

Participatory Budgeting with Multiple Resources

Nima Motamed¹, Arie Soeteman², Simon Rey², and Ulle Endriss²

¹ Intelligent Systems, Utrecht University

² Institute for Logic, Language and Computation (ILLC), University of Amsterdam

Abstract. We put forward a formal model of participatory budgeting where projects can incur costs with respect to several different resources, such as money, energy, or emission allowances. We generalise several well-known mechanisms from the usual single-resource setting to this multi-resource setting and analyse their algorithmic efficiency, the extent to which they are immune to strategic manipulation, and the degree of proportional representation they can guarantee. We also prove a general impossibility theorem establishing the incompatibility of proportionality and strategyproofness for this model.

Keywords: Computational social choice · Participatory budgeting

1 Introduction

Participatory budgeting (PB) is an important development in deliberative grassroots democracy now used in hundreds of cities across the globe [27]. PB allows citizens to vote directly on the funding of projects proposed by their peers. Each project is associated with a cost and the projects selected must not exceed a given budget limit. Both in the theoretical literature and in current practice, such costs are expressed in monetary terms only. In this paper, we argue for and define a richer model of PB that can account for costs with respect to resources other than money—such as energy, spatial demands, or allowances for the emission of certain pollutants.

Such a richer model has several advantages. As noted by Goldfrank [10] and Rose and Omolo [24], governmental officials often need to interfere in the PB process to determine the technical feasibility of projects and to ensure their alignment with public policy, thereby reducing transparency. For example, a proposed water fountain may require significant energy resources or a proposed cultural event might breach noise regulations in a residential neighbourhood. A multi-resource model would allow us to make such costs (in terms of energy or noise) explicit and to take them into account when tallying the votes. As we shall see, allowing for multiple resources also permits us to encode additional constraints of practical interest. For instance, to specify that at most \$100*k* (out of a total budget of, say, \$500*k*) may be spent on cultural projects, we could introduce a new resource (“culture-dollars”) with the appropriate budget limit and assign a nonzero cost in terms of this resource to culture-related projects.

Our model is a natural generalisation of the standard single-resource model of PB in which you vote by approving any subset of the projects on the ballot sheet [2]. But now, we require the selected projects not to exceed the budget limit relative to every single one of the resources (rather than just in terms of money). Adopting the methodology of computational social choice [4], we analyse several mechanisms for selecting projects from both an axiomatic and an algorithmic perspective. Regarding the former, we focus on the axioms of proportionality and strategyproofness, and show that no PB mechanism can satisfy both of them, although there are simple mechanisms that perform reasonably well with respect to either one of these desiderata. Regarding algorithmic concerns, we analyse the extent to which the computational complexity of standard mechanisms increases when we move from the single-resource to the multi-resource setting.

Related work. As is well known, PB generalises multiwinner voting [5], a connection we will be using on multiple occasions.

Prior work on PB itself that is of a formal nature has been concerned with the analysis of strategic incentives [9], axioms encoding various fairness requirements [1,21,28], and the computational complexity of PB [6]. Other authors have proposed different extensions of the basic model, e.g., by considering other types of ballots [3,17], allowing for additional constraints [12,13,22], integrating the so-called shortlisting phase—where citizens propose projects—into the basic model [23], modelling several PB exercises running concurrently in districts of the same city [11], and modelling several PB exercises running in consecutive years [16]. Note that Rey et al. [22] also considered multiple resources, although this aspect is not central to their work. We shall discuss specific contributions that are directly relevant to our work in the body of this paper.

We note that the design of PB mechanisms that can account for multi-dimensional constraints, i.e., budget constraints relative to multiple resources has previously been mentioned as an important challenge by Aziz and Shah [2] in their survey on formal approaches to PB.

Paper outline. The remainder of this paper is organised as follows. We introduce our model of multi-resource PB in Section 2, where we also define three mechanisms for this model, formulate suitable axioms of proportionality and strategyproofness, and illustrate the richness of the model by showing how it can accommodate additional constraints on feasible outcomes. We then present our axiomatic results in Section 3 and our algorithmic results in Section 4.

2 The Model

In this section we define our model of multi-resource PB. We also define three simple mechanisms for selecting projects that are directly inspired by familiar mechanisms for single-resource PB, as well as a number of axioms encoding important normative requirements for such mechanisms. Finally, we briefly discuss how the availability of multiple resources allows us to easily encode various additional constraints directly within our model.

2.1 Scenarios and Profiles

A *PB scenario* with m projects and d resources is a tuple $\langle P, \mathbf{c}, \mathbf{b} \rangle$, where $P = \{p_1, \dots, p_m\}$ is a set of *projects*, $\mathbf{c} = (c_1, \dots, c_d)$ is a vector of *cost functions* $c_k : P \rightarrow \mathbb{N} \cup \{0\}$, and $\mathbf{b} = (b_1, \dots, b_d)$ is a vector of *budget limits* $b_k \in \mathbb{N}$. Here, c_k maps each project to its cost in terms of the k -th resource, while b_k is the total number of units of that resource we can spend.³ We extend the definition of each c_k to sets $S \subseteq P$ and write $c_k(S)$ for $\sum_{p \in S} c_k(p)$. Such a set $S \subseteq P$ is *feasible* if $\mathbf{c}(S) \leq \mathbf{b}$, i.e., if $c_k(S) \leq b_k$ for all $k \in \{1, \dots, d\}$, meaning that S does not exceed our budget for any resource. Let $\text{FEAS}(P, \mathbf{c}, \mathbf{b}) = \{S \subseteq P \mid \mathbf{c}(S) \leq \mathbf{b}\}$ be the set of all feasible sets in this scenario.

W.l.o.g., we shall make two assumptions: (i) every project has a nonzero cost in terms of at least one resource (for all $p \in P$ there exists a k with $c_k(p) > 0$); and (ii) there exists at least one feasible set of projects (i.e., $\text{FEAS}(P, \mathbf{c}, \mathbf{b}) \neq \emptyset$).

During a PB exercise, we ask a group $N = \{1, \dots, n\}$ of *voters* to express their preferences by indicating which of the projects in P they approve of. So a *ballot* for a voter i is a set $A_i \subseteq P$. A *profile* is a vector $\mathbf{A} = (A_1, \dots, A_n)$ of such ballots, one for each voter. On the basis of such a profile of approval ballots, we want to select a feasible set of projects to implement.

2.2 Mechanisms

A *mechanism* is a function F that takes as input a scenario $\langle P, \mathbf{c}, \mathbf{b} \rangle$ and a profile \mathbf{A} , and that returns a nonempty set $F(\langle P, \mathbf{c}, \mathbf{b} \rangle, \mathbf{A}) \in \text{FEAS}(P, \mathbf{c}, \mathbf{b})$ of projects that is feasible. The scenario is sometimes omitted when clear from context. We now define three mechanisms for multi-resource PB, all of which are simple generalisations of well-known mechanisms for the single-resource case. Together they cover the main types of approaches to the design of mechanisms considered in the PB literature to date.

