

A Framework for Approval-Based Budgeting Methods

Nimrod Talmon

Ben-Gurion University
Be'er Sheva, Israel
talmonn@bgu.ac.il

Piotr Faliszewski

AGH University
Krakow, Poland
faliszew@agh.edu.pl

Abstract

We define and study a general framework for approval-based budgeting methods and compare certain methods within this framework by their axiomatic and computational properties. Furthermore, we visualize their behavior on certain Euclidean distributions and analyze them experimentally.

Introduction

Participatory budgeting (Cabannes 2004), initiated by the Brazil workers' party (Wainwright 2003), is gaining increased attention, and is currently applied on many continents, including North America (Gilman 2012) and Europe (e.g., Paris is organizing one of the largest citywide participatory budgeting processes¹). The general premise of participatory budgeting is to let residents of a municipality influence the way by which their common funds are being distributed, through a deliberative grassroots process. Concretely, residents are participating in constructing the municipal budget, by acting as voters and specifying their preferences over a set of available items; then, an aggregation mechanism (i.e., a budgeting method) is applied to decide upon the exact set of items to be funded.

Even though more and more funds are decided through participatory budgets, not many budgeting methods have been proposed, and no mathematical frameworks to allow for a systematic comparison between budgeting methods are available, rendering their use somewhat ad-hoc. Here we describe such a general framework for the approval-based setting, in which voters specify subsets of the available items which they approve of. Corresponding methods within our framework differentiate in two aspects: (1) The way by which voter satisfaction from a given set of funded items is defined, modeled through a *satisfaction function*; and (2) The way these satisfaction functions are used.

We consider several concrete methods within our framework, including some which are used in practice and some which generalize known multiwinner voting rules. To compare these methods, we (1) consider their computational complexity; (2) define several axioms, relevant to budgeting methods, and study how well these axioms are satisfied

by the methods at hand; and (3) report on three experiments: In the first two, we visualize the behavior of these methods on certain Euclidean preferences, by adapting the methodology of Elkind et al. (2017a), originally developed for multiwinner voting rules; we are specifically interested in the behavior of these methods with respect to their preference for choosing either cheap items or expensive ones. In the third experiment we assess how well our budgeting methods deal with local and global items.

The main contributions of our work are: (1) a general framework of approval-based budgeting methods with several methods within; (2) several useful axiomatic properties which are relevant to budgeting methods at large; (3) an adaptation of the methodology of Elkind et al. (2017a) to participatory budgeting; and (4) an evaluation of certain methods within our framework according to their axiomatic, computational, and visual properties, as well as their ability to deal with local and global items.

Due to space constraints, some proof details are omitted.

Related Work

Researchers have considered ordinal-based budgeting methods, in which voters rank budget items (Goel et al. 2015; Shapiro and Talmon 2017); utility-based budgeting methods, in which voters have numerical utilities over budget items (Fluschnik et al. 2017; Benade et al. 2017); and, as we do, approval-based budgeting methods, in which each voter approves a set of items (Goel et al. 2015; Goel, Krishnaswamy, and Sakshuwong 2016; Benade et al. 2017; Aziz, Lee, and Talmon 2017). Specifically, Goel et al. (2015; 2016) study k -Approval, where each voter approves exactly k items, and Knapsack voting, where each voter approves items while respecting the budget limit. As shown below, the aggregation method used by Goel et al., in which the winning bundle is selected by greedily considering the items in decreasing number of approvals, fits within our framework.

Benade et al. (2017) consider the implicit utilitarian model (Caragiannis et al. 2017) and analyze the distortion achieved by eliciting the preferences of the voters by Knapsack voting and by Threshold voting (in which voters are asked to approve those items which they evaluate above a given threshold); their distortion-based aggregation methods do not fit within our framework. Aziz et al. (2017) generalize multiwinner proportionality axioms to the setting of partic-

Copyright © 2019, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹<https://budgetparticipatif.paris.fr/bp>

ipatory budgeting. Their budgeting method, which is a generalization of Phragmén’s sequential method for committee elections, does not fit within our framework.

As we are interested in axiomatic properties of budgeting methods, we mention the work of Shapiro and Talmon (2017), which considers a generalization of the Condorcet principle, the paper of Aziz et al. (2017), which considers axioms of representation, and the papers of Fluschnik et al. (2017) and Fain et al. (2018), which also consider representation. Our framework of approval-based budgeting methods is a generalization of the framework studied by Lackner and Skowron (2017), which regards approval-based multiwinner methods. Our framework, as well as theirs, might be seen as the approval-based variant of the framework of committee scoring rules (Elkind et al. 2017b; Faliszewski et al. 2016b; 2016a). For a general survey of approval-based voting, we point the reader to the work of Laslier and Sanver (2010).

Budgeting Methods

Budgeting Scenarios

We consider the following model of participatory budgeting. A *budgeting scenario* E is a tuple $E = (A, V, c, \ell)$ where $A = \{a_1, \dots, a_m\}$ is a set of items, $c: A \rightarrow \mathbb{N}$ is a *cost function*, so that the *cost* of item $a \in A$ is $c(a)$, $V = \{v_1, \dots, v_n\}$ is a set of voters, where each voter $v \in V$ specifies her approval set $A_v \subseteq A$, containing those items which she approves of, and $\ell \in \mathbb{N}$ is a budget limit.

A *budgeting method* \mathcal{R} is a function which takes a budgeting scenario $E = (A, V, c, \ell)$ and returns a *bundle* $B \subseteq A$, such that the total cost of the items of B respects the budget limit; i.e., slightly abusing notation, it must hold that $c(B) = \sum_{b \in B} c(b) \leq \ell$. The winning bundle (i.e., the set of funded items) for a budgeting scenario E , under a budgeting method \mathcal{R} , is denoted by $\mathcal{R}(E)$. With respect to a winning bundle B , an item is *budgeted* (or *funded*) if it is contained in B . For simplicity, we ignore issues related to tie-breaking, which can be dealt with using standard techniques.

