles of this type [17].

Thus, we can make Equal Shares exhaustive by perturbing the input utilities so that all utility values are positive. Specifically, for a small $\varepsilon > 0$ ($\varepsilon \ll \min_{u_i(c)>0} u_i(c)$), and for each $i \in N, c \in C$ such that $u_i(c) = 0$ in the initial instance we set $u_i^\varepsilon(c) = \varepsilon$. Then we run Equal Shares on the modified instance $\{u_i^\varepsilon\}_{i \in N}$. We call this strategy "*completion by perturbation*".

Other completions Users of Equal Shares could also choose to use the budget not spent by Equal Shares to further other goals. For example, they could use "*completion by utilitarian greedy*" which selects unselected projects in order of their total utility $\sum_{i \in N} u_i(c)$, mirroring the voting rule commonly used in practice. Peters and Skowron [15] proposed to complete Equal Shares with the use of Phragmén's sequential rule, but this only works for approval-based utilities.

4 Extended Justified Representation (EJR)

The formal notion of proportionality that we examine is Extended Justified Representation (EJR). This axiom was first proposed for approval-based committee elections [3]. Even for the special case of approval-based committee elections, only few rules are known to satisfy EJR [3, 4, 15], but the Method of Equal Shares is one of them. In this section, we introduce a generalization of EJR to the PB model and show that our rule continues to satisfy EJR (at least up to one project).

We first recall the definition of EJR for approval-based committee elections. Intuitively, this axiom ensures that every large enough group of voters whose approval sets have a large enough intersection must obtain a fair number of representatives. For example, if a group of voters forms an α-fraction of the whole population and if this group agrees on sufficiently many candidates, then it should be allowed to decide about an α-fraction of the elected candidates. Formally, this is achieved by excluding the possibility that each member of the group approves less than $\lfloor \alpha k \rfloor$ elected candidates.

Definition 3 (Extended Justified Representation for approval-based committee elections)**.** We say that a group of voters S is *ℓ-cohesive* for $\ell \in \mathbb{N}$ if $|S| \geqslant \ell/k \cdot n$ and $|\bigcap_{i\in S} A(i)| \geqslant \ell$.

A rule $\mathcal{R}$ satisfies *extended justified representation* if for each election instance E, each $\ell \in \mathbb{N}$, and each ℓ-cohesive group S of voters, there exists a voter $i \in S$ such that $|A(i) \cap \mathcal{R}(E)| \geqslant \ell$.

At first sight it is unintuitive that we only require that at least one voter obtains ℓ representatives. However, the strengthening of EJR that requires each member of S to obtain ℓ representatives is impossible even on very small instances [3]. Still, even with only the at-least-one guarantee, EJR has plenty of bite, in particular implying that the members of a cohesive group have high utility on average [4, 19, 15].

We will generalize this axiom to the PB model. To warm up, let's first relax the unit cost assumption, but stay in the approval-based setting. Then EJR should state the following.

Definition 4 (Extended Justified Representation for approval-based elections)**.** We say that a group of voters S is T-cohesive for $T \subseteq C$ if $|S|/n \geqslant \mathrm{cost}(T)/b$ and $T \subseteq \bigcap_{i\in S} A(i)$. A rule $\mathcal{R}$ satisfies *extended justified representation* if for each election instance E, each $T \subseteq C$, and each T-cohesive group S of voters, there exists a voter $i \in S$ such that $|A(i) \cap \mathcal{R}(E)| \geqslant |T|$.

Thus, cohesiveness now requires that the group S can identify a collection of projects T that they all approve and that is affordable with their fraction of the budget ($|S|/n \geqslant \mathrm{cost}(T)/b$). Note that voters $i \in S$ obtain utility $u_i(T) = |T|$ from T; EJR requires that at least one member of S must attain this utility in the election outcome.

Generalizing EJR beyond approvals is more difficult, because we cannot talk about a candidate being approved by all members of S. Instead, we quantify cohesion by calculating the minimum utility that any member of S assigns to each project in T. (For approvals, this value is 1 if the candidate is approved by all members in S and 0 otherwise.)