Two of our mechanisms are defined in terms of so-called *approval scores*. Given a profile $\mathbf{A} = (A_1, \dots, A_n)$, the approval score of a project p is defined as $s_{\mathbf{A}}(p) = |\{i \in N : p \in A_i\}|$. The approval score of a set $S \subseteq P$ is the sum of the approval scores of the projects in S , i.e. $s_{\mathbf{A}}(S) = \sum_{p \in S} s_{\mathbf{A}}(p) = \sum_{i \in N} |S \cap A_i|$.

Greedy-approval. The *greedy-approval mechanism* F_G goes through all projects in order of their approval scores, with ties being broken by the index of projects in P . Projects are added to the outcome set S one by one, with any project that would render S infeasible being skipped.

For $d = 1$, this is the mechanism most commonly used in practice, though often with certain restrictions on either the size or the cumulative cost of ballots. In case ballots are restricted to feasible sets, the greedy-approval mechanism has been termed *knapsack voting* by Goel et al. [9].

³ Note that *negative costs* can be appropriate as well (e.g., planting trees has “negative environmental cost”). We shall occasionally comment on the effects of doing so.

Max-approval. The *max-approval mechanism* F_M returns a feasible set that maximises the approval score. In case of a tie, we use *lexicographic tie-breaking* based on the projects’ indices to select the final outcome. For $d = 1$, this mechanism and some of its variants have been studied by Talmon and Faliszewski [28].

Sequential load-balancing. The *sequential load-balancing mechanism* F_L is parametrised by a set $R \subseteq \{1, \dots, d\}$ of *relevant resources*. It builds an outcome set S by adding projects one at a time (in a greedy fashion), always picking a project that maintains the feasibility of S and minimises $\max_{k \in R} y_k$, where each y_k is computed by a linear program specific to k and S :

$$\begin{aligned} \min y_k \text{ where } y_k \geq \frac{1}{b_k} \cdot \sum_{p \in S} x_{i,k,p} \text{ for all } i \in N \text{ with} \\ \sum_{i \in N} \mathbb{1}_{p \in A_i} \cdot x_{i,k,p} = c_k(p) \text{ for all } p \in S \text{ (and } x_{i,k,p} \geq 0). \end{aligned}$$

Intuitively, for any voter i with $p \in A_i$, the quantity $x_{i,k,p}$ is the part of $c_k(p)$ shouldered by that voter. Only voters approving of p contribute to its realisation, and the loads across projects are balanced so as to minimise the total load carried by the worst-off voter. Then y_k represents the highest proportion of b_k shouldered by any one voter. Ties between projects again are broken by project index.

F_L is inspired by voting rules for committee elections advocated by Phragmén in the 1890s [14] and closely related to the so-called *maximin support method* recently proposed by Sánchez-Fernández et al. [25]. For PB, a similar mechanism was also proposed by Aziz et al. [1].

2.3 Axioms

In social choice theory, an *axiom* is a formal property of mechanisms that encodes certain normative desiderata. Axioms might relate to the economic efficiency of a mechanism, various notions of fairness, or strategic incentives.

Exhaustiveness. Recall that mechanisms, by definition, return sets that are feasible. But they need not exhaust the budget. This failure to make use of available funds might be considered undesirable. So our first axiom is exhaustiveness. A mechanism F is *exhaustive* if for every scenario $\langle P, \mathbf{c}, \mathbf{b} \rangle$ and profile \mathbf{A} there exists no feasible set $S \in \text{FEAS}(P, \mathbf{c}, \mathbf{b})$ with $S \supsetneq F(\langle P, \mathbf{c}, \mathbf{b} \rangle, \mathbf{A})$.

Proportionality. Intuitively speaking, a mechanism provides proportional representation (or simply: is proportional) if it ensures that sufficiently large groups of voters with sufficiently similar preferences receive adequate representation in the outcome. A range of proportionality axioms has been proposed in the literature, both for PB itself and for the simpler model of approval-based committee elections [1,5]. We define both a strong and weak proportionality axiom. Both are parametrised by a nonempty set R of “relevant” resources (with respect to which we require proportionality).

We call a mechanism F *strongly R -proportional* if, for every scenario $\langle P, \mathbf{c}, \mathbf{b} \rangle$, profile \mathbf{A} , and set $S \subseteq P$, the following two conditions together imply that all of the projects in S get selected, i.e., that $S \subseteq F(\langle P, \mathbf{c}, \mathbf{b} \rangle, \mathbf{A})$:

- (i) $|\{i \in N : A_i = S\}| \geq n \cdot \frac{c_k(S)}{b_k}$ for all $k \in R$;
 (ii) $c_k(S \cup F((P, \mathbf{c}, \mathbf{b}), \mathbf{A})) \leq b_k$ for all $k \notin R$.

Condition (i) says that there is a coalition of voters approving of precisely S that is large enough to “deserve” the proportion of the budget b_k needed to realise S for every relevant resource k .⁴ Condition (ii) expresses that realising S (in situations where it is not yet fully realised) would not exceed the budget for any of the other resources.⁵ A mechanism F is *weakly R -proportional* if it satisfies the above conditions for all singleton sets $S = \{p\}$. We stress that the very narrow conditions for the applicability of the axiom make the axiom logically particularly weak and thus normatively particularly appealing.

In the single-resource case (with $d = 1$ and $R = \{1\}$), weak R -proportionality is the natural generalisation of the basic proportionality axiom formulated by Peters [19] for multiwinner voting (except that Peters also restricts the axiom to so-called “party-list profiles”). This proportionality axiom is particularly attractive due to its simplicity and the weak requirements it imposes. We refer the reader to Peters [19] for a discussion of how it relates to some of the myriad of other proportionality axioms found in the literature. Note that for the single-resource case condition (ii) becomes vacuous.

Strategyproofness. We would like voters to vote truthfully. To make this precise, we need to make assumptions about their incentives. We assume that every voter i has a *preference relation* \succsim_i , which is a reflexive and transitive binary relation on feasible sets of projects (i.e. a preorder). We use \succ_i to refer to its strict part. We further assume that \succsim_i is induced by some set $S_i^* \subseteq P$ of projects voter i *truthfully* approves of. We consider two types of voters; the manner in which \succsim_i is induced by S_i^* depends on the voter’s type:

- For a given nonempty set R of relevant resources, voter i has *R -Paretian preferences* in case $S \succsim_i S'$ holds for two sets $S, S' \in \text{FEAS}(P, \mathbf{c}, \mathbf{b})$ if and only if $c_k(S_i^* \cap S) \geq c_k(S_i^* \cap S')$ for all $k \in R$. That is, such a voter weakly prefers S to S' if the cumulative cost of her truthfully approved projects in S is at least as high as for those in S' with respect to each relevant resource. Thus, $S \succ_i S'$ holds if and only if $c_k(S_i^* \cap S) \geq c_k(S_i^* \cap S')$ for all $k \in R$ and this inequality is strict in at least one case.
- Voter i has *subset preferences* in case $S \succsim_i S'$ holds for $S, S' \in \text{FEAS}(P, \mathbf{c}, \mathbf{b})$ if and only if $S_i^* \cap S \supseteq S_i^* \cap S'$.