Satisfaction Functions

Budgeting methods in our framework operate by considering the satisfaction of the voters from possible bundles. To measure the satisfaction of a voter from a possible winning bundle, we use satisfaction functions, defined below.

Definition 1 (Satisfaction function). A *satisfaction function* f is a function $f: 2^A \times 2^A \rightarrow \mathbb{N}$, where A is a set of items. Given a budgeting scenario with a set of items A , a voter $v \in V$ with her approval set A_v , and a bundle $B \subseteq A$, the value $f(A_v, B)$ is the *satisfaction* of v from the bundle B .

Example 1. Consider a budgeting scenario with $A = \{a, b, c\}$, a voter v with $A_v = \{a, b\}$, and a bundle $B = \{b, c\}$. For a satisfaction function f , the value $f(A_v, B) = f(\{a, b\}, \{b, c\})$ is the satisfaction of v from the bundle B .

Below we give some satisfaction functions. For a bundle B and a voter v with her approval set A_v , let $B_v := A_v \cap B$.

1. $f(A_v, B) = |B_v|$: The satisfaction of a voter is the number of funded items she approves of.

2. $f(A_v, B) = \sum_{a \in B_v} c(a) = c(B_v)$: A voter’s satisfaction is the total cost of her approved items which are funded. This follows the intuition that there is some positive correlation between the satisfaction of a voter and the funds spent on items approved by her.
3. $f(A_v, B) = \mathbb{1}_{|B_v| > 0}$: A voter has satisfaction 0 if none of her approved items is funded, and 1 otherwise (i.e., if at least one of her approved items is funded).

Using Satisfaction Functions

A given satisfaction function can be used in various ways. We consider the following three approaches.

1. **Max rules:** For a satisfaction function f , the rule \mathcal{R}_f^m selects, as a winning bundle, a bundle which maximizes the sum of voters’ satisfaction, according to f . Formally, \mathcal{R}_f^m selects $\arg \max_{B \subseteq A} \sum_{v \in V} f(A_v, B)$.
2. **Greedy rules:** For a satisfaction function f , the rule \mathcal{R}_f^g proceeds in iterations, maintaining a partial bundle B , where in each iteration it adds an item a to B which maximizes the value $\sum_{v \in V} f(A_v, B \cup \{a\})$. Notice that it halts only when it is not possible to add more items.
3. **Proportional greedy rules:** The rule \mathcal{R}_f^p is similar to \mathcal{R}_f^g , except that in each iteration it adds an item a to B which maximizes:

$$\left(\sum_{v \in V} f(A_v, B \cup \{a\}) - \sum_{v \in V} f(A_v, B) \right) / c(a).$$

Remark 1. So, while the Max rules optimally maximize the sum of satisfactions of the voters, the Greedy rules operate by greedily maximizing the sum of satisfactions, without considering the costs of the items, and the Proportional Greedy rules differ from the Greedy ones in that they operate by greedily maximizing the sum of “satisfaction divided by cost” (i.e., they greedily maximize “bang-per-the-buck”).

Example 2. Consider a budgeting scenario with items a_2, a_3, a_4, a_5 , and a_6 , where, for $2 \leq i \leq 6$, item a_i costs i , budget limit $\ell = 10$, and voters $v_1 = \{a_2, a_5, a_6\}$, $v_2 = \{a_2, a_3, a_4, a_5\}$, $v_3 = \{a_3, a_4, a_5\}$, $v_4 = \{a_4, a_5\}$, and $v_5 = \{a_6\}$. Then, the winning bundle for this budgeting scenario under $\mathcal{R}_{|B_v|}^m$ is $\{a_2, a_3, a_5\}$ with total score 8, under $\mathcal{R}_{|B_v|}^g$ and $\mathcal{R}_{|B_v|}^p$ is $\{a_4, a_5\}$ (but $\mathcal{R}_{|B_v|}^g$ selects first a_5 and then a_4 while $\mathcal{R}_{|B_v|}^p$ selects a_4 and then a_5), under $\mathcal{R}_{\mathbb{1}_{|B_v| > 0}}^m$ is $\{a_4, a_6\}$ which gets the highest total satisfaction possible, under $\mathcal{R}_{\mathbb{1}_{|B_v| > 0}}^g$ is $\{a_2, a_3, a_5\}$ (where a_5 is selected first and then perhaps a_3 and a_2 , depending on the tie breaking), under $\mathcal{R}_{\mathbb{1}_{|B_v| > 0}}^p$ is $\{a_2, a_3, a_4\}$ (where first a_2 is selected, then a_4 , and then perhaps a_3 , depending on the tie breaking), and under $\mathcal{R}_{c(B_v)}^m$, $\mathcal{R}_{c(B_v)}^g$ and $\mathcal{R}_{c(B_v)}^p$ is $\{a_4, a_5\}$.

Remark 2. The above satisfaction functions and approaches result in nine budgeting methods which we discuss throughout the paper. Indeed, studying other functions and approaches is an immediate future work direction. We chose those functions and approaches as they are natural, generalize known multiwinner voting rules, and include some known budgeting methods: $\mathcal{R}_{\mathbb{1}_{|B_v| > 0}}^m$ generalizes approval-based Chamberlin–Courant (CC) rule (Chamberlin and

Courant 1983; Procaccia, Rosenschein, and Zohar 2008) and $\mathcal{R}_{\mathbb{1}_{|B_v|>0}}^g$ generalizes the greedy approximation of this rule (Lu and Boutilier 2011). Furthermore, $\mathcal{R}_{|B_v|}^g$ is similar to the popular k -Approval and Knapsack voting (Goel et al. 2015) (the aggregation method is the same, albeit k -Approval and Knapsack restrict the voter approval sets; we do not consider such restrictions in our framework, as we are interested in the aggregation).