Definition 5 (Extended Justified Representation)**.** A group of voters S is (α, T)-*cohesive*, where $\alpha\colon C \to [0;1]$ and $T \subseteq C$, if $|S|/n \geqslant \mathrm{cost}(T)/b$ and if $u_i(c) \geqslant \alpha(c)$ for all $i \in S$ and $c \in T$. A rule $\mathcal{R}$ satisfies *extended justified representation* if for each election instance E, each $\alpha\colon C \to [0;1]$, $T \subseteq C$, and each (α, T)-cohesive group of voters S, there exists a voter $i \in S$ such that $u_i(\mathcal{R}(E)) \geqslant \sum_{c\in T} \alpha(c)$.

Again, an (α, T)-cohesive group of voters S can propose the projects in T, since they are affordable with S's share of the budget. The values $(\alpha(c))_{c\in T}$ denote how much the coalition S agrees about the desirability of the projects in T. In particular, we have $u_i(T) \geqslant \sum_{c\in T}\alpha(c)$ for each $i \in S$. Definition 5 prohibits any outcome in which every voter in S gets utility strictly lower than $\sum_{c\in T}\alpha(c)$; hence there must exist $i \in S$ such that $u_i(\mathcal{R}(E)) \geqslant \sum_{c\in T}\alpha(c)$.

EJR is a demanding property in the PB model. Consider the special case where there is only one voter, $N = \{1\}$. Then any outcome W satisfying EJR must solve the knapsack problem, i.e., it must maximize $\sum_{c\in W} u_1(c)$ subject to the budget constraint, since otherwise an optimum knapsack T witnesses an EJR violation. Because the knapsack problem is weakly NP-hard, this presents a difficulty for a rule to satisfy EJR.[4]

Proposition 1. *Unless P = NP, no aggregation rule that can be computed in strongly polynomial time can satisfy EJR in the general PB model.*

Equal Shares can be computed in strongly polynomial time, and indeed it fails EJR in the general PB model.[5] In fact one can prove that an EJR outcome always exists [17] but we do not know an efficiently computable method that finds such an outcome. However, we can show that Equal Shares selects an outcome that satisfies a mild relaxation of EJR, which requires that EJR holds "up to one project".

Definition 6 (Extended Justified Representation Up To One Project).[6] A rule $\mathcal{R}$ satisfies extended justified representation *up to one project* if for each election instance E and each (α, T)-cohesive group of voters S there exists a voter $i^* \in S$ such that either $u_{i^*}(\mathcal{R}(E)) \geqslant \sum_{c\in T}\alpha(c)$ or for some $c^* \in T$ it holds that $u_{i^*}(\mathcal{R}(E) \cup \{c^*\}) > \sum_{c\in T}\alpha(c)$.

It is worth noting that in the approval-based model, Definitions 5 and 6 are actually equivalent, because the "up to one project" option never applies: Consider an (α, T)-cohesive group of voters S. Since voters' utilities are 0/1, we may assume that for each $c \in T$ we have $\alpha(c) = 1$: if $\alpha(c) > 0$ this is clear; otherwise we can remove c from T without losing cohesiveness. Thus, the cohesiveness condition requires that every voter in S approves every candidate in T. Finally, note that in the approval model, due to the strict inequality, both conditions $u_{i^*}(\mathcal{R}(E)) \geqslant \sum_{c\in T}\alpha(c)$ and $\exists_{c^*\in T}. u_{i^*}(\mathcal{R}(E) \cup \{c^*\}) > \sum_{c\in T}\alpha(c)$ boil down to $|A(i^*) \cap \mathcal{R}(E)| \geqslant \sum_{c\in T}\alpha(c) = |T|$.

Our main result is that the Method of Equal Shares satisfies EJR up to one project in the general PB model. By the previous observation, it hence satisfies EJR in the approval-based model (even when not imposing unit costs), see Table 1.

Theorem 1. *The Method of Equal Shares satisfies EJR up to one project in the participatory budgeting model.*

[4]Knapsack is also weakly NP-hard when assuming that the value of each item equals its weight (this is subset sum), so satisfying EJR is weakly NP-hard even for cost utilities, where every voter assigns a project c utility either 0 or cost(c). Aziz et al. [5, Prop. 3.8] prove a similar result for their BPJR-L notion.