⁴ Observe that for condition (i) it is important to count the number of voters who vote for S exactly rather than those who vote for a (not necessarily proper) superset of S . Indeed, weakening the conditions for the applicability of the axiom in this sense would immediately render it impossible to satisfy in general. To see this, consider a single-resource scenario in which we need to divide a budget of $b = 2$ amongst three projects of cost 1, and in which there are two voters, with approval ballots $A_1 = \{p_1, p_2\}$ and $A_2 = \{p_3\}$. Then each project forms a singleton set S for which $n \cdot \frac{c(S)}{b} = 1$, while $|\{i \in N : A_i \supseteq S\}| = 1$. But we cannot select all three projects.

⁵ Note that dropping condition (ii) would render this axiom unsatisfiable in general, since sets satisfying the first condition can exceed the budget for some $k \notin R$.

Let (\mathbf{A}_{-i}, S_i^*) denote the profile \mathbf{A} in which the ballot of voter i has been replaced by S_i^* . We can now define strategyproofness in the familiar manner. A mechanism F is *strategyproof* against voters with either R -Paretian or subset preferences if, for every scenario $\langle P, \mathbf{c}, \mathbf{b} \rangle$ and profile \mathbf{A} , we get $F(\langle P, \mathbf{c}, \mathbf{b} \rangle, \mathbf{A}) \not\prec_i F(\langle P, \mathbf{c}, \mathbf{b} \rangle, (\mathbf{A}_{-i}, S_i^*))$ for all voters $i \in N$ with these preferences.

Following Goel et al. [9], we furthermore define F to be *approximately strategy-proof* against voters with R -Paretian or subset preferences if, for every $\langle P, \mathbf{c}, \mathbf{b} \rangle$ and \mathbf{A} , we get $F(\langle P, \mathbf{c}, \mathbf{b} \rangle, \mathbf{A}) \not\prec_i F(\langle P, \mathbf{c}, \mathbf{b} \rangle, (\mathbf{A}_{-i}, S_i^*) \cup \{p\})$ for all $i \in N$ with these preferences and some $p \in P$. This allows for the possibility that a truthful vote might result in a worse outcome—provided the difference is, in some sense, bounded by the value of the most attractive project.

2.4 Modelling Additional Constraints

Recent work on PB has emphasised the importance of enriching the basic model with the possibility of expressing additional constraints the projects selected for funding must satisfy [2,12,13,22]. As we are going to see now, an advantage of working with a multi-resource PB model is that it allows us to encode such constraints directly within the basic framework.

Distributional constraints. For many real-world PB exercises there are upper bounds on the funding that may be spent on projects belonging to a given category (say, culture or the environment). Suppose $X \subseteq P$ represents a specific category of projects, and that for a certain resource k , we want to limit the part of b_k going to projects in X to $\lfloor \alpha \cdot b_k \rfloor$ for some $\alpha \in [0, 1]$. To achieve this, we can introduce a new resource k^* with $b_{k^*} = \lfloor \alpha \cdot b_k \rfloor$ and $c_{k^*}(p) = \mathbb{1}_{p \in X} \cdot c_k(p)$.

Jain et al. [13] develop a PB model centred around such distributional constraints, and Rey et al. [22] show how to encode them in judgment aggregation. Patel et al. [18] study a different variant of this model, where the distributional constraints relate to the score rather than the costs of the selected projects.

Incompatibility constraints. Some projects might be incompatible with one another. Suppose we want to express that we cannot realise all of the projects in some nonempty set $X \subseteq P$ together. To do so, we can introduce a new resource k^* with budget limit $b_{k^*} = |X| - 1$ and fix $c_{k^*}(p) = \mathbb{1}_{p \in X}$ for each project p . That is, projects in X cost 1 unit and all others do not cost anything. Then respecting the budget constraint for k^* implies never accepting all of the projects in X .

In the single-resource setting incompatibility constraints are a special case of distributional constraints. But in general this is not the case, since two incompatible projects might not both have a nonzero cost for the same resource.

Dependency constraints. Realising a given project might be possible only if certain other projects are implemented as well. If we were to allow for negative costs, we could easily encode such dependency constraints. Suppose we want to express that project p^* should be selected only if all projects in $X \subseteq P$ are

selected as well. We again create a new resource k^* , and set $b_{k^*} = 1$, $c_{k^*}(p^*) = |X| + 1$, $c_{k^*}(p) = -1$ for all $p \in X$, and $c_{k^*}(p) = 0$ for all other projects. Then selecting p^* and thus spending an amount of $|X| + 1$ is possible only if we also select all of the projects in X and thus push the total amount spent down to the budget limit of 1. Rey et al. [22] also discuss modelling such constraints.

The fact that encoding constraints involves introducing some purely technical resources lends additional support to the idea of parametrising mechanisms and axioms by a set of relevant resources R . For example, for F_L we may want to put all “real” resources in R but leave all “technical” resources aside.

3 Axiomatic Analysis

In this section we first analyse the concrete mechanisms defined earlier in view of the axiomatic requirements of proportionality and strategyproofness, and we then show that it is impossible to satisfy both requirements at the same time.

3.1 Proportionality

Unfortunately, neither the greedy-approval mechanism nor the max-approval mechanism can guarantee weak proportionality, and thus certainly not strong proportionality. To see this, consider the following example.

Example 1. Take a single-resource scenario $\langle P, c, b \rangle$ with $P = \{p_1, p_2\}$, $b = 3$, $c(p_1) = 1$, and $c(p_2) = 3$. For profile $\mathbf{A} = (\{p_1\}, \{p_2\}, \{p_2\})$ both F_G and F_M return the outcome $\{p_2\}$. However, weak proportionality (with $R = \{1\}$) would require p_1 to be part of that outcome. \triangle

This kind of counterexample also works for multi-resource scenarios: simply add any number of dummy resources with budget 1 and cost 0 for both projects (as long as $R \supseteq \{1\}$). On the other hand, the sequential load-balancing mechanism F_L satisfies even our strong proportionality axiom.

Proposition 1. *The sequential load-balancing mechanism F_L is strongly R -proportional for any set R of relevant resources.*

Proof. The proof is similar to that of Proposition 3.13 in the work of Aziz et al [1]. Suppose F_L is *not* strongly R -proportional, for some R . Then there must be a scenario $\langle P, \mathbf{c}, \mathbf{b} \rangle$, a profile \mathbf{A} , and a subset of projects $S^* \subseteq P$ satisfying the requirements of strong R -proportionality but for which there exists a project $p^* \in S^*$ not selected by F_L . Let $N^* \subseteq N$ be the set of voters such that for all $i \in N^*$, we have $A_i = S^*$. Recall that F_L works in iterations. Let ℓ be the first iteration for which selecting p^* would violate the budget constraint of at least one relevant resource which we call $k^* \in R$. Thanks to condition (ii) of strong R -proportionality, we know that $\{p^*\} \cup F_L(\langle P, \mathbf{c}, \mathbf{b} \rangle, \mathbf{A})$ cannot exceed the budget of a non-relevant resource. This implies that such an ℓ always exists as otherwise

F_L would not have terminated. In the following we will prove a contradiction, namely that F_L should have selected p^* at an iteration before ℓ .