Below, we study the computational and axiomatic properties of our budgeting methods, and report on simulations.

Budgeting Algorithms

We consider the computational complexity of identifying winning bundles. First, it follows from the definitions of the Greedy rules and the Proportional greedy rules that computing their winners, given that the functions used can be computed efficiently, can be done in polynomial time; this holds as these rules are defined through efficient iterative processes. This is not the case for Max rules, which in general are NP-hard. (To be concrete, next we consider a specific Max rule; to be formally correct, we show NP-hardness for deciding whether a bundle with at least a given total satisfaction, i.e., sum of satisfaction values, exists.)

Observation 1. Given a budgeting scenario and a bound s , deciding whether a feasible bundle B for which $\sum_{v \in V} \mathbb{1}_{|B_v|>0} \geq s$ exists is NP-hard.

Proof. Notice that $\mathcal{R}_{\mathbb{1}_{|B_v|>0}}^m$ generalizes the approval-based CC rule, which is NP-hard (Procaccia, Rosenschein, and Zohar 2008). Also, when all items are of unit cost, the problem is equivalent to Max Cover (Garey and Johnson 1979; Skowron and Faliszewski 2017). \square

Next we consider other Max rules. The proof of the next results follows by the similarity of the problem to the Knapsack problem which is solvable in polynomial-time whenever one of the dimensions is given in unary.

Proposition 1. Identifying a winning bundle under $\mathcal{R}_{|B_v|}^m$ can be done in polynomial time.

Proposition 2. Identifying winning bundles under $\mathcal{R}_{c(B_v)}^m$ can be done in pseudopolynomial time. Further, given a budgeting scenario and a bound s , deciding whether there is a feasible bundle B with $\sum_{v \in V} f(A_v, B) \geq s$, for $f(A_v, B) = \sum_{a \in B_v} c(a)$, is weakly NP-hard.

Proof. Weak NP-hardness follows by a straightforward reduction from the SUBSET SUM problem (see, e.g., the text of Garey and Johnson (1979)). Given a Subset Sum instance with integers $X = \{x_1, \dots, x_n\}$, where the existence of a subset $X' \subseteq X$ with $\sum_{x \in X'} x = Z$ is to be decided, we construct the following budgeting scenario: For each integer x_i , we create a voter v_i approving an item a_i of cost $c(a_i) = x_i$, and set the limit to be $\ell = Z$. For a yes-instance of Subset Sum, a winning bundle shall be of total cost Z , thus weak NP-hardness follows.

The reduction above also hints on the pseudopolynomial-time algorithm for $\mathcal{R}_{c(B_v)}^m$: Apply dynamic programming

similar to that for Subset Sum, by iterating over the items and remembering the maximum total satisfaction that can be achieved for each amount of money. \square

Coping with Intractability

We describe a simple Integer Linear Program (ILP) for Max rules over satisfaction functions which can be defined using ILPs; note that all satisfaction functions considered here (i.e., $|B_v|$, $\mathbb{1}_{|B_v|>0}$, and $c(B_v)$) can be defined using ILPs. This is useful due to the availability of efficient ILP solvers. Indeed, the simulations reported below were performed using the Gurobi ILP solver (Gurobi Optimization 2018) on such ILP formulations.

Observation 2. Let f be a satisfaction function which can be formulated using linear constraints. Then, identifying a winning bundle under \mathcal{R}_f^m can be done using an ILP.

Next is a general, parameterized complexity result.

Observation 3. For all Max rules, the problem of deciding whether a winning bundle of total cost at least a given value is fixed-parameter tractable for the number m of items, but there are some Max rules for which this problem is hard even when $n = 1$, where n is the number of voters.

For $\mathcal{R}_{c(B_v)}^m$, which is weakly NP-hard (see Proposition 2) even when there is only one voter (see the proof of Observation 3), we have a pseudopolynomial time algorithm (see the proof of Proposition 2) and an FPTAS, as we show next.

Observation 4. There is an FPTAS for $\mathcal{R}_{c(B_v)}^m$.

As for $\mathcal{R}_{\mathbb{1}_{|B_v|>0}}^m$, which is NP-hard (see Observation 1), there is no approximation algorithm with approximation ratio better than $1 - 1/e$, as it generalizes the multiwinner voting rule CC for approval elections, which itself is equivalent to the Max Cover problem (Feige 1998; Skowron and Faliszewski 2017); next we show that it is fixed-parameter tractable for the number n of the voters.

Proposition 3. $\mathcal{R}_{\mathbb{1}_{|B_v|>0}}^m$ is fixed-parameter tractable for n .

Proof. Guess the partition of the voters with the intended meaning that each group of voters in the partition is represented by the same item. For each such group, guess the number of voters which would be satisfied. Then, go over all items and pick as representative the cheapest item that makes exactly this number of voters in this group satisfied. \square

Budgeting Axioms

In this section we suggest several axiomatic properties which are relevant to budgeting methods. In particular, we focus on axioms which relate to the costs of the items. For each axiom, after providing the definition we check which of the rules in our framework satisfy it. We stress that we view axioms as properties and not necessarily as normative recommendations; thus, while we demonstrate certain situations in which each of the axioms is desirable, there are situations in which failing an axiom is not a drawback.