[5]Counterexamples: Suppose the budget is $b = 3$, there is just 1 voter, and two projects c_1 and c_2 with $\mathrm{cost}(c_1) = 1$ and $\mathrm{cost}(c_2) = 3$. Suppose the utilities are $u_1(c_1) = \frac{1}{2}$ and $u_1(c_2) = 1$. Then c_1 is 2-affordable and c_2 is 3-affordable. Thus Equal Shares selects c_1 and the outcome is $W = \{c_1\}$. (Note that this is exhaustive.) But the grand coalition $S = \{1\}$ can propose $T = \{c_2\}$ witnessing a violation of EJR. A counterexample with unit costs: $b = 2$, two voters, three projects, and utilities $u_1(c_1) = u_2(c_1) = 2$, $u_1(c_2) = u_2(c_3) = 3$, and $u_1(c_3) = u_2(c_2) = 0$. Then c_1 is $\frac{1}{4}$-affordable and c_2 and c_3 are both $\frac{1}{3}$-affordable. So c_1 is elected (with both voters paying equal amounts). Now nothing is affordable, and thus $W = \{c_1\}$. But then $S = \{1\}$ and $T = \{c_2\}$ is an EJR violation.

[6]Compared to earlier versions of this paper, this definition of EJR up to one is slightly stronger: we now require that the "extra" project c^* is a member of T while previous versions allowed c^* to be any project, not necessarily in T. As motivation for the new stronger version, consider a setting where there is some project d that is too expensive to ever be affordable, but that provides very high utility to everyone. In those cases, the weaker version of EJR up to one does not provide any guarantees because we can choose $c^* = d$.

	Approval utilities	Additive utilities
Unit costs	EJR	EJR up to one project
General costs	EJR	EJR up to one project†

Table 1: Equal Shares and Extended Justified Representation (see Theorem 1 and Footnote 5). †: Unless P = NP, no strongly polynomial time method (such as Equal Shares) satisfies EJR.

Proof. Let $S \subseteq N$ be a non-empty group of voters, and let $T \subseteq C$ be a proposal with $\mathrm{cost}(T)/b \leqslant |S|/n$. For each $c \in C$, write $\alpha_c = \min_{i \in S} u_i(c)$, and write $\alpha = \sum_{c \in T} \alpha_c$. We assume that $\alpha_c > 0$ for all $c \in T$ (otherwise we can delete c from T). If W is the output of Equal Shares, we will show that there exists a voter $i^* \in S$ such that either $u_{i^*}(W) \geqslant \alpha$, or there is a candidate $c^* \in T$ such that $u_{i^*}(W \cup \{c^*\}) > \alpha$.

In this proof, we will consider three runs of Equal Shares in different variations:

(A) Equal Shares run on the original instance (thus, outputting W).

(B) Equal Shares run so that voters in S are not bound by their budget constraint when paying for candidates in T. To make this formal, in the definition of Equal Shares, we redefine the notion of ρ-affordability so that $c \in T$ is ρ-affordable if

$$\underbrace{\sum_{i \in S} \rho \cdot u_i(c)}_{\text{no budget limit}} + \underbrace{\sum_{i \in N \setminus S} \min\{b_i, \rho \cdot u_i(c)\}}_{\text{with budget limit}} = \mathrm{cost}(c), \tag{1}$$

and $c \in C \setminus T$ is ρ-affordable if

$$\sum_{i \in S} \min\{\underbrace{\max\{b_i, 0\}}_{b_i \text{ may be } < 0}, \rho \cdot u_i(c)\} + \sum_{i \in N \setminus S} \min\{b_i, \rho \cdot u_i(c)\} = \mathrm{cost}(c),$$

with payments defined as these equations suggest (namely, i's payment is the value of the ith term of the sum).

(C) Equal Shares run on a smaller instance where only candidates in T and only voters in S exist, and each voter has an unlimited budget $b_i = \infty$. In addition, we set $u_i(c) = \alpha_c$ for all $i \in S$ and $c \in T$.

Note that in variations (B) and (C), all candidates in T will be elected (eventually) as $\alpha_c > 0$ for all $c \in T$, and the voters in S have unrestricted budgets when buying candidates in T.

If at the end of the execution of (B) all voters in S have spent strictly less than b/n, then (B) has selected all candidates in T and no voter has overshot their budget. Thus (A) also elects all of T, so $u_i(W) \geqslant u_i(T) \geqslant \alpha$ for all $i \in S$, and we are done. Otherwise, let $i^* \in S$ be the first voter in S who during the execution of (B) spends at least b/n. Suppose this happens just after (B) adds candidate c^* to the outcome and the voters pay for it. Write $W_{(\mathrm{B})}$ for the set of candidates selected by (B) up to but excluding c^*. Note that (A) has also selected all candidates in $W_{(\mathrm{B})}$, because until that point the two rules behave identically.