Let S be the set of projects selected by F_L at iteration ℓ . Use $x_{i,k,p}$ to represent the part of $c_k(p)$ shouldered by voter i at iteration ℓ (see the definition of F_L). Define $x_{i,k} = \sum_{p \in A_i} \mathbb{1}_{p \in A_i} \times x_{i,k,p}$ to be the total load of voter i for resource k at iteration ℓ . Given that all voters in N^* approve only of S^* , we see that the cost $c_k(S) - c_k(S \cap S^*)$ for every relevant resource $k \in R$ should be spread across $n - |N^*|$ voters. By averaging, there must then be a voter $i \in N \setminus N^*$ for which:

$$\forall k \in R : x_{i,k} \geq \frac{c_k(S) - c_k(S \cap S^*)}{n - |N^*|} \times \frac{1}{b_k}.$$

From the definition of ℓ , we know that for k^* we have $c_{k^*}(S \cup \{p^*\}) > b_{k^*}$. The equation above thus implies for resource k^* that:

$$x_{i,k^*} > \frac{b_{k^*} - c_{k^*}(p^*) - c_k(S \cap S^*)}{n - |N^*|} \times \frac{1}{b_{k^*}}.$$

Now from the definition of strong R -proportionality, we know that $\frac{|N^*|}{n} \times b_{k^*} \geq c_{k^*}(p^*) + c_{k^*}(S \cap S^*)$. This implies that:

$$x_{i,k^*} > \frac{b_{k^*} - \frac{|N^*|}{n} \times b_{k^*}}{n - |N^*|} \times \frac{1}{b_{k^*}} = \frac{1}{n}. \quad (1)$$

Thus at iteration ℓ , the maximum load of a voter is at least $1/n$.

We now assume that at iteration $\ell - 1$, it is project p^* that is selected, instead of the other project that F_L selected. We distinguish between two cases.

First, if all voters have load no more than $1/n$ for all resources $k \in R$ at the “new” iteration ℓ , then we are done: to minimise the maximum load, F_L should have selected p^* at iteration $\ell - 1$ because of Equation (1).

Suppose now that there is a voter $i \in N$ and a resource $k \in R$ such that i 's load for k at the “new” iteration ℓ exceeds $1/n$. Note first that we must have $i \notin N^*$. Indeed, the costs of projects in S^* can always be distributed across voters in N^* , keeping their load for every $k \in R$ at most $\frac{|S^*| \times b_{k^*}}{n \times |S^*|} \times \frac{1}{b_{k^*}} = 1/n$. Then, since $i \notin N^*$, i 's load did not increase by selecting p^* , and so there must then be a smallest iteration $\ell' < \ell$ after which i 's load for k exceeded $1/n$. But then, since the maximum load after ℓ' exceeded $1/n$, we find that F_L should have selected p^* at iteration $\ell' - 1$. Indeed, by selecting p^* , all voters in $N \setminus N^*$ would have a load of less than $1/n$, while the load of voters in N^* still would not exceed $1/n$.

Overall, by definition of F_L , project p^* should have been selected before iteration ℓ . By contradiction, F_L is thus proven to be strongly R -proportional. \square

We note that this positive result ceases to hold when we allow for negative costs. Indeed, as the following example demonstrates, in that case satisfying strong proportionality is impossible for any mechanism.

Example 2. Consider a single-resource scenario $\langle P, c, b \rangle$ with $P = \{p_1, p_2, p_3\}$, $b = 2$, $c(p_1) = c(p_2) = 2$, and $c(p_3) = -1$. Then under profile $\mathbf{A} = (A_1, A_2)$ with

$A_1 = \{p_1, p_3\}$ and $A_2 = \{p_2, p_3\}$, both voters approve sets with cumulative cost $c(A_1) = c(A_2) = 2 - 1 = 1$, and so strong proportionality requires us to accept both sets. But this would exceed the budget: $c(A_1 \cup A_2) = 2 + 2 - 1 = 3 > b$. Δ

3.2 Strategyproofness

None of our mechanisms are strategyproof against voters with either Paretian or subset preferences (not even for $d = 1$). We again provide an example.

Example 3. Take a scenario with budget $b = 2$ and three projects with $c(p_1) = 1$, $c(p_2) = 2$, and $c(p_3) = 1$. Suppose we receive two ballots, $S_1^* = \{p_3\}$ and $A_2 = \{p_2\}$. Both F_G and F_M pick p_2 (due to lexicographic tie-breaking). If voter 1 instead (untruthfully) votes $A_1 = \{p_1, p_3\}$, both mechanisms return $\{p_1, p_3\}$. The same applies to F_L if we add a third voter with $A_3 = \{p_2\}$. But $\{p_1, p_3\} \succ_1 \{p_2\}$, for both Paretian and subset preferences. Δ

So let us focus on approximate strategyproofness instead. As we shall see, F_G guarantees approximate strategyproofness for voters with subset preferences, but not (in general) for voters with Paretian preferences. As we shall also see, unfortunately, neither F_M nor F_L can guarantee approximate strategyproofness in either case. For the positive result we first prove a simple lemma.

Lemma 1. *If projects cost 1 unit of one resource and 0 units of all others, then under F_G , voters with R -Paretian preferences weakly prefer the outcome obtained by voting truthfully over any obtained by voting untruthfully (for any R).*

Proof. In this setting, F_G picks, for each $k \in \{1, \dots, d\}$, the b_k most approved projects costing 1 unit of resource k . Now, since every project in S_i^* is approved at least as often in (\mathbf{A}_{-i}, S_i^*) as in \mathbf{A} , we have $F(\mathbf{A}_{-i}, S_i^*) \succeq_i F(\mathbf{A})$. \square

Proposition 2. *F_G is approximately strategyproof against voters with R -Paretian preferences for $R = \{1, \dots, d\}$ and against voters with subset preferences.*

Proof. Let $R = \{1, \dots, d\}$ and consider a scenario $\langle P, \mathbf{c}, \mathbf{b} \rangle$. Construct a second scenario $\langle P', \mathbf{c}', \mathbf{b} \rangle$ where all projects in P have been decomposed into smaller projects that cost 1 unit of one resource and 0 units of all other resources, such that the costs of any $p \in P$ equals the sum of costs of the corresponding projects $p' \in P'$. By Lemma 1, no voter with R -Paretian preferences has an incentive to manipulate in the second scenario. F_G accepts, in the first scenario, projects $p \in P$ in the same order as it accepts the corresponding projects $p' \in P'$ in the second scenario, until the first of the d budget limits is reached. For that resource k for which the limit is reached first and that project p that is not accepted, the difference in cost between the two outcomes is at most $c_k(p)$. If we give project p to an R -Paretian voter on top of the outcome of the first scenario, then the amount of resource k spent as she desires is at least as much as in the second scenario, which in turn is at least as much as in any manipulated version of the second scenario, which is at least as much as in any manipulated version

of the first scenario. Hence, she will not have any incentive to manipulate in the first scenario either.

