

Maxmin Participatory Budgeting

Gogulapati Sreedurga, Mayank Ratan Bhardwaj and Yadati Narahari

Indian Institute of Science

{gogulapatis, mayankb, narahari}@iisc.ac.in

Abstract

Participatory Budgeting (PB) is a popular voting method by which a limited budget is divided among a set of projects, based on the preferences of voters over the projects. PB is broadly categorised as divisible PB (if the projects are fractionally implementable) and indivisible PB (if the projects are atomic). Egalitarianism, an important objective in PB, has not received much attention in the context of indivisible PB. This paper addresses this gap through a detailed study of a natural egalitarian rule, Maxmin Participatory Budgeting (MPB), in the context of indivisible PB. Our study is in two parts: (1) computational (2) axiomatic. In the first part, we prove that MPB is computationally hard and give pseudo-polynomial time and polynomial-time algorithms when parameterized by certain well-motivated parameters. We propose an algorithm that achieves for MPB, additive approximation guarantees for restricted spaces of instances and empirically show that our algorithm in fact gives exact optimal solutions on real-world PB datasets. We also establish an upper bound on the approximation ratio achievable for MPB by the family of exhaustive strategy-proof PB algorithms. In the second part, we undertake an axiomatic study of the MPB rule by generalizing known axioms in the literature. Our study leads to the proposal of a new axiom, *maximal coverage*, which captures fairness aspects. We prove that MPB satisfies maximal coverage.

1 Introduction

Participatory budgeting (PB) is an appealing voting paradigm used for distributing a limited budget (divisible resource such as money, time, etc.) among proposed alternatives (also called projects). It has been deployed in many applications [Röcke, 2014] due to its practical relevance. The PB setting in which projects are fractionally implementable is called divisible PB whereas the one in which projects are indivisible, each having a certain cost, is called indivisible PB [Aziz and Shah, 2021]. A PB rule aggregates the preferences of voters and proposes a budget division for the projects. Among the

various preference elicitation methods, approval votes (each agent specifies a subset of projects to be funded) have received much attention in the PB literature due to their cognitive simplicity [Benade *et al.*, 2018].

In approval-based PB, one way to define the utility of a voter is to base it on the number of approved projects of the voter that are funded [Rey *et al.*, 2020; Pierczyński *et al.*, 2021]. This notion, however, fails to capture the role played by the size of a project: e.g., suppose the budget is \$100M and a PB rule selects two projects approved by voter 1, costing \$4M and \$6M, and a project approved by voter 2 costing \$90M. The utility notion in question gives higher utility to voter 1 though 90% of the budget is allocated favorably to voter 2. This immediately motivates the well-studied notion of utility - the total amount allocated to the approved projects of the voter [Bogomolnaia *et al.*, 2005; Duddy, 2015; Talmon and Faliszewski, 2019; Aziz *et al.*, 2019; Goel *et al.*, 2019; Freeman *et al.*, 2021]. We work with this utility notion.

Major goals that are usually sought in voting rules include proportionality, individual excellence, and diversity [Faliszewski *et al.*, 2017]. Proportionality assures that every group of voters receives an allocation proportional to its size and is well studied in the PB literature [Aziz *et al.*, 2018b; Aziz and Lee, 2021; Pierczyński *et al.*, 2021; Fairstein *et al.*, 2021]. Individual excellence refers to maximizing the sum of utilities of all agents and is hence achieved by the well studied utilitarian rules [Talmon and Faliszewski, 2019; Goel *et al.*, 2019; Freeman *et al.*, 2021]. Both proportionality and individual excellence are majority-focused rules. Diversity, on the other hand, is a fairness notion that aims to give representation to every voter, irrespective of whether she belongs to majority or minority.