Next, we will lower bound the utility that i^* receives under (B) by the time i^* has spent at least b/n. To do so, we define a function $f_{(\mathrm{B})}$ so that for a number x, $f_{(\mathrm{B})}(x)$ is the amount of money that i^* had to spend during the execution of (B) until i^* receives utility x. We make this into a continuous, piecewise linear function, so that to get a β-fraction of the utility of a candidate one needs to spend a β-fraction of the total spending for that candidate. See Figure 2. In other words, $f_{(\mathrm{B})}$ consists of a sequence of line segments where

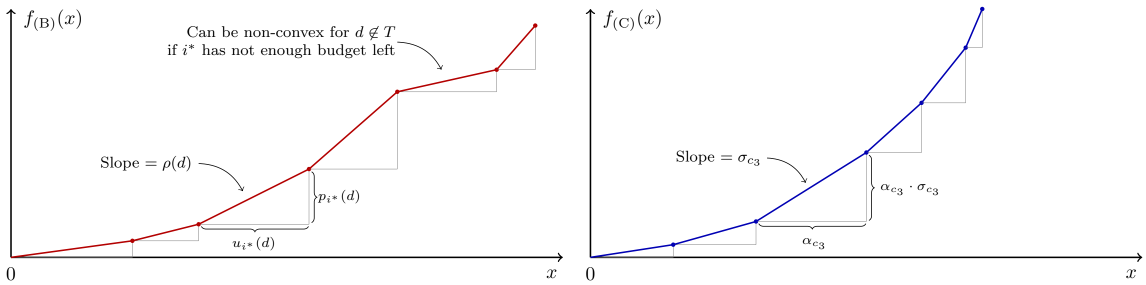

Figure 2: Illustration of functions $f_{(\mathrm{B})}(x)$ and $f_{(\mathrm{C})}(x)$.

the segment corresponding to $d \in C$ has width $u_i(d)$ and height $p_i(d)$ (the amount that i^* paid for d). Note that the segment has a slope of usually $\rho(d)$, but it can be lower than $\rho(d)$ in case $d \notin T$ and i^*'s budget was not enough to pay the full amount $\rho(d) \cdot u_{i^*}(d)$ for d.

We can define a function $f_{(\mathrm{C})}$ in exactly the same way with respect to the execution of (C), based on the same voter i^*. The function $f_{(\mathrm{C})}$ is easy to understand. For each $c \in T$, let us write $\sigma_c = \mathrm{cost}(c)/(|S|\alpha_c)$, and let us label $T = \{c_1, \dots, c_r\}$ such that $\sigma_{c_1} \leqslant \dots \leqslant \sigma_{c_r}$. Note that under (C) each not-yet-selected $c \in T$ is σ_c-affordable, because all voters have unlimited budgets and

$$\sum_{i \in S} \sigma_c \cdot \alpha_c = |S| \cdot \sigma_c \cdot \alpha_c = \mathrm{cost}(c).$$

It follows that $f_{(\mathrm{C})}$ consists of a sequence of line segments of width α_c and slope σ_c, one for each $c \in T$. These line segments come in increasing order of σ_c, i.e. in the order $c_1, \dots, c_r$, because Equal Shares always selects the ρ-affordable candidate with lowest ρ.

Recall that $\alpha = \sum_{c \in T} \alpha(c)$. We claim that

$$f_{(\mathrm{C})}(x) \geqslant f_{(\mathrm{B})}(x) \quad \text{for all } x \in [0, \alpha]. \tag{2}$$

(Intuitively, this says that under (C) the money of i^* is used less efficiently for i^* than under (B).) The inequality certainly holds at $x = 0$ because both functions take the value 0. To establish (2) for other x, we will show that the slope of $f_{(\mathrm{C})}$ is always at least as high as the slope of $f_{(\mathrm{B})}$ (except of course when x is a point joining two line segments, where the slope is not defined, but this only applies to finitely many points).