Finally, note that strict subset preferences imply strict $\{1, \dots, d\}$ -Paretian preferences. Thus, approximate strategyproofness under $\{1, \dots, d\}$ -Paretian preferences implies the same under subset preferences, completing the proof. \square

In the single-resource setting, F_G has been shown to guarantee approximate strategyproofness against (what we call) Paretian voters using a similar approach [9]. However, as our next example illustrates, for multi-resource PB in which voters might not care about all resources, this is no longer the case.

Example 4. Take a scenario with $P = \{p_1, p_2, p_3, p_4, p_5\}$, three resources, $R = \{1\}$, and the following costs and budget limits:

Cost	p_1	p_2	p_3	p_4	p_5	Budget limit
c_1	0	0	0	1	1	$b_1 = 2$
c_2	2	1	1	0	0	$b_2 = 2$
c_3	1	3	3	4	4	$b_3 = 9$

Let us consider the two-voter profile $(S_1^*, A_2) = (\{p_2, p_3, p_4, p_5\}, \{p_1, p_2, p_3\})$. Then the greedy-approval mechanism F_G selects the set $\{p_2, p_3\}$. If voter 1 instead votes $A_1 = \{p_1, p_4, p_5\}$, F_G returns $\{p_1, p_4, p_5\}$, which is better for her—in terms of resource 1—than $\{p_2, p_3\} \cup \{p\}$, for every $p \in P$. \triangle

The next two examples demonstrate that neither F_M nor F_L can guarantee approximate strategyproofness against voters with subset preferences (and, thus, certainly not against Paretian voters).

Example 5. Consider the single-resource scenario $\langle P, c, b \rangle$ with projects $P = \{p_1, p_2, p_3, p_4, p_5\}$. Let the first four projects cost 1, while $c(p_5) = 4 = b$. Now let $(S_1^*, A_2, A_3, A_4) = (\{p_1, p_2\}, \{p_5\}, \{p_5\}, \{p_5\})$. Then F_M returns $\{p_5\}$. However, if voter 1 switches to $A_1 = \{p_1, p_2, p_3, p_4\}$, then F_M returns $\{p_1, p_2, p_3, p_4\}$, increasing the set of accepted projects she truly likes from \emptyset to $\{p_1, p_2\}$. \triangle

Example 6. Consider a two-resource scenario with $P = \{p_1, p_2, p_3, p_4, p_5\}$ with these costs and budget limits:

Costs	p_1	p_2	p_3	p_4	p_5	Budget limit
c_1	2	2	0	3	0	$b_1 = 5$
c_2	2	0	2	0	3	$b_2 = 5$

For the profile $(S_1^*, A_2, A_3) = (\{p_1, p_2, p_3\}, \{p_1, p_4, p_5\}, \{p_1, p_4, p_5\})$, the sequential load-balancing mechanism F_L picks the set $\{p_1, p_4, p_5\}$. However, if voter 1 switches to $A_1 = \{p_2, p_3\}$, then F_L still selects p_1 , but also p_2 and p_3 , since voter 1 can no longer carry any load for p_1 . Thus, by manipulating she can add two projects she cares about to the outcome, without losing any others. \triangle

3.3 An Impossibility Result

We now show it is impossible to guarantee both weak proportionality and strategyproofness together. Our result mirrors (and is inspired by) an impossibility result for multiwinner voting due to Peters [19], although there are subtle differences (meaning that our result is not implied by that of Peters). In particular, Peters requires a (very weak) efficiency axiom (for a discussion of this point, refer to Peters [20]).⁶ We are going to prove the following result (for single-resource PB) and then generalise to full multi-resource PB.

Theorem 1. *Let $b \geq 3$, $m > b$ and $n = q \cdot b$ for some integer $q \geq 1$. Then no mechanism can guarantee both weak proportionality and strategyproofness against voters with Paretian preferences for PB scenarios with a single resource, budget b , m projects, and n voters.*

For ease of reading, let us call a single-resource mechanism F *good* if it satisfies both weak proportionality and strategyproofness against voters with Paretian preferences. We first prove Theorem 1 for the special case of $(b, m, n) = (3, 4, 3)$, and then generalise using induction.

Lemma 2. *No mechanism for $(b, m, n) = (3, 4, 3)$ is good.*

Proof. For the sake of contradiction, suppose F is such a mechanism. Let $\langle P, c, b \rangle$ be a single-resource scenario with $P = \{a, b, c, d\}$, $c(p) = 1$ for all $p \in P$, and $b = 3$. Consider profile $\mathbf{A}^1 = (ab, c, d)$, where we omitted set brackets to improve readability. By weak proportionality, we must have $cd \subseteq F(\mathbf{A}^1)$. Furthermore, by strategyproofness, either a or b must be in the selected project set as well, since otherwise voter 1 can manipulate by removing a single project from her ballot. Thus $F(\mathbf{A}^1)$ is either acd or bcd . W.l.o.g., let us assume the former is the case. Table 1 shows how to derive a contradiction from $F(\mathbf{A}^1) = acd$ by means of a sequence of steps involving 14 different profiles.⁷ \square

Next, we prove three inductive lemmas.

Lemma 3. *If there exists a good mechanism for (b, m, n) with $n = q \cdot b$ for some integer $q \geq 1$, then a good mechanism also exists for (b, m, b) .*

⁶ We are able to circumvent the need for this additional efficiency requirement because we do not impose exhaustiveness (which in multiwinner voting is an implicit part of the basic model). This gives us more freedom for the inductive lemmas we need to prove. At the same time, our result is weaker than that of Peters in other respects: his proportionality axiom is subtly weaker (as it needs to be imposed only for so-called party-list profiles) and his result applies even under subset preferences.

⁷ We found these 14 profiles and the derivation of Table 1 by first encoding the requirements of F as a set of clauses in propositional logic, and then applying a SAT-solver to that set to compute a minimally unsatisfiable set exhibiting the impossibility of finding a mechanism of the required kind. For an introduction to this approach, the reader may wish to consult the expository article of Geist and Peters [8].