Diversity is well achieved by egalitarian rules. For example, suppose a budget of \$50M is to be allocated to school construction projects in three villages, X, Y, and Z, with populations of 10000, 6000, and 2000, respectively. Village X proposes schools at localities $\{X_1, X_2, X_3, X_4\}$ costing $\{\$10M, \$20M, \$20M, \$60M\}$ respectively, Y proposes schools at $\{Y_1, Y_2, Y_3\}$ costing $\{\$14M, \$14M, \$16M\}$ respectively, and Z proposes a school at Z_1 costing \$6M. Suppose each voter in a village approves all and only the projects proposed by the village. Utilitarian rules select $\{X_1, X_2, X_3\}$. Majority focused rules ignore village Z. However, an egalitarian solution, $\{X_1, X_2, Y_1, Z_1\}$, has a more

diverse representation and thus promotes universal literacy.

Contributions and Outline

Egalitarian rules have been studied in the PB context by Aziz and Stursberg [2014], Aziz et al. [2019], Airiau et al. [2019], and Tang et al. [2020]. All these works focus only on divisible PB. In the context of indivisible PB, egalitarian rules are studied only for a very restricted special case, multi-winner voting, where the budget is k and the projects cost 1 each (e.g., Brams et al. [2007]). Thus, surprisingly, egalitarian rules for indivisible PB have not been studied systematically, barring a case-study by Laruelle [2021] that experimentally evaluates a sub-optimal greedy algorithm for an egalitarian objective. In this paper, we address this important gap by systematically studying a popular egalitarian objective, maxmin, for indivisible PB (MPB).

Our choice of maxmin for a detailed study is motivated by the natural appeal, simplicity, and understandability of the maxmin objective in social choice contexts. The objective does bring with it a couple of limitations. The first limitation is applicable not only to maxmin but to all egalitarian rules: presence of outlier voters affects the outcome adversely. However, in many real-life PB elections, it is seen that the voters can be partitioned into a small number (10 to 20) of buckets based on their preferences (e.g., Poland (2018), Aleksandrow (2019) PB elections [Stolicki et al., 2020]) and outliers are usually not present. The second limitation is that the utility of a voter with inexpensive approval vote is low by default. However, lower cost of the approval vote could be an indication that the voter is not excited by the project proposals, justifying the low utility. Thus, in our view, these two apparent limitations do not take away the importance and relevance of the maxmin objective. In fact, in real-life, maxmin is found to be preferred over the existing proportional PB rules by the voters [Rosenfeld and Talmon, 2021].¹

Our study of MPB for indivisible PB is in two parts: (1) computational (2) axiomatic. In the first part (Section 3), we prove that MPB is computationally hard and give pseudo-polynomial time and polynomial-time algorithms when parameterized by certain well-motivated parameters. We propose an algorithm that achieves, for MPB, additive approximation guarantees for some restricted spaces and empirically show that our algorithm in fact gives exact optimal solutions on real-world PB datasets. We also establish an upper bound on the approximation ratio achievable for MPB by the family of exhaustive strategy-proof PB algorithms. In the second part (Section 4), we undertake an axiomatic study of the MPB rule by generalizing known axioms in the literature. Our study leads us to propose a new axiom, *maximal coverage*, which captures the fairness notion of diversity. We prove that MPB satisfies maximal coverage and achieves diversity.

2 Notations

Let $N = \{1, \dots, n\}$ be the set of voters and $P = \{p_1, \dots, p_m\}$ be the set of projects. Let $c : P \rightarrow \mathbb{N}$ be the cost function and $b \in \mathbb{N}$ be the total budget available where \mathbb{N} is the set of natural numbers. The approval vote profile of all

¹Considers cardinal utilities. Utilitarian and Nash rules, though preferred, are unfair and computationally harder respectively.

voters is represented by \mathcal{A} where each $A_i \in \mathcal{A}$ is the set of projects approved by the voter i . An instance \mathcal{I} of participatory budgeting is $\langle N, P, c, b, \mathcal{A} \rangle$. With a slight abuse of notation, we represent the cost of a set S of projects, $\sum_{p \in S} c(p)$, by $c(S)$. A set S is said to be feasible if $c(S) \leq b$. A feasible approval vote A_i is also called a *knapsack vote*. Let \mathcal{E} represent the set of all feasible subsets of projects.