Our first axiom models the very natural expectation that if within a budgeting scenario we can afford to fund more

	$\mathcal{R}_{ B_v }^m$	$\mathcal{R}_{ B_v }^g$	$\mathcal{R}_{ B_v }^p$	$\mathcal{R}_{\mathbb{1}_{ B_v >0}}^m$	$\mathcal{R}_{\mathbb{1}_{ B_v >0}}^g$	$\mathcal{R}_{\mathbb{1}_{ B_v >0}}^p$	$\mathcal{R}_{c(B_v)}^m$	$\mathcal{R}_{c(B_v)}^g$	$\mathcal{R}_{c(B_v)}^p$
Complexity	P	P	P	NP-h	P	P	weak NP-h	P	P
Inclusion Maximality	✓	✓	✓	✓	✓	✓	✓	✓	✓
Discount Monotonicity	✓	✓	✓	✓	✓	✓	×	×	×
Splitting Monotonicity	✓	✓	✓	✓	✓	✓	✓	×	✓
Merging Monotonicity	×	×	×	✓	✓	×	✓	✓	✓
Limit Monotonicity	×	×	×	×	×	×	×	×	×

Table 1: Computational and axiomatic properties of certain approval-based budgeting methods.




Table 2: Results of Experiment 1. Cheap items are on the left, expensive on the right.

items, then we should; i.e., voters might wish to use all the available funds. Formally, we express it as follows.

Definition 2 (Inclusion Maximality). A budgeting method \mathcal{R} satisfies Inclusion Maximality if for each budgeting scenario $E = (A, V, c, \ell)$ and each pair of feasible bundles B and B' such that $B \subset B'$ it holds that if B is winning then also B' is winning.

Inclusion Maximality seems desirable in all cases of participatory budgeting. Yet, a somewhat artificial example shows that not all rules in our framework satisfy it.

Example 3. Let $f(A_v, B) = \min(\{c(a) : a \in B_v\})$; that is, the satisfaction of a voter equals the cost of her cheapest approved item which is funded. Then, \mathcal{R}_f^m does not satisfy Inclusion Maximality (to see it, consider two items a, b of cost 1, 2 respectively, one voter approving both a and b , and a pair of bundles $B = \{b\}$ and $B' = \{b, a\}$).

Nevertheless, notice that all three satisfaction functions we consider here (namely $f = |B_v|$, $f = \mathbb{1}_{|B_v|>0}$, and $f = c(B_v)$) are super-set monotone; that is, for each of them, we have that $f(B') \geq f(B)$ for each $B \subset B'$. Thus, since one can verify that Max rules are Inclusion Maximal for super-set monotone satisfaction functions, we have the following.

Corollary 1. $\mathcal{R}_{|B_v|}^m$, $\mathcal{R}_{c(B_v)}^m$, and $\mathcal{R}_{\mathbb{1}_{|B_v|>0}}^m$, satisfy Inclusion Maximality.

For Greedy rules and Proportional greedy rules, it is never the case that two feasible bundles B, B' with $B \subset B'$ are both winning, so we have the following.

Corollary 2. $\mathcal{R}_{|B_v|}^g$, $\mathcal{R}_{c(B_v)}^g$, $\mathcal{R}_{\mathbb{1}_{|B_v|>0}}^g$, $\mathcal{R}_{|B_v|}^p$, $\mathcal{R}_{c(B_v)}^p$, and $\mathcal{R}_{\mathbb{1}_{|B_v|>0}}^p$, satisfy Inclusion Maximality.

In the next axiom we consider the response of our rules to increasing the available budget limit. Specifically, we require that if we increase the limit, then all funded items remain funded, provided that no new item becomes affordable (this last condition is quite natural; if there is an item that all the voters approve, which is above the budget limit before its extension but is within the limit after the extension, then it is quite natural that this item might be funded and might remove many previously funded ones).

Definition 3 (Limit Monotonicity). We say that a budgeting method \mathcal{R} satisfies Limit Monotonicity if for each pair of budgeting scenarios $E = (A, V, c, \ell)$, $E' = (A, V, c, \ell + 1)$ with no item which costs exactly $\ell + 1$, for each $a \in A$, it holds that $a \in \mathcal{R}(E) \implies a \in \mathcal{R}(E')$.

Interestingly, while Limit Monotonicity generalizes the Committee Monotonicity axiom of multiwinner voting rules (Elkind et al. 2017b), which is satisfied by many such rules, Limit Monotonicity is not satisfied by any of the rules we consider here. This shows that even though the frameworks of participatory budgeting and multiwinner voting are strongly related, the fact that the former deals with items of different costs can lead to significant differences.

Proposition 4. None of $\mathcal{R}_{|B_v|}^m$, $\mathcal{R}_{c(B_v)}^m$, and $\mathcal{R}_{\mathbb{1}_{|B_v|>0}}^m$ satisfies Limit Monotonicity.

x	$\mathcal{R}_{ B_v }^m$	$\mathcal{R}_{ B_v }^g$	$\mathcal{R}_{ B_v }^p$	$\mathcal{R}_{\mathbb{1}_{ B_v >0}}^m$	$\mathcal{R}_{\mathbb{1}_{ B_v >0}}^g$	$\mathcal{R}_{\mathbb{1}_{ B_v >0}}^p$	$\mathcal{R}_{c(B_v)}^m$	$\mathcal{R}_{c(B_v)}^g$	$\mathcal{R}_{c(B_v)}^p$
0									
10									
50									
70									
100									

Table 3: Results of Experiment 2. Cheap items are on the left, expensive on the right.

Proof. Below we consider each of the three rules separately.

$\mathcal{R}_{\mathbb{1}_{|B_v|>0}}^m$: Consider a budgeting scenario E with items a , b , and c , all of unit cost, and four voters: $v_1 : \{a\}$; $v_2 : \{a, b\}$; $v_3 : \{b, c\}$; $v_4 : \{c\}$. Then, with budget limit 1 a winning bundle might be $\{b\}$, while with budget limit 2 the only winning bundle is $\{a, c\}$.