Profile	Strategyproofness	Proportionality	Outcome
$\mathbf{A}^2 = (ab, ac, d)$	$ac \subsetneq F(\mathbf{A}^2) \Rightarrow M(2, \mathbf{A}^2, \mathbf{A}^1)$	$d \subseteq F(\mathbf{A}^2)$	$F(\mathbf{A}^2) = acd$
$\mathbf{A}^3 = (b, ac, d)$	$ab \subseteq F(\mathbf{A}^3) \Rightarrow M(1, \mathbf{A}^2, \mathbf{A}^3)$ $ac \not\subseteq F(\mathbf{A}^3) \Rightarrow M(2, \mathbf{A}^3, (b, a, d))$	$bd \subseteq F(\mathbf{A}^3)$	$F(\mathbf{A}^3) = bcd$
$\mathbf{A}^4 = (b, ac, cd)$	$cd \subsetneq F(\mathbf{A}^4) \Rightarrow M(3, \mathbf{A}^4, \mathbf{A}^3)$	$b \subseteq F(\mathbf{A}^4)$	$F(\mathbf{A}^4) = bcd$
$\mathbf{A}^5 = (b, a, cd)$	$ac \subseteq F(\mathbf{A}^5) \Rightarrow M(2, \mathbf{A}^4, \mathbf{A}^5)$ $cd \not\subseteq F(\mathbf{A}^5) \Rightarrow M(3, \mathbf{A}^5, (b, a, c))$	$ab \subseteq F(\mathbf{A}^5)$	$F(\mathbf{A}^5) = abd$
$\mathbf{A}^6 = (b, ad, cd)$	$ad \subsetneq F(\mathbf{A}^6) \Rightarrow M(2, \mathbf{A}^6, \mathbf{A}^5)$	$b \subseteq F(\mathbf{A}^6)$	$F(\mathbf{A}^6) = abd$
$\mathbf{A}^7 = (b, ad, c)$	$cd \subseteq F(\mathbf{A}^7) \Rightarrow M(3, \mathbf{A}^6, \mathbf{A}^7)$ $ad \not\subseteq F(\mathbf{A}^7) \Rightarrow M(2, \mathbf{A}^7, (b, a, c))$	$bc \subseteq F(\mathbf{A}^7)$	$F(\mathbf{A}^7) = abc$
$\mathbf{A}^8 = (b, ad, ac)$	$ac \subsetneq F(\mathbf{A}^8) \Rightarrow M(3, \mathbf{A}^8, \mathbf{A}^7)$	$b \subseteq F(\mathbf{A}^8)$	$F(\mathbf{A}^8) = abc$
$\mathbf{A}^9 = (b, d, ac)$	$ad \subseteq F(\mathbf{A}^9) \Rightarrow M(2, \mathbf{A}^8, \mathbf{A}^9)$ $ac \not\subseteq F(\mathbf{A}^9) \Rightarrow M(3, \mathbf{A}^9, (b, d, c))$	$bd \subseteq F(\mathbf{A}^9)$	$F(\mathbf{A}^9) = bcd$
$\mathbf{A}^{10} = (b, cd, ac)$	$cd \subsetneq F(\mathbf{A}^{10}) \Rightarrow M(2, \mathbf{A}^{10}, \mathbf{A}^9)$	$b \subseteq F(\mathbf{A}^{10})$	$F(\mathbf{A}^{10}) = bcd$
$\mathbf{A}^{11} = (b, cd, a)$	$ac \subseteq F(\mathbf{A}^{11}) \Rightarrow M(3, \mathbf{A}^{10}, \mathbf{A}^{11})$ $cd \not\subseteq F(\mathbf{A}^{11}) \Rightarrow M(2, \mathbf{A}^{11}, (b, c, a))$	$ab \subseteq F(\mathbf{A}^{11})$	$F(\mathbf{A}^{11}) = abd$
$\mathbf{A}^{12} = (b, cd, ad)$	$ad \subsetneq F(\mathbf{A}^{12}) \Rightarrow M(3, \mathbf{A}^{12}, \mathbf{A}^{11})$	$b \subseteq F(\mathbf{A}^{12})$	$F(\mathbf{A}^{12}) = abd$
$\mathbf{A}^{13} = (b, c, ad)$	$cd \subseteq F(\mathbf{A}^{13}) \Rightarrow M(2, \mathbf{A}^{12}, \mathbf{A}^{13})$ $ad \not\subseteq F(\mathbf{A}^{13}) \Rightarrow M(3, \mathbf{A}^{13}, (b, c, a))$	$bc \subseteq F(\mathbf{A}^{13})$	$F(\mathbf{A}^{13}) = abc$
$\mathbf{A}^{14} = (ab, c, ad)$	$ab \subsetneq F(\mathbf{A}^{14}) \Rightarrow M(1, \mathbf{A}^{14}, \mathbf{A}^{13})$ $abc = F(\mathbf{A}^{14}) \Rightarrow M(3, \mathbf{A}^{14}, \mathbf{A}^1)$	$c \subseteq F(\mathbf{A}^{14})$	Contradiction

Table 1. Derivation for Lemma 2. $M(i, \mathbf{A}, \mathbf{A}')$ means that voter i can successfully manipulate by moving from profile \mathbf{A} to profile \mathbf{A}' , while $S \not\subseteq S'$ signifies that $S \cap S' = \emptyset$.

Proof. Let F be a good mechanism for (b, m, n) . We construct F' for (b, m, b) as follows. Given a profile \mathbf{A} with b voters, copy each ballot q times to construct profile \mathbf{A}^q , and let $F'(\mathbf{A}) = F(\mathbf{A}^q)$. We show that F' satisfies both axioms, starting with proportionality. Note that, due to $d = 1$, the second proportionality condition is vacuously satisfied. Suppose that for some project p with cost $c(p)$, we have $|\{i \in N : A_i = \{p\}\}| \geq b \cdot \frac{c(p)}{b}$. Then q times as many (i.e., at least $n \cdot \frac{c(p)}{b}$) voters have ballot $\{p\}$ in \mathbf{A}^q . Since $p \in F(\mathbf{A}^q)$, also $p \in F'(\mathbf{A})$. For strategyproofness, suppose for the sake of contradiction that $F'(\mathbf{A}) \succ_i F'(\mathbf{A}_{-i}, S_i^*)$ for some voter i with Paretian preferences. Then $F(\mathbf{A}^q) \succ_i F((\mathbf{A}_{-i}, S_i^*)^q)$. Now, in $(\mathbf{A}_{-i}, S_i^*)^q$, let the q voters corresponding to i switch, one by one, to the untruthful ballot A_i . This results in a sequence of q profiles, each of which is not strictly preferred over the former by i , since F is strategyproof. As for $d = 1$ the relation \succ_i is transitive, we get $F(\mathbf{A}^q) \not\succeq_i F((\mathbf{A}_{-i}, S_i^*)^q)$, a contradiction. \square

Lemma 4. *If there exists a good mechanism for $(b, m + 1, n)$, then a good mechanism also exists for (b, m, n) .*

Proof. Let F be a good mechanism for $(b, m + 1, n)$. We construct F' for (b, m, n) . Add a dummy project p^* so that $F'(\mathbf{A}, P) = F(\mathbf{A}, P \cup \{p^*\}) \setminus \{p^*\}$ for every profile \mathbf{A} .⁸ To show that F' satisfies proportionality, note that any project in P satisfying the two conditions is selected by F if and only if it is selected by F' . The strategyproofness of F' follows directly from that of F since no voter approves of p^* in \mathbf{A} . \square

Lemma 5. *If there exists a good mechanism for $(b, m, n) = (k + 1, k + 2, k + 1)$, then a good mechanism also exists for $(k, k + 1, k)$.*

Proof. Let F be a good mechanism for $(k + 1, k + 2, k + 1)$. We construct F' for $(k, k + 1, k)$. Given profile \mathbf{A}^k with k voters and $k + 1$ projects in P^{k+1} , add a dummy project p^* with cost $c(p^*) = 1$ to form P^{k+2} and a singleton ballot $\{p^*\}$ to form \mathbf{A}^{k+1} . Now let $F'(\mathbf{A}^k, P^{k+1}) = F(\mathbf{A}^{k+1}, P^{k+2}) \setminus \{p^*\}$. Note that, since F is proportional and $|\{i \in N : A_i = \{p^*\}\}| = 1 \geq \frac{k+1}{k+1} \cdot c(p^*)$, we always have $p^* \in F(\mathbf{A}^{k+1}, P^{k+2})$, and so F' does not violate the budget constraint (i.e., F' is well-defined). For the proportionality of F' , note that if a project $p \in P^{k+1}$ is approved of $\frac{k}{k+1} \cdot c(p)$ times in \mathbf{A}^k , it is also approved of $\frac{k+1}{k+1} \cdot c(p)$ times in \mathbf{A}^{k+1} . Since p is then selected by F , it is also selected by F' . For strategyproofness, again note that a strict preference between two outcomes of F' for a voter $i \in \{1, \dots, k\}$ implies the same strict preference for the associated outcomes of F , since i does not approve of $\{p^*\}$. Hence, the strategyproofness of F' follows from the strategyproofness of F . \square

We are now ready to prove our theorem.