Given an instance \mathcal{I} , a **participatory budgeting (PB) rule** R outputs a set of feasible subsets of projects, i.e., $R(\mathcal{I}) \subseteq \mathcal{E}$. A **PB algorithm** is a rule such that for all \mathcal{I} , $|R(\mathcal{I})| = 1$.

The utility of a voter i from a set of projects S is defined as the amount of money allocated to the projects approved by i , i.e., $u_i(S) = c(S \cap A_i)$.

Definition 1 (Maxmin Participatory Budgeting (MPB)). *Given a PB instance \mathcal{I} , the MPB rule outputs a set of all subsets $S \subseteq P$ such that*

$$S \in \arg \max_{S \in \mathcal{E}} \min_{i \in N} u_i(S).$$

For ease of presentation, we call the objective optimized by the MPB rule the **MPB objective** or simply **MPB**. Given a PB instance \mathcal{I} and a score s , the problem of determining if there exists $S \in \mathcal{E}$ such that $\min_{i \in N} u_i(S) \geq s$ is called the **decision version of MPB**.

3 Maxmin PB: Computational Results

We first prove the hardness of MPB and present some tractable special cases based on certain well-motivated parameters. We then give an approximation algorithm for MPB and show empirically that, in fact, it gives exact optimal solutions on real-world PB datasets. We conclude by establishing an upper bound on the approximation ratio achieved for MPB by any exhaustive strategy-proof PB algorithm. We start by formulating MPB as the following integer linear program where each variable corresponds to selection of a project.

$$\begin{aligned} & \max q \\ & \text{subject to } q \leq \sum_{p \in A_i} c(p) x_p \quad \forall i \in N \quad (1) \\ & \sum_{p \in P} c(p) x_p \leq b \\ & x_p \in \{0, 1\} \quad \forall p \in P \quad (2) \\ & q \geq 0 \end{aligned}$$

3.1 NP-Hardness of Maxmin PB

Theorem 1. *The decision version of MPB is strongly NP-hard.*

Proof. We reduce the SET COVER problem to our problem. Given a set U of elements, a collection S of subsets of U , and a positive integer k , the SET COVER problem is to find if there exists $F \subseteq S$ such that $\cup_{P \in F} P = U$ and $|F| = k$. It is known to be strongly NP-hard [Garey and Johnson, 1979]. Given an instance $\langle U, S, k \rangle$ of SET COVER, we construct an MPB instance as follows: For each $C \in S$, create a project p_C with unit cost. For each $a \in U$, create a voter i with $A_i = \{p_C : i \in C\}$. Set $b = k$ and $s = 1$. Both these instances are equivalent. Detailed proof is provided in the long version [Sreedurga et al., 2022]. \square

3.2 Tractable Special Cases

Here, we discuss some tractable special cases of MPB.

Constant Number of Projects

We look at this parameter, since, in many scenarios, there can be an upper bound on the number of projects that can be funded due to logistical reasons. The following proposition follows from the brute-force algorithm and also the fact that solving the integer linear program is in FPT when parameterized by the number of variables [Lenstra Jr, 1983].

Proposition 1. *MPB can be solved in polynomial time when the number of projects is constant.*

Constant Number of Distinct Votes

In many real-world scenarios, though the number of voters is large, the number of distinct votes is small. That is, the set of voters can be partitioned into a small number of equivalence classes such that all voters in an equivalence class approve the same set of projects. For example, the 2018 PB elections held in Powazki (Warsaw, Poland) had 3482 voters, out of which there were only 16 distinct approval votes. In fact, it has been found that in many real-life PB datasets [Stolicki *et al.*, 2020], the number of distinct votes is less than 20% of the number of voters. Information on some such datasets is included at <https://github.com/Participatory-Budgeting/maxmin-2022>.