$\mathcal{R}_{c(B_v)}^m$: Consider a budgeting scenario E with items a, b, c , and d , with costs 2, 3, 3, and 5, respectively, and one voter: $v_1 : \{a, b, c, d\}$. Then, with budget limit 6 the only winning bundle is $\{b, c\}$, while with budget limit 7 the only winning bundle is $\{a, d\}$.

$\mathcal{R}_{|B_v|}^m$: Consider a budgeting scenario E with the same items as above and five voters: $v_1 : \{a, b, c, d\}$; $v_2 : \{a, b, c, d\}$; $v_3 : \{b, c, d\}$; $v_4 : \{d\}$; $v_5 : \{d\}$. Then, with budget limit 6 the only winning bundle is $\{b, c\}$, while with budget limit 7 the only winning bundle is $\{a, d\}$. \square

Greedy and Proportional greedy rules also fail Limit Monotonicity. Consider a budgeting scenario with items a, b , and c , where a has the largest value according to the relevant satisfaction function; thus, a is selected in the first iteration. Set the cost of a so that, after selecting it, the remaining budget limit is such that, for the original budget limit ℓ , only b can be selected, while for the budget limit $\ell + 1$, also c can be selected. Set the value of c to be higher than that of b . Thus, b is selected in the first case and c in the second case.

Corollary 3. $\mathcal{R}_{|B_v|}^g$, $\mathcal{R}_{c(B_v)}^g$, $\mathcal{R}_{\mathbb{1}_{|B_v|>0}}^g$, $\mathcal{R}_{|B_v|}^p$, $\mathcal{R}_{c(B_v)}^p$, and $\mathcal{R}_{\mathbb{1}_{|B_v|>0}}^p$, do not satisfy Limit Monotonicity.

As neither of our rules satisfies Limit Monotonicity, it is natural to wonder if the axiom is truly useful. One argument in its favor is that it would be surprising for voters to learn that due to increasing the available funds their approved item

ceases to be funded. Thus if one uses a budgeting method that fails Limit Monotonicity then one should make sure that the budget is fixed ahead of voting and is neither increased nor decreased later (as it leads to suspicion of manipulation).

A budgeting method \mathcal{R} satisfies Discount Monotonicity if any funded item remains funded when its price decreases. This is a very desirable property as failing it means that people proposing new items for the participatory budget have to think strategically about the item's price, instead of trying to minimize it.

Definition 4 (Discount Monotonicity). A budgeting method \mathcal{R} satisfies Discount Monotonicity if for each budgeting scenario $E = (A, V, c, \ell)$ and for each $b \in \mathcal{R}(E)$, it holds that $b \in \mathcal{R}(E')$ for $E' = (A, V, c', \ell)$, where for each item $a \in A$, we have that $c'(a) = c(a)$ whenever $a \neq b$, and $c'(b) = c(b) - 1$.

Proposition 5. $\mathcal{R}_{|B_v|}^m$, $\mathcal{R}_{\mathbb{1}_{|B_v|>0}}^m$, $\mathcal{R}_{|B_v|}^g$, $\mathcal{R}_{\mathbb{1}_{|B_v|>0}}^g$, $\mathcal{R}_{|B_v|}^p$, and $\mathcal{R}_{\mathbb{1}_{|B_v|>0}}^p$, satisfy Discount Monotonicity, while $\mathcal{R}_{c(B_v)}^m$, $\mathcal{R}_{c(B_v)}^g$, and $\mathcal{R}_{c(B_v)}^p$ fail it.

Proof sketch. Intuitively, for $f = |B_v|$ and $f = \mathbb{1}_{|B_v|>0}$, decreasing the cost only increases the attractiveness of the item, while for $f = c(B_v)$, decreasing the cost makes the item less attractive. \square

The next two axioms regard the robustness of a voting rules with respect to the situation of a person proposing a new item, provided that this new item has some internal structure and can be presented either as a single one or as several items (e.g., renovation of a school can either be a single project, or several ones, including painting the interior, painting the exterior, buying new furniture etc.). We consider splitting and merging items.

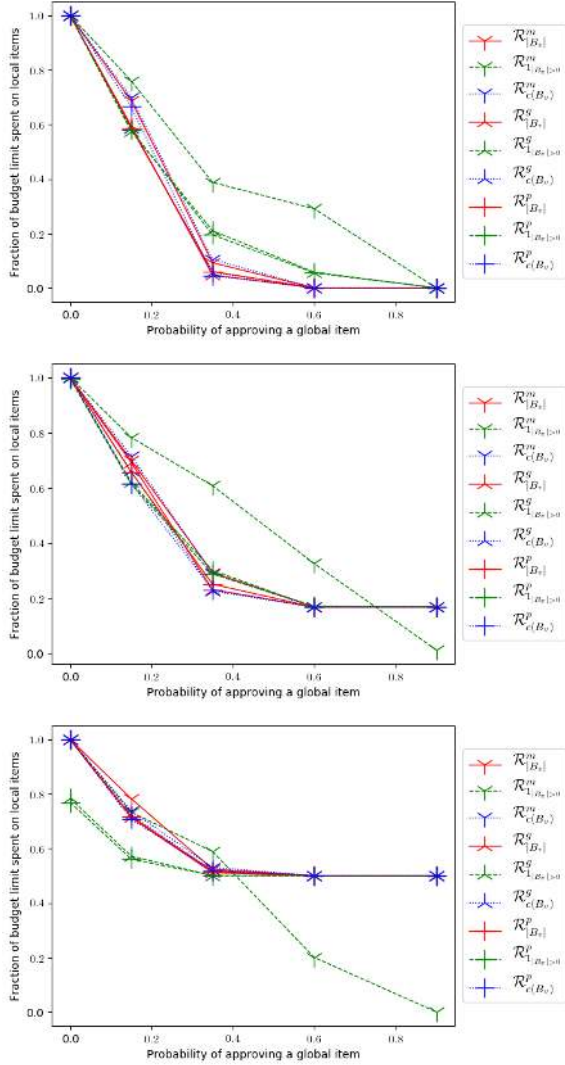


Table 4: Results of Experiment 3. The top, middle, and bottom figure are for $\ell = 20, 30$, and 50 , respectively.