Proof (of Theorem 1). For the sake of contradiction, suppose there exists a good mechanism for some (b, m, n) with $b \geq 3$, $m > b$, and $n = q \cdot b$. Then, by Lemma 3, there exists such a mechanism for (b, m, b) . Further, by repeated applications of Lemma 4 and Lemma 5, we can get a good mechanism for $(b, b + 1, b)$ and then for $(3, 4, 3)$. But this contradicts Lemma 2. \square

Using a straightforward induction over the number of resources, we can generalise to the multi-resource setting and obtain the following corollary.

Corollary 1. *Let $d \geq 1$, $R \subseteq \{1, \dots, d\}$, $m > b_k \geq 3$ for some $k \in R$, and $n = q \cdot b_k$ for some $q \geq 1$. Then no mechanism can guarantee both weak R -proportionality and strategyproofness against voters with R -Paretian preferences for d -resource PB scenarios with relevant resources R , budgets $\mathbf{b} = (b_1, \dots, b_d)$, m projects, and n voters.*

To what extent this impossibility result can be strengthened further as well as whether relaxing some of our assumptions might allow for the design of attractive

⁸ Observe that F' might not be exhaustive, with the implications discussed above.

mechanisms are interesting open problems. For example, we do not know whether the impossibility persists for voters with subset preferences (the counterexample used for the proof of the base case still works, but some of the arguments used in the inductive lemmas do not). Similarly, we do not have a full picture regarding the impact of the constraints on the numerical parameters involved (such as n being a multiple of one of the budget limits) on the impossibility.⁹ Finally, we do not know whether there are mechanisms for multi-resource PB that are weakly proportional and approximately strategyproof.¹⁰

4 Algorithms and Complexity

We now analyse each of the three mechanisms defined in Section 2.2 from a computational point of view. We also comment on how allowing for negative costs would affect our results.

4.1 The Greedy-Approval Mechanism

The greedy-approval mechanism F_G clearly can be executed in polynomial time. This remains true when we allow for negative costs. However, as illustrated by the following example, it is questionable whether a greedy mechanism is appropriate in the presence of negative costs.

Example 7. Consider a PB scenario with one resource and three projects, where $b_1 = 5$, $c_1(p_1) = c_1(p_2) = 3$, and $c_1(p_3) = -1$. Suppose p_1 has a higher approval score than p_2 , which in turn has a higher score than p_3 . Then F_G would first accept p_1 (reducing the budget to $5 - 3 = 2$), then skip p_2 (as it costs more than 2), and finally accept p_3 . At this point, the remaining budget is $2 + 1 = 3$, so accepting p_2 would now be feasible. But that would amount to a form of backtracking (given that we now accept a project we previously rejected), which is not allowed under greedy algorithms in general and F_G in particular. \triangle

4.2 The Max-Approval Mechanism

For single-resource PB, Talmon and Faliszewski [28] sketch a polynomial-time algorithm implementing the max-approval mechanism F_M . As we shall see next, for multi-resource PB there can be no such algorithm, unless $P = NP$.

First, let us formally define a decision variant of the problem of maximising the approval score (for a fixed dimension d).¹¹

⁹ The question of whether these constraints can be relaxed is of some technical interest, but arguably less relevant to practice. Indeed, we would want our mechanism to work for *arbitrary* numbers of voters (including those that are multiples of a budget limit).

¹⁰ When all resources are relevant (in the single-resource case for instance), there is a trivial mechanism of this kind: simply return the union of all singletons satisfying condition (i) in the definition of proportionality. To see this, recall that condition (ii) is vacuous if there are only relevant resources.

¹¹ Recall that the approval score of a set S for a given profile \mathbf{A} is defined as $s_{\mathbf{A}}(S) = \sum_{i \in N} |S \cap A_i|$, and that F_M seeks to maximise that score.

 MAXAPPSCORE_{*d*}

Instance: *d*-resource scenario $\langle P, \mathbf{c}, \mathbf{b} \rangle$, profile \mathbf{A} , target $K \in \mathbb{N}$
Question: Is there a set $S \in \text{FEAS}(P, \mathbf{c}, \mathbf{b})$ with $s_{\mathbf{A}}(S) \geq K$?

This problem is closely related to the *d*-dimensional knapsack problem [15]. In particular, in the setting where $d = 2$ and there is just a single voter who approves of all projects, our problem is equivalent to the problem referred to as CARDINALITY (2-KP) by Kellerer et al. [15], which is a weakly NP-hard problem. This insight immediately implies the next result.

Proposition 3. *For any number $d \geq 2$ of resources, there exists no polynomial-time algorithm to compute outcomes under the max-approval mechanism F_M , unless $P = NP$.*

But note that weak NP-hardness still allows for the existence of pseudopolynomial-time algorithms. Indeed, the dynamic programming algorithm of Kellerer et al. [15] for the multidimensional knapsack problem can be applied directly (after translating the input profile into a vector of approvals per project). This yields the following observation.

Proposition 4. *For any fixed number d of resources, outcomes under the max-approval mechanism F_M can be computed in pseudopolynomial time.*

Mapping to a *d*-dimensional knapsack problem works only when d is a constant. This assumption is often reasonable: we typically have to deal with just a small number of resources (money, space, pollutants). However, we saw in Section 2.4 that encoding distributional or incompatibility constraints results in additional “technical” resources, the number of which grows with the number of projects. So it is important to also understand the complexity of F_M relative to d . To this end, we now introduce a variant of the decision problem defined earlier. Instances of this new problem are PB scenarios for arbitrary numbers of resources (rather than for some fixed dimension d).

 MAXAPPSCORE

Instance: scenario $\langle P, \mathbf{c}, \mathbf{b} \rangle$, profile \mathbf{A} , target $K \in \mathbb{N}$
Question: Is there a set $S \in \text{FEAS}(P, \mathbf{c}, \mathbf{b})$ with $s_{\mathbf{A}}(S) \geq K$?

To analyse the complexity of this problem, we employ a similar construction as the one we used to encode incompatibility constraints in the basic model (see Section 2.4). Observe that the following result rules out the possibility of the existence of a pseudopolynomial-time algorithm.