Theorem 2. *The decision version of MPB is weakly NP-hard when the number of distinct votes is constant.*

Proof. We reduce the SUBSET SUM problem to our problem. Given an integer Z and a set of integers $X = \{x_1, x_2, \dots, x_n\}$, the SUBSET SUM problem is to determine if there exists a subset $X' \subseteq X$ such that $\sum_{x \in X'} x = Z$. It is known to be weakly NP-hard [Garey and Johnson, 1979]. Given an instance $\langle Z, X \rangle$ of SUBSET SUM, we construct an MPB instance as follows: For each $x_i \in X$, create a project p_i with cost x_i . Create a single voter who approves all the projects. Set $b = s = Z$. Both these instances are equivalent. Detailed proof is provided in [Sreedurga *et al.*, 2022]. \square

From Theorem 1, we know that it is impossible to have a pseudo-polynomial time algorithm if the number of distinct votes is large. We provide a pseudo-poly time algorithm to solve MPB when the number of distinct votes is small.

Theorem 3. *MPB can be solved in pseudo-polynomial time when the number of distinct votes is constant.*

Proof. Let the number of distinct votes be \hat{n} . Let $A_1, \dots, A_{\hat{n}}$ represent these distinct votes. We propose a dynamic programming algorithm. Construct a $\hat{n} + 2$ dimensional binary matrix Q such that $Q(i, j, u_1, \dots, u_{\hat{n}})$ takes a value 1 if and only if there exists a subset $S \subseteq \{p_1, \dots, p_i\}$ such that $c(S) \leq j$ and $c(S \cap A_t) = u_t$ for all $t \in [\hat{n}]$. Here, i takes values from 1 to m , j takes values from 0 to b , and the remaining entries take values from 0 to j . Let the collection of all such $\hat{n} + 2$ sized tuples be X . We fill the first row of the matrix as follows:

$$Q(1, j, u_1, \dots, u_{\hat{n}}) = \begin{cases} 1 & \text{if } j \geq c(p_1), u_t \in \{0, c(p_1)\}, \\ & u_t = 0 \iff p_1 \notin A_t \forall t \in [\hat{n}] \\ 0 & \text{otherwise} \end{cases}$$

Now, we fill the matrix recursively as follows:

$$Q(i, j, u_1, \dots, u_{\hat{n}}) = \max\{Q(i-1, j, u_1, \dots, u_{\hat{n}}), \\ Q(i-1, j-c(p_i), v_1, \dots, v_{\hat{n}})\}$$

where for all $t \in [\hat{n}]$, v_t is u_t if $p_i \notin A_t$ and $u_t - c(p_i)$ otherwise. We know that there are $|X|$ entries in our matrix. The solution of MPB is as follows:

$$\max_{(m, b, u_1, \dots, u_{\hat{n}}) \in X} \left(Q(m, b, u_1, \dots, u_{\hat{n}}) \cdot \min_{t \in [\hat{n}]} u_t \right)$$

There are at most $m(b+1)^{\hat{n}+1}$ tuples in X and computing each entry of our matrix takes constant time. The computation of MPB solution from the matrix takes $O(\hat{n}b^{\hat{n}} \log b)$ time. The total running time is $O(m\hat{n}b^{\hat{n}} \log b)$, which is pseudo-polynomial if \hat{n} is constant. Correctness follows from the definition of Q . \square

Constant Scalable Limit

We introduce a new natural parameter, *scalable limit*, that is reasonably small in several real-world PB elections.