Definition 5 (Splitting Monotonicity). A budgeting method \mathcal{R} satisfies Splitting Monotonicity if for each budgeting scenario $E = (A, V, c, \ell)$, for each $a \in \mathcal{R}(E)$, and for each budgeting scenario E' which is formed by splitting a into a set of items A' with $c(a) = c(A')$, and such that the voters which approve a in E approve all items of A' in E' and no other voters approve items of A' , it holds that $\mathcal{R}(E') \cap A' \neq \emptyset$.

Remark 3. Note that in the above definition we assume that if a voter approves an item, then this voter also approves all the projects that result from splitting it. While in some cases this is true (e.g., when the voters recognize that several items form a larger project), in some others it would be false (e.g., when voters decide that completing only a part of a larger project is sufficient and remaining funds should be used for other purposes). In practice we expect both types of behavior to occur and it would be natural to study variants

of Splitting Monotonicity that take such more complicated behaviors into account.

Proposition 6. $\mathcal{R}_{|B_v|}^m, \mathcal{R}_{\mathbb{1}_{|B_v|>0}}^m, \mathcal{R}_{|B_v|}^g, \mathcal{R}_{\mathbb{1}_{|B_v|>0}}^g, \mathcal{R}_{|B_v|}^p, \mathcal{R}_{\mathbb{1}_{|B_v|>0}}^p, \mathcal{R}_{c(B_v)}^m$, and $\mathcal{R}_{c(B_v)}^p$, satisfy Splitting Monotonicity, while $\mathcal{R}_{c(B_v)}^g$ does not satisfy Splitting Monotonicity.

Proof sketch. Intuitively, $\mathcal{R}_{c(B_v)}^g$ does not satisfy Splitting Monotonicity as the new items' value is less than the original item's value. For other rules, the new items' value is at least as the original item's value, thus at least one is selected. \square

Definition 6 (Merging Monotonicity). A budgeting method \mathcal{R} satisfies Merging Monotonicity if for each budgeting scenario $E = (V, A, c, \ell)$, and for each $A' \subseteq \mathcal{R}(E)$ such that for each $v \in V$ we either have $v \cap A' = \emptyset$ or $A' \subseteq A_v$, it holds that $a \in \mathcal{R}(E')$ for $E' = (A \setminus A' \cup \{a\}, V', c', \ell)$, $c'(a) = \sum_{a \in A'} c(a)$, and each voter $v \in V$ for which $A' \subseteq A_v$ in E , approves a in E' , and no other voter approves a .

Proposition 7. $\mathcal{R}_{\mathbb{1}_{|B_v|>0}}^m, \mathcal{R}_{c(B_v)}^m, \mathcal{R}_{\mathbb{1}_{|B_v|>0}}^g, \mathcal{R}_{c(B_v)}^g$, and $\mathcal{R}_{c(B_v)}^p$, satisfy Merging Monotonicity, while $\mathcal{R}_{|B_v|>0}^p, \mathcal{R}_{|B_v|}^m, \mathcal{R}_{|B_v|}^g$, and $\mathcal{R}_{|B_v|}^p$ fail it.

Proof sketch. Intuitively, for the satisfaction function $f(A_v, B) = |B_v|$, the value of the original items decreases but the merged item is as expensive. For $f(A_v, B) = c(B_v)$, the value of the merged item equals the total value of the merged items. For $f(A_v, B) = \mathbb{1}_{|B_v|>0}$ the fact that all the original items were funded means that the merged item still satisfies the same voters. \square

As in the case of Discount Monotonicity, the main advantage of rules satisfying Splitting and Merging Monotonicity axioms is that it reduces the need for strategic thinking on the part of the project proposers.

Experiments on Budgeting Methods

In this section we report on three experiments: Two experiments generalize the technique of Elkind et al. (2017a), used quite extensively for multiwinner elections (Faliszewski, Szufa, and Talmon 2018; Aziz et al. 2018; Faliszewski and Talmon 2018), to the setting of approval-based budgeting scenarios. One experiment focuses on locality issues, assuming certain projects are more relevant to certain voters and others are relevant to all.

In each experiment we consider the 2-dimensional Euclidean domain, where both voters and items correspond to ideal points on a 2-dimensional plane; for clarity, we set this 2-dimensional plane to be of width 1 and height 1, where $(0, 0)$, $(0, 1)$, $(1, 0)$, and $(1, 1)$, denote, respectively, the left-bottom point, the left-top point, the right-bottom point, and the right-top point. Concretely, each particular simulation setting consists of: (1) a distribution of the ideal points of the voters; (2) a distribution of the ideal points of the items; (3) a distribution of the item costs; (4) a budget limit; and (5) a threshold function which, based on the positions and costs, creates approval sets for the voters.

For the first two experiments, for each particular simulation setting we generate several corresponding budgeting scenarios by sampling from these distributions, compute the winning bundle in each of them, and aggregate the results into a 2-dimensional histogram. These histograms are formed by first partitioning the square from $(0, 0)$ to $(1, 1)$ into bins (we use 50×50 bins). Then, we draw a pixel in each of these bins, where the more funds spent on this bin, the brighter the pixel is drawn. Specifically, denoting the total funds used by y and the funds used in a certain bin by x , we normalize these values using the formula $\frac{\arctan(x/0.0005*y)}{\pi/2}$.