Proposition 5. *MAXAPPSCORE is strongly NP-hard.*

Proof. We proceed by reduction from the INDEPENDENT SET problem, asking whether a given graph $G = \langle V, E \rangle$ has an independent set of size K . This problem is known to be strongly NP-complete [7].

Given $G = \langle V, E \rangle$ and K , construct a d -resource PB scenario $\langle V, \mathbf{c}, \mathbf{b} \rangle$ with $d = |E|$: $c_k(p) = 1$ if the k th edge contains vertex p (and $c_k(p) = 0$ otherwise) and $\mathbf{b} = (1, \dots, 1)$. So a project set S is feasible if and only if S is an independent set in the original graph. Now consider a profile in which a single voter approves of all projects. Then an approval score of K is attainable if and only if the graph has an independent set of size K . \square

Jain et al. [13] use the same kind of reduction to prove hardness for their model of PB with distributional constraints, which they call “project groups”. As their model is a special case of ours, this thus entails Proposition 5. We nevertheless included our proof above as it is much easier to follow.

Finally, let us note that, while Propositions 3 and 5 clearly continue to hold when we allow for negative costs, the dynamic programming algorithm of Kellerer et al. [15] does not generalise to this setting.

4.3 The Sequential Load-Balancing Mechanism

Even though the definition of the sequential load-balancing mechanism F_L is rather involved, it is not difficult to show that it is a tractable mechanism.

Proposition 6. *Outcomes under the sequential load-balancing mechanism F_L can be computed in polynomial time.*

Proof. The claim follows immediately from the definition of the mechanism, given that executing F_L boils down to solving a polynomial number of linear programs, each of which is solvable in polynomial time [26]. \square

Proposition 6 remains valid when we permit negative costs. But given the “greedy” nature of F_L , it is debatable whether it should be considered appropriate to use F_L in the presence of negative costs (just as it is debatable for F_G). Indeed, conceptually, a core feature of F_L , which arguably makes it a natural mechanism, is the fact that the load of each individual voter never decreases as we accept additional projects. This property is lost once we allow for negative costs.

5 Conclusion

We initiated the systematic study of PB with multiple resources. Our results indicate that—despite the significant increase in expressive power when moving from the single-resource to the multi-resource setting—devising attractive mechanisms does not become insurmountably harder, in either axiomatic or algorithmic terms. We hope that this will encourage others to further develop this approach and to, eventually, field it in real-world PB exercises.

Acknowledgements

This paper grew out of a student project that started in the context of a course on computational social choice at the University of Amsterdam in 2020. We acknowledge the contribution of Gerson Foks during that initial stage. We also would like to thank several anonymous reviewers for their feedback.

References

1. Aziz, H., Lee, B.E., Talmon, N.: Proportionally representative participatory budgeting: Axioms and algorithms. In: Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS) (2018)
2. Aziz, H., Shah, N.: Participatory budgeting: Models and approaches. In: Pathways between Social Science and Computational Social Science: Theories, Methods and Interpretations. Springer (2020)
3. Benade, G., Nath, S., Procaccia, A.D., Shah, N.: Preference elicitation for participatory budgeting. *Management Science* **67**(5), 2813–2827 (2021)
4. Brandt, F., Conitzer, V., Endriss, U., Lang, J., Procaccia, A.D. (eds.): Handbook of Computational Social Choice. Cambridge University Press (2016)
5. Faliszewski, P., Skowron, P., Slinko, A., Talmon, N.: Multiwinner voting: A new challenge for social choice theory. In: Endriss, U. (ed.) Trends in Computational Social Choice, chap. 2, pp. 27–47. AI Access (2017)
6. Fluschnik, T., Skowron, P., Triphaus, M., Wilker, K.: Fair knapsack. In: Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI) (2019)
7. Garey, M.R., Johnson, D.S.: Computers and Intractability: A Guide to the Theory of NP-Completeness. W.H. Freeman and Co. (1979)
8. Geist, C., Peters, D.: Computer-aided methods for social choice theory. In: Endriss, U. (ed.) Trends in Computational Social Choice, chap. 13, pp. 249–267. AI Access (2017)
9. Goel, A., Krishnaswamy, A.K., Sakshuwong, S., Aitamurto, T.: Knapsack voting: Voting mechanisms for participatory budgeting. *ACM Transaction on Economics and Computation* **7**(2), 8 (2019)
10. Goldfrank, B.: Lessons from Latin America’s experience with participatory budgeting. In: Shah, A. (ed.) Participatory Budgeting, chap. 3, pp. 91–126. Public Sector Governance and Accountability Series, The World Bank, Washington, DC (2007)
11. Hershkowitz, D.E., Kahng, A., Peters, D., Procaccia, A.D.: District-fair participatory budgeting. In: Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI) (2021)
12. Jain, P., Sornat, K., Talmon, N.: Participatory budgeting with project interactions. In: Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI) (2020)
13. Jain, P., Sornat, K., Talmon, N., Zehavi, M.: Participatory budgeting with project groups. In: Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI) (2021)
14. Janson, S.: Phragmén’s and Thiele’s election methods. arXiv preprint 1611.08826 (2016)
15. Kellerer, H., Pferschy, U., Pisinger, D.: Multidimensional knapsack problems. In: Knapsack Problems, chap. 9, pp. 235–283. Springer (2004)
16. Lackner, M., Maly, J., Rey, S.: Fairness in long-term participatory budgeting. In: Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI) (2021)
17. Laruelle, A.: Voting to select projects in participatory budgeting. *European Journal of Operational Research* **288**(2), 598–604 (2021)
18. Patel, D., Khan, A., Louis, A.: Group fairness for knapsack problems. In: Proceedings of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS) (2021)

19. Peters, D.: Proportionality and strategyproofness in multiwinner elections. In: Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS) (2018)
20. Peters, D.: Fair Division of the Commons. DPhil thesis, University of Oxford (2019)
21. Peters, D., Pierczynski, G., Skowron, P.: Proportional participatory budgeting with additive utilities. In: Proceedings of the 35th Annual Conference on Neural Information Processing Systems (NeurIPS) (2021)
22. Rey, S., Endriss, U., de Haan, R.: Designing participatory budgeting mechanisms grounded in judgment aggregation. In: Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning (KR) (2020)
23. Rey, S., Endriss, U., de Haan, R.: Shortlisting rules and incentives in an end-to-end model for participatory budgeting. In: Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI) (2021)
24. Rose, J., Omolo, A.: Six Case Studies of Local Participation in Kenya. The World Bank, Washington, DC (2013), available at <http://hdl.handle.net/10986/17556>
25. Sánchez-Fernández, L., Fernández-García, N., Fisteus, J.A., Brill, M.: The maximin support method: An extension of the d'Hondt method to approval-based multiwinner elections. *Mathematical Programming* (2022), in press
26. Schrijver, A.: *Theory of Linear and Integer Programming*. John Wiley & Sons (1998)
27. Shah, A. (ed.): *Participatory Budgeting*. Public Sector Governance and Accountability Series, The World Bank, Washington, DC (2007)
28. Talmon, N., Faliszewski, P.: A framework for approval-based budgeting methods. In: Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI) (2019)