Definition 2 (Scalable Limit). *Given a PB instance $\langle N, P, c, b, A \rangle$, we refer to the ratio $\frac{\max\{c(p_1), \dots, c(p_m)\}}{\text{GCD}(c(p_1), \dots, c(p_m), b)}$ as *scalable limit*, denoted by δ .*

Often in many real-world settings, the costs of the projects and the budget are expressed as multiples of some large value. For example, suppose a budget of 10 billion dollars is to be distributed among a set of projects costing hundreds of millions each. That is, the cost of each project is a multiple of \$100M. This PB instance could be scaled down by dividing the costs and budget with \$100M to derive a new instance with a budget of 100. If the cost of the most expensive project originally was \$900M, it would now cost 9 in the scaled down instance. This number 9 is what we call the scalable limit. In other words, the scalable limit of an instance is the cost of the most expensive project after scaling down the costs and budget to values as low as possible. This parameter takes quite low values in many real-world PB election datasets, e.g., PB elections in Boston, New York District 8, Seattle District 1 (2019) etc. (see <https://pbstanford.org/>).

From Theorem 2, we know that MPB is not poly-time solvable even if the number of distinct votes is small. We now prove that, if the scalable limit is also small in conjunction with the number of distinct votes being small, then MPB is poly-time solvable.

Theorem 4. *MPB can be solved in polynomial time when the number of distinct votes and the scalable limit are constant.*

Proof. Let \hat{n} be the number of distinct votes and let $A_1, \dots, A_{\hat{n}}$ represent these distinct votes. Divide the costs and budget of the instance by $\text{GCD}(c(p_1), \dots, c(p_m), b)$ to obtain a new instance \mathcal{I}' with costs c' and budget b' . Clearly, \mathcal{I}' has the same optimal MPB solution as that of \mathcal{I} .

It is known that the problem of solving an ILP is in FPT when parameterized by the number of constraints and the highest value in coefficient matrix [Ganian and Ordyniak, 2019]. We modify the integer linear program for MPB by replacing N with $[\hat{n}]$ and c with c' . Now, the highest value in the coefficient matrix is $\max\{c'(p)\}$, i.e., δ , and the number of constraints is $\hat{n} + 1$. Clearly, the modified ILP is equivalent to the initial ILP and the theorem follows. \square

3.3 A High Performing PB Algorithm

In this section, we first introduce a family of PB algorithms called as Ordered-Fill algorithms. We then propose an algorithm, called ORDERED-RELAX, from this family, that is based on LP-rounding. We prove that it achieves approximation guarantees for MPB for some restricted spaces of instances. We show empirically that it provides exact optimal solutions for MPB on real-world PB datasets.

Definition 3 (Ordered-Fill Algorithms). *Given a participatory budgeting instance \mathcal{I} , an ordered-fill algorithm w.r.t. a complete order \succ over P selects the projects in the decreasing order of their ranks in \succ until the next ranked project does not fit within the budget.²*

Example 1. *Consider an instance where $P = \{p_1, p_2, p_3\}$, $b = 4$, $c(p_1) = c(p_3) = 2$, and $c(p_2) = 3$. An ordered-fill algorithm w.r.t. $p_1 \succ p_2 \succ p_3$ outputs $\{p_1\}$.³*

We consider the LP-relaxation of integer linear program for MPB objective by relaxing Eq. (2) to $0 \leq x_p \leq 1$. Consider the following LP-rounding algorithm: Solve the relaxed LP to get (q^*, x^*) . Let $S = \phi$ be the initial outcome. Add the project with the highest value of $c(p)x_p^*$ to S , followed by the one with the second highest value, and so on, till the next project does not fit. Call this algorithm ORDERED-RELAX.

Lemma 1. *For any PB instance \mathcal{I} , ORDERED-RELAX outputs a set $S \in \mathcal{E}$ such that $ALG \geq OPT - \frac{|A_j \setminus S|}{|S \setminus A_j|} \cdot (b - OPT)$ where $j = \arg \min_{i \in N} c(A_i \cap S)$, $ALG = c(A_j \cap S)$, and OPT is the minimum utility in the optimal solution for MPB.*

Proof. Let S be the outcome of ORDERED-RELAX, $j = \arg \min_{i \in N} c(A_i \cap S)$, and $\eta = \frac{|A_j \setminus S|}{|S \setminus A_j|}$.