Experiment 1 and 2 (depicted in Table 2 and 3, resp.). We describe Experiment 1 and 2 together, as they share a lot in common. We have:

1. voters, positioned uniformly on a disc of radius 0.3, centered at position $(0.5, 0.5)$; we have 50 such voters for Experiment 1 and 100 for Experiment 2;
2. 50 *cheap* items, positioned uniformly on a disc of radius 0.2, centered at $(0.3, 0.5)$, and 50 *expensive* items, positioned uniformly on a disc of radius 0.2, centered at $(0.7, 0.5)$;
3. the cheap items cost 10 each, while the expensive items cost 100 each for Experiment 2, and for Experiment 1, we use a parameter x for the cost of the expensive items (so Table 2 shows histograms for various values of x);
4. the budget limit is 1000 for Experiment 1 and 200 for Experiment 2;
5. for Experiment 1: Each voter approves the 10 items which are the closest to her; for Experiment 2: The approval sets of the voters are generated with respect to a parameter x , as follows: for each cheap item, we identify the 5 voters which are the closest to it, and add the item to their approval sets; for each expensive item, we identify the x voters which are the closest to it, and add the item to their approval sets.

Each histogram is an aggregation of 100 single elections. Informally speaking, the approval sets in Experiment 1 are generated “from the point-of-view of the voters”, while in Experiment 2 they are generated “from the point-of-view of the items”.

Experiment 3 focuses on the issue of global and local items: Consider, e.g., a budgeting scenario with some city-level projects and some neighborhood-level projects. Ideally, we would want some mix of city-level projects and neighborhood-level projects to be funded. While it is possible to achieve some mix artificially (e.g., requiring voters to select both city-level projects and neighborhood-level projects), here we are interested in finding out the natural mix that rules in our framework achieve.

Experiment 3 (depicted in Table 4). We have: (1) 20 voters, positioned uniformly on the whole 1×1 square; (2) 5 items, termed *global items*, which are also positioned uniformly on the square; and another 30 items, termed *local items*, also positioned uniformly on the square; (3) each global item and each local item costs 5; (4) we vary the budget limit between 20 and 50; (5) the approval sets of the voters are populated with respect to a parameter p , as follows:

For each pair of a voter and a global item, we let the voter approve the global item with some probability p . For each pair of a voter and a local item, we let the voter approve the local item if and only if their Euclidean distance is at most 0.2. In Table 4 each datapoint is averaged over 100 repetitions and we consider the average funds spent on local items as a function of the probability p of approving a global item.

Experimental Results

Next we discuss the results of our experiments, depicted in Tables 2, 3, and 4.

In Table 2, and focusing on the Max rules, it is visible that $\mathcal{R}_{c(B_v)}^m$ prefers expensive items the most, $\mathcal{R}_{|B_v|}^m$ prefers expensive items the least, while $\mathcal{R}_{\mathbb{1}_{|B_v|>0}}^m$ is somehow in between. This can be seen specifically by noticing that $\mathcal{R}_{|B_v|}^m$ ceases to select expensive items as soon as $x = 30$, $\mathcal{R}_{\mathbb{1}_{|B_v|>0}}^m$ ceases to select expensive items only when $x = 190$, while $\mathcal{R}_{c(B_v)}^m$ keeps on selecting expensive items even when $x = 190$. Intuitively, this is so as expensive items are useful for $\mathcal{R}_{c(B_v)}^m$ and not useful for $\mathcal{R}_{|B_v|}^m$; for $\mathcal{R}_{\mathbb{1}_{|B_v|>0}}^m$, while the costs of the items are not useful, their positions are, as the rule’s goal is to satisfy as many voters as possible. The greedy rules are also somehow in between the extremes, while, as expected, the proportional greedy rules prefer cheaper items.

This general behavior is consistent with Experiment 2, as can be seen in Table 3. Specifically, observe that $\mathcal{R}_{c(B_v)}^m$ switches to funding only expensive items as soon as $x = 10$, while $\mathcal{R}_{|B_v|}^m$ switches to funding only expensive items only when $x = 70$; due to the positions of the expensive items, and due to the fact that they cover different sets of voters, $\mathcal{R}_{\mathbb{1}_{|B_v|>0}}^m$ interleave cheap items with expensive items (except for the corner cases of $x = 0$ and $x = 100$).

As for locality issues, the results depicted in Table 4 for $\ell = 20$ demonstrate that $\mathcal{R}_{c(B_v)}^m$ starts to consider global items first, with $\mathcal{R}_{|B_v|}^m$ after it, and $\mathcal{R}_{\mathbb{1}_{|B_v|>0}}^m$ being the last to consider global items. As ℓ increases, $\mathcal{R}_{c(B_v)}^m$ and $\mathcal{R}_{|B_v|}^m$ select more local items while $\mathcal{R}_{\mathbb{1}_{|B_v|>0}}^m$ is not affected.

One particularly interesting observation from our experiments is that the greedy rule used by Goel (2015), $\mathcal{R}_{|B_v|}^g$ in our language, behaves substantially differently from its Max variant, $\mathcal{R}_{|B_v|}^m$. This is quite visible in the first two columns of Tables 2 and 3. The Max variant gives much more attention to the cheap items than the greedy one. Thus the choice between these two rules—already made by many users of Goel’s work—may have nonnegligible consequences.