Let us first note the following useful fact:

$$\text{Under (B), at each step, any not-yet-selected } c \in T \text{ is } \rho\text{-affordable for some } \rho \leqslant \sigma_c. \tag{3}$$

Informally, the fact holds because there are extra voters in (B) compared to (C), and the voters in S have weakly higher utility for c in (B). Formally, looking at the definition (1) of affordability in (B), fact (3) follows because

$$\sum_{i \in S} \sigma_c \cdot u_i(c) + \sum_{i \in N \setminus S} \min\{b_i, \sigma_c \cdot u_i(c)\} \geqslant \sum_{i \in S} \sigma_c \cdot u_i(c) \geqslant \sum_{i \in S} \sigma_c \cdot \alpha_c = \mathrm{cost}(c),$$

where the last step holds because c is σ_c-affordable during the execution of (C).

Now, let $x' \in [0, \alpha]$ be any point that is not a boundary point (for either $f_{(\mathrm{B})}$ or $f_{(\mathrm{C})}$). Say that x' lies in the interior of the line segment corresponding to $d \in C$ of $f_{(\mathrm{B})}$ (call this Segment 1) and in the interior of the line segment corresponding to $c_s \in T$ of $f_{(\mathrm{C})}$ (call this Segment 2). See Figure 3 for an illustration. Consider the time point t when (B) chose to add d to its outcome (but before it actually added d). At time t, i^*'s utility under (B) was equal to the x-coordinate of the left endpoint of Segment 1, and thus less than x'. Further,

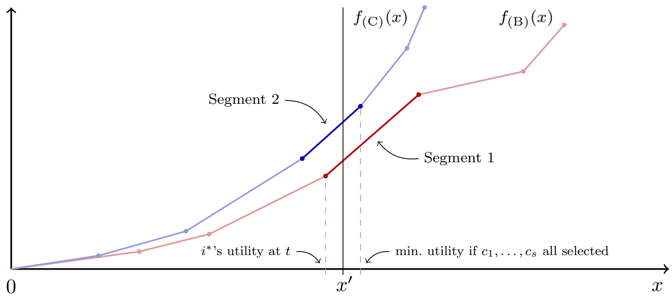

Figure 3: Illustration of the proof of claim (2).

at time t, it cannot be the case that (B) has already selected all of the candidates $c_1, \dots, c_s$, because then i^*'s utility would be at least $\alpha_{c_1} + \dots + \alpha_{c_s}$ which is the x-coordinate of the right endpoint of Segment 2 and thus more than x'. Hence there is some c_p, $p \in [s]$, such that (B) has not selected c_p before time t. By fact (3), c_p is ρ'-affordable in (B) at time t for some $\rho' \leqslant \sigma_{c_p}$. Since (B) always selects a candidate that is ρ-affordable for the smallest ρ, it must be the case that d is ρ-affordable for some $\rho \leqslant \rho'$. Thus, the slope of Segment 1 is at most ρ and hence at most σ_{c_p}. On the other hand, Segment 2 has slope σ_{c_s}. Note that $\sigma_{c_p} \leqslant \sigma_{c_s}$ because $p \leqslant s$. Thus Segment 1 has slope weakly lower than Segment 2. Since this is true for all x' (not on boundary points), our claim (2) follows.

Next, note that

$$f_{(C)}(\alpha) = \sum_{c \in T} \sigma_c \cdot \alpha_c = \sum_{c \in T} \frac{\mathrm{cost}(c)}{|S|} = \frac{\mathrm{cost}(T)}{|S|} \leqslant \frac{b}{n}.$$

Using (2), it follows that

$$f_{(B)}(\alpha) \leqslant \frac{b}{n}. \tag{4}$$

Finally, consider the point in time just after (B) adds c^* to its output and voters have paid for it. By definition of c^*, voter i^* has spent at least b/n at this point. There are two cases:

(i) i^* has spent exactly b/n. In this case c^* is also selected by (A) because the rules behave identically until this point. Now (4) implies that $u_{i^*}(W_{(B)} \cup \{c^*\}) \geqslant \alpha$. Hence $u_{i^*}(W) \geqslant \alpha$.

(ii) i^* has spent strictly more than b/n. In this case, by definition of (B), we have $c^* \in T$. Now (4) implies that $u_{i^*}(W_{(B)} \cup \{c^*\}) > \alpha$. Because $W_{(B)} \subseteq W$, this implies $u_{i^*}(W \cup \{c^*\}) > \alpha$.