Let the solution of the relaxed LP be (q^*, x^*) . Let $Y_i = S \cap A_i$ for each voter i . Since $S \setminus Y_i \subseteq P \setminus A_i$, for each i ,

$$\begin{aligned} \sum_{p \in S \setminus Y_i} c(p)x_p^* &\leq \sum_{p \in P \setminus A_i} c(p)x_p^* \\ \sum_{p \in S \setminus Y_i} c(p)x_p^* &\leq b - \sum_{p \in A_i} c(p)x_p^* \end{aligned} \quad (3)$$

From the design of this algorithm, every project p not in S has $c(p)x_p^*$ not more than that of projects in S .

$$\begin{aligned} \forall p \in S \setminus Y_i \quad c(p)x_p^* &\geq \max_{p' \in A_i \setminus Y_i} c(p')x_{p'}^* \\ &\geq \frac{\sum_{p' \in A_i \setminus Y_i} c(p')x_{p'}^*}{|A_i| - |Y_i|} \end{aligned}$$

$$\sum_{p \in S \setminus Y_i} c(p)x_p^* \geq \frac{|S| - |Y_i|}{|A_i| - |Y_i|} \sum_{p \in A_i \setminus Y_i} c(p)x_p^*$$

Since $x_p^* \leq 1 \forall p$ and $u_i(S) = \sum_{p \in Y_i} c(p)$, $\sum_{p \in Y_i} c(p)x_p^* \leq u_i(S)$.

$$\sum_{p \in S \setminus Y_i} c(p)x_p^* \geq \frac{|S| - |Y_i|}{|A_i| - |Y_i|} \left(\sum_{p \in A_i} c(p)x_p^* - u_i(S) \right) \quad (4)$$

²Greedy rules [Talmon and Faliszewski, 2019] are ordered-fill algorithms where \succ is based on utility from each affordable project.

³Note that the outcome is not maximal since $\{p_1, p_3\} \in \mathcal{E}$.

From Eq. (3) and Eq. (4):

$$\begin{aligned} \frac{|S| - |Y_i|}{|A_i| - |Y_i|} \left(\sum_{p \in A_i} c(p)x_p^* - u_i(S) \right) &\leq b - \sum_{p \in A_i} c(p)x_p^* \\ \sum_{p \in A_i} c(p)x_p^* &\leq \frac{b + \frac{|S| - |Y_i|}{|A_i| - |Y_i|} u_i(S)}{\frac{|S| - |Y_i|}{|A_i| - |Y_i|} + 1} \\ \sum_{p \in A_i} c(p)x_p^* &\leq \frac{\eta b + ALG}{1 + \eta} \end{aligned} \quad (5)$$

Since the optimal solution also belongs to the feasible region of the relaxed LP, we know that $OPT \leq q^*$. From Eq. (1), we have $q^* \leq \sum_{p \in A_j} c(p)x_p^*$. Combining these observations with Eq. (5), we get,

$$\begin{aligned} OPT &\leq \frac{\eta b + ALG}{1 + \eta} \\ ALG &\geq OPT - \eta(b - OPT) \end{aligned}$$

This proves the result. \square

Given an instance \mathcal{I} , let l_o and h_o denote respectively the minimum and maximum cardinalities of the outputs produced by all ordered-fill algorithms on \mathcal{I} . They are poly-time computable [Sreedurga et al., 2022]. Let l_A and h_A respectively denote the lowest and highest cardinalities of the approval votes, i.e., $l_A = \min_{i \in N} |A_i|$ and $h_A = \max_{i \in N} |A_i|$.