Outlook

We have defined a framework for approval-based budgeting methods and studied nine rules within it, considered their computational and axiomatic properties, and reported on simulations to evaluate them experimentally. Our framework, and the axiomatic properties we consider, can be used to better evaluate known budgeting methods, as well as propose new budgeting methods, which might prove to have better theoretical guarantees and better practical behavior.

E.g., our results show that, while $\mathcal{R}_{|B_v|}^g$ is used extensively in practice (see, e.g., Goel et al. 2015), it produces significantly different results than the rule which it approximates, namely $\mathcal{R}_{|B_v|}^m$.

An immediate future research direction would be to study more satisfaction functions and ways of using them, which would correspond to more rules within our framework. Furthermore, defining and studying more axiomatic properties which are relevant to budgeting methods, as well as performing more extensive experimental analysis on various budgeting methods would help in better understanding these rules.

Another avenue for future research is to seek general results for rules in the framework, such as identifying classes of rules within the framework which satisfy certain axiomatic properties, and better understanding which budgeting methods reside within the framework and which do not.

Acknowledgments. Piotr Faliszewski was supported by the National Science Centre, Poland, under project 2016/21/B/ST6/01509.

References

- Aziz, H.; Faliszewski, P.; Grofman, B.; Slinko, A.; and Talmon, N. 2018. Egalitarian committee scoring rules. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI '18)*, 56–62.
- Aziz, H.; Lee, B.; and Talmon, N. 2017. Proportionally representative participatory budgeting: Axioms and algorithms. *arXiv preprint arXiv:1711.08226*.
- Benade, G.; Nath, S.; Procaccia, A. D.; and Shah, N. 2017. Preference elicitation for participatory budgeting. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI '17)*, 376–382.
- Cabannes, Y. 2004. Participatory budgeting: a significant contribution to participatory democracy. *Environment and urbanization* 16(1):27–46.
- Caragiannis, I.; Nath, S.; Procaccia, A. D.; and Shah, N. 2017. Subset selection via implicit utilitarian voting. *Journal of Artificial Intelligence Research* 58:123–152.
- Chamberlin, B., and Courant, P. 1983. Representative deliberations and representative decisions: Proportional representation and the Borda rule. *American Political Science Review* 77(3):718–733.
- Elkind, E.; Faliszewski, P.; Laslier, J.-F.; Skowron, P.; Slinko, A.; and Talmon, N. 2017a. What do multiwinner voting rules do? An experiment over the two-dimensional Euclidean domain. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI '17)*, 494–501.
- Elkind, E.; Faliszewski, P.; Skowron, P.; and Slinko, A. 2017b. Properties of multiwinner voting rules. *Social Choice and Welfare* 48(3):599–632.
- Fain, B.; Munagala, K.; and Shah, N. 2018. Fair allocation of indivisible public goods. In *Proceedings of the 2018 ACM Conference on Economics and Computation (EC' 18)*, 575–592.
- Faliszewski, P., and Talmon, N. 2018. Between proportionality and diversity: Balancing district sizes under the Chamberlin-Courant rule. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, 14–22.
- Faliszewski, P.; Skowron, P.; Slinko, A.; and Talmon, N. 2016a. Committee scoring rules: Axiomatic classification and hierarchy. In *Proceedings of the 27th Joint Conference on Artificial Intelligence (IJCAI '16)*, 250–256.
- Faliszewski, P.; Skowron, P.; Slinko, A. M.; and Talmon, N. 2016b. Multiwinner analogues of the Plurality rule: Axiomatic and algorithmic perspectives. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI '16)*, 482–488.
- Faliszewski, P.; Szufa, S.; and Talmon, N. 2018. Optimization-based voting rule design: The closer to Utopia the better. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, 32–40.
- Feige, U. 1998. A threshold of $\ln n$ for approximating set cover. *Journal of the ACM* 45(4):634–652.
- Fluschnik, T.; Skowron, P.; Triphaus, M.; and Wilker, K. 2017. Fair Knapsack. *arXiv preprint arXiv:1711.04520*.
- Garey, M. R., and Johnson, D. S. 1979. Computers and intractability: A guide to the theory of NP-completeness.
- Gilman, H. R. 2012. Transformative deliberations: Participatory budgeting in the united states. *Journal of Public Deliberation* 8(11).
- Goel, A.; Krishnaswamy, A. K.; Sakshuwong, S.; and Aitamurto, T. 2015. Knapsack voting. *Collective Intelligence*.
- Goel, A.; Krishnaswamy, A. K.; and Sakshuwong, S. 2016. Budget aggregation via Knapsack voting: welfare-maximization and strategy-proofness. *Collective Intelligence* 783–809.
- Gurobi Optimization, L. 2018. Gurobi optimizer reference manual.
- Lackner, M., and Skowron, P. 2017. Consistent approval-based multi-winner rules. *arXiv preprint arXiv:1704.02453*.
- Laslier, J.-F., and Sanver, R., eds. 2010. *Handbook on Approval Voting*. Springer.
- Lu, T., and Boutilier, C. 2011. Budgeted social choice: From consensus to personalized decision making. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI '11)*, volume 11, 280–286.
- Procaccia, A. D.; Rosenschein, J. S.; and Zohar, A. 2008. On the complexity of achieving proportional representation. *Social Choice and Welfare* 30(3):353–362.
- Shapiro, E., and Talmon, N. 2017. A Condorcet-consistent participatory budgeting algorithm. *arXiv preprint arXiv:1709.05839*.
- Skowron, P., and Faliszewski, P. 2017. Chamberlin-Courant rule with approval ballots: Approximating the MaxCover problem with bounded frequencies in FPT time. *Journal of Artificial Intelligence Research* 60:687–716.
- Wainwright, H. 2003. Making a people's budget in Porto Alegre. *NACLA Report on the Americas* 36(5):37–42.