In both cases, we conclude that W satisfies EJR up to one project. □

Theorem 1 establishes a strong sense in which the Method of Equal Shares can be said to provide proportional representation. It is the first known rule that can give such strong proportionality guarantees in the model with additive utilities and arbitrary costs.

In the literature on approval-based committee elections, another rule has received much attention: Proportional Approval Voting (PAV), which selects a feasible outcome maximizing $\sum_{i \in N} \mathrm{H}(|A(i) \cap W|)$, where $\mathrm{H}(r) = \sum_{j=1}^{r} 1/j$ is the rth harmonic number.

This rule satisfies EJR when assuming unit costs [3]. But without unit costs, for each $r \geqslant 0$, PAV does not even satisfy EJR up to r projects:

Example 1 (PAV fails EJR)**.** Take $r \geqslant 2$, $b = r^3$, and the following approval-based profile, which is a so-called party list profile where all approval sets are either disjoint or equal

$$\begin{aligned} &r^2 - 1 \text{ voters:} \qquad && \{a_1, a_2, \dots, a_r\}, \\ &1 \text{ voter:} && \{b_1, b_2, \dots, b_r\}. \end{aligned}$$

The projects $a_1, a_2, \dots, a_r$ cost r^2 dollars each; the projects $b_1, b_2, \dots, b_r$ cost 1 dollar each. EJR requires that the 1 voter who approves projects $b_1, \dots, b_r$ must approve at least r projects in the outcome. But PAV selects $\{a_1, a_2, \dots, a_r\}$, leaving the voter with nothing. □

5 Outlook

Since we first introduced the Method of Equal Shares in earlier versions of this paper, other researchers have built on our results. Los et al. [13] considered some other proportionality axioms for PB such as laminar proportionality and PJR. Brill et al. [7] showed that in settings based on approvals (for example, where voters' utility from an approved project is the cost of that project, "cost utilities"), Equal Shares in fact satisfies EJR up to *any* good (EJR-X). For approval settings, they also showed that Equal Shares run based on the induced 0/1-utilities performs well for other utility functions including cost utilities, in a PJR sense. Brill and Peters [6] showed that Equal Shares run with cost utilities satisfied a stronger version of EJR-1 which does not need an explicit cohesiveness requirement.

Fairstein et al. [11] ran an experiment with participants form MTurk to compare different ballot input formats. Running Equal Shares on their data, they report that Equal Shares is robust to changes in the input format. It would be desirable to run further experiments, in particular to see which aggregation methods are preferred by participants.

We have begun actively advocating for cities to try out the Method of Equal Shares. We have built a website that explains the voting method in detail and that includes a discussion of benefits and of implementation details for using the method: equalshares.net. Two small cities will be using Equal Shares in their Participatory Budgets in 2023: Wieliczka in Poland and Aarau in Switzerland. We have recommended that these cities use a utility model based on cost utilities, because this ensures that Equal Shares selects projects roughly in order of their total vote count (instead of vote count divided by cost), which more closely mirrors the standard practice of most cities using the simple greedy method.

Interesting questions for future work remain. Can one characterize Equal Shares as the only one satisfying EJR-1 within a certain class of rules [1]? Can an EJR outcome be computed in strongly polynomial time for additive utilities but unit costs? Can one give more formal reasons to prefer particular completion strategies for Equal Shares? What is the best input format to use together with Equal Shares? In particular, do voters who approve different numbers of projects have predictably different amounts of influence on the outcome, justifying limits on the cardinality of approval sets? What can one say about strategic voting, beyond it being possible in the worst case [14]? Do voters on average have more or less influence on the outcome than under the simple greedy rule, in the sense of being pivotal? Are there simpler rules than Equal Shares that are similarly good? How big a deal is the failure of Equal Shares to be Pareto-optimal and to satisfy the core? What are the incentives facing project proposers in a system that uses Equal Shares (e.g. choice of cost, bundling different projects, splitting into subprojects) [20]? How to best explain the outcome of Equal Shares to voters? Which axioms besides proportionality-type axioms does Equal Shares satisfy? Computing the outcome of Equal Shares can be slow when completing by varying the budget; can it be sped up? How can Equal Shares be adapted to contexts with additional constraints [10]? How can it be adapted to non-additive utilities, and to allow for negative votes? How can we allow for projects with flexible sizes?