Lemma 1 can be used to establish theoretical approximation guarantees for some restricted spaces of instances, e.g., let us look at a family of instances that satisfy a property that we call *High Cardinality Budget Property* (HCBP) which intuitively is a cardinal extension of all votes being knapsack votes. That is, it requires that the budget is high enough to fund more projects than the number of projects in any single approval vote. In other words, it requires that $l_o > h_A$. Note that this property is natural in scenarios where budget is high enough to accommodate all projects approved by one voter, e.g., when the budget with the federal government is high enough to fund all proposed projects of one county or when each member in the audience approves less number of talks than the total capacity of the workshop, etc. For the instances satisfying HCBP, our algorithm ensures that $ALG \geq OPT - h_A(b - OPT)$ ⁴.

As another example, we note that the algorithm yields an optimal solution for the instances which guarantee $A_j \subseteq S$. One such family of instances will be the one that includes instances satisfying the following condition for all $S \in \mathcal{E}$: if there exists some $p \in P$ such that $c(S) + c(p) > b$, then all approved projects of the worse-off voter from S are included in S . Note that this is a sufficient but not necessary condition for the algorithm to yield an optimal solution. We acknowledge that, for some instances, the approximation ratio given by the lemma could be trivial theoretically, but our empirical analysis compensates for this.

⁴Further, if we define the disutility of a voter i from a set S to be $b - u_i(S)$ and minimize the maximum disutility, our algorithm achieves $\left(2 - \frac{1}{h_o}\right)$ approximation for HCBP instances. Refer [Sreedurga et al., 2022].

Empirical Analysis

Though the theoretical guarantee of our algorithm is limited in terms of approximation ratios and the spaces of instances it covers, empirically it is found to exhibit remarkable performance to give *exact optimal solutions* on all the PB datasets available in [Stolicki *et al.*, 2020]. The datasets and corresponding results are included at https://github.com/Participatory-Budgeting/maxmin_2022. This indicates that, although theoretically the worst cases do not allow the approximation ratio to get better, such worst cases are seldom encountered in real-life, resulting in the excellent performance of our algorithm.

3.4 Bound on the Performance of Exhaustive and Strategy-Proof PB Algorithms

It would be useful to know how the requirement of strategy-proofness will degrade the performance of PB algorithms with respect to MPB objective. We establish an upper bound on the approximation ratios of exhaustive strategy-proof PB algorithms by invoking an argument used in the mechanism design literature [Procaccia and Tennenholtz, 2009; Caragiannis *et al.*, 2010]. We construct a specific profile whose outcome can be derived using strategy-proofness and establish an upper bound on its approximation ratio.

Definition 4 (Strategy-proofness). *An algorithm \mathcal{G} is said to be strategy-proof, if and only if, for any instance \mathcal{I} ,*

$$\forall i \quad \forall S \subseteq P \quad c(A_i \cap \mathcal{G}(A_i, \mathcal{A}_{-i})) \geq c(A_i \cap \mathcal{G}(S, \mathcal{A}_{-i}))$$

where \mathcal{A}_{-i} denotes the vote profile of all voters except i and $\mathcal{G}(\mathcal{A})$ denotes the outcome of \mathcal{G} at profile \mathcal{A} .

That is, if all voters except i continue to approve the same sets, i should not obtain a better outcome by approving any set other than A_i . An algorithm is exhaustive if it is impossible to fund an unselected project with the remaining budget.

Definition 5 (Exhaustiveness). *An algorithm \mathcal{G} , is said to be exhaustive if and only if for every instance \mathcal{I} and every $p \notin \mathcal{G}(\mathcal{I})$, it holds that $c(p) + c(\mathcal{G}(\mathcal{I})) > b$.*

Theorem 5. *No exhaustive strategy-proof algorithm can achieve an approximation guarantee better than $\frac{2}{3}$ for MPB even for instances with only knapsack votes and unit costs.*