References

[1] H. Aziz and B. E. Lee. Proportionally representative participatory budgeting with ordinal preferences. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, pages 5110–5118, 2021. [→ p. 12]

[2] H. Aziz and N. Shah. Participatory budgeting: Models and approaches. In Tamás Rudas and Gábor Péli, editors, *Pathways Between Social Science and Computational Social Science: Theories, Methods, and Interpretations*. Springer, 2020. [→ p. 1]

[3] H. Aziz, M. Brill, V. Conitzer, E. Elkind, R. Freeman, and T. Walsh. Justified representation in approval-based committee voting. *Social Choice and Welfare*, 48(2): 461–485, 2017. [→ p. 2, 7, 11]

[4] H. Aziz, E. Elkind, S. Huang, M. Lackner, L. Sánchez-Fernández, and P. Skowron. On the complexity of extended and proportional justified representation. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, pages 902–909, 2018. [→ p. 7]

[5] H. Aziz, B. Lee, and N. Talmon. Proportionally representative participatory budgeting: Axioms and algorithms. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 23–31, 2018. [→ p. 2, 6, 8]

[6] M. Brill and J. Peters. Robust and verifiable proportionality axioms for multiwinner voting. *arXiv:2302.01989*, 2023. [→ p. 12]

[7] M. Brill, S. Forster, M. Lackner, J. Maly, and J. Peters. Proportionality in approval-based participatory budgeting. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, 2023. [→ p. 2, 12]

[8] Y. Cabannes. Participatory budgeting: A significant contribution to participatory democracy. *Environment and Urbanization*, 16(1):27–46, 2004. [→ p. 1]

[9] M. S. De Vries, J. Nemec, and D. Špaček. *International Trends in Participatory Budgeting: Between Trivial Pursuits and Best Practices*. Palgrave Macmillan, 2021. [→ p. 1]

[10] B. Fain, K. Munagala, and N. Shah. Fair allocation of indivisible public goods. In *Proceedings of the 2018 ACM Conference on Economics and Computation (EC)*, pages 575–592, 2018. Extended version arXiv:1805.03164. [→ p. 2, 12]

[11] R. Fairstein, G. Benadè, and K. Gal. Participatory budgeting design for the real world. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, 2023. [→ p. 12]

[12] M. Lackner and P. Skowron. *Multi-Winner Voting with Approval Preferences*. SpringerBriefs in Intelligent Systems. Springer, 2022. [→ p. 2]

[13] M. Los, Z. Christoff, and D. Grossi. Proportional budget allocations: Towards a systematization. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*, pages 398–404, 2022. [→ p. 12]

[14] D. Peters. Proportionality and strategyproofness in multiwinner elections. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1549–1557, 2018. [→ p. 12]

[15] D. Peters and P. Skowron. Proportionality and the limits of welfarism. In *Proceedings of the 2020 ACM Conference on Economics and Computation (EC)*, pages 793–794, 2020. Extended version arXiv:1911.11747. [→ p. 2, 3, 6, 7]

[16] D. Peters, G. Pierczynski, N. Shah, and P. Skowron. Market-based explanations of collective decisions. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, pages 5656–5663, 2021. [→ p. 6]

[17] D. Peters, G. Pierczyński, and P. Skowron. Proportional participatory budgeting with additive utilities. *Advances in Neural Information Processing Systems*, 34:12726–12737, 2021. [→ p. 6, 8]

[18] S. Rey and J. Maly. The (computational) social choice take on indivisible participatory budgeting. *arXiv:2303.00621*, 2023. [→ p. 1]

[19] P. Skowron. Proportionality degree of multiwinner rules. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*, pages 820–840, 2021. [→ p. 7]

[20] N. Talmon and P. Faliszewski. A framework for approval-based budgeting methods. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, pages 2181–2188, 2019. [→ p. 12]

[21] B. Wampler, S. McNulty, and M. Touchton. *Participatory budgeting in global perspective*. Oxford University Press, 2021. [→ p. 1]

Dominik Peters
CNRS, LAMSADE, Université Paris Dauphine - PSL
Paris, France
Email: mail@dominik-peters.de

Grzegorz Pierczyński and Piotr Skowron
Faculty of Mathematics, Informatics and Mechanics
University of Warsaw
Warsaw, Poland
Email: g.pierczynski@mimuw.edu.pl
Email: p.skowron@mimuw.edu.pl