Proof. Consider an exhaustive strategy-proof algorithm \mathcal{G} . Consider an instance \mathcal{I}_1 with a budget b , $2b$ projects each costing 1, two voters, and a vote profile $\mathcal{A}_1 = \{A_1, A_2\}$ such that $A_1 \cap A_2 = \emptyset$ and $c(A_1) = c(A_2) = b$. Let $\mathcal{G}(\mathcal{I}_1) = S$. Without loss of generality, assume that $c(A_1 \cap S) \leq c(A_2 \cap S)$. Modify the vote profile of \mathcal{I}_1 to get \mathcal{I}_2 as follows: $\mathcal{A}_2 = \{A_1, S\}$. For \mathcal{I}_2 , \mathcal{G} should output S . Else, voter 2 would have the incentive to misreport A_j instead of the true vote S and get the preferred outcome S . Hence, $\mathcal{G}(\mathcal{I}_2) = S$. Therefore, the minimum utility of the algorithm for \mathcal{I}_2 , $ALG(\mathcal{I}_2)$, is $c(A_1 \cap S)$. Let $OPT(\mathcal{I}_2)$ be the optimal minimum utility for \mathcal{I}_2 . That is,

$$OPT(\mathcal{I}_2) \geq \min(c(A_1 \cap X), c(S \cap X)) \quad \forall X \in \mathcal{E} \quad (6)$$

Since \mathcal{G} is exhaustive, $c(S) = b$. Consider a set $X \subseteq P$ such that $(A_1 \cap S) \subseteq X$, $|X \cap (A_1 \setminus S)| = \frac{b - c(A_1 \cap S)}{2}$, and

$$|X \cap (S \setminus A_1)| = \frac{b - c(A_1 \cap S)}{2}. \text{ Clearly, } X \in \mathcal{E}.$$

$$\begin{aligned} c(A_1 \cap X) &= c(S \cap X) = c(A_1 \cap S) + \frac{b - c(A_1 \cap S)}{2} \\ &= \frac{b + c(A_1 \cap S)}{2} \end{aligned} \quad (7)$$

Since $c(A_1 \cap S) \leq c(A_2 \cap S)$ and $c(S) = b$, $c(A_1 \cap S) \leq \frac{b}{2}$. Substituting $b \geq 2c(A_1 \cap S)$ in Eq. (7), we get,

$$OPT(\mathcal{I}_2) \geq \frac{3c(A_1 \cap S)}{2} \quad (\text{From Eq. (6)})$$

$$\frac{ALG(\mathcal{I}_2)}{OPT(\mathcal{I}_2)} \leq \frac{2}{3}$$

Hence, the approximation guarantee of the algorithm is at most $\frac{2}{3}$. Clearly in both the instances \mathcal{I}_1 and \mathcal{I}_2 , all the votes are knapsack votes and the theorem follows. \square

4 Maxmin PB: An Axiomatic Analysis

An axiomatic study of a voting rule provides valuable insights into the characteristics of the voting rule. We now undertake an axiomatic study of the MPB rule by exploring several axiomatic properties available in the literature. Given an instance \mathcal{I} and a PB rule R , we say a project p wins if there exists $S \in \mathcal{R}(\mathcal{I})$ such that $p \in S$. Let $\mathcal{W}_R(\mathcal{I})$ represent the set of all projects which win, i.e., $\mathcal{W}_R(\mathcal{I}) = \{p \in P : \exists S \in \mathcal{R}(\mathcal{I}) p \in S\}$. Throughout this section, we represent the MPB rule by M .

First, we examine axioms in the PB literature for rules that output a single feasible subset of projects [Talmon and Faliszewski, 2019]. We then extend the axioms to rules (like MPB) that output multiple feasible subsets. We start by defining all the axioms. The first axiom requires that if a winning project is replaced by a set of projects whose total cost is the same, then at least one project of this set should continue to win.

Definition 6 (Splitting Monotonicity). *A PB rule R is said to satisfy splitting monotonicity iff, for any instance \mathcal{I} , $\mathcal{W}_R(\mathcal{I}') \cap P' \neq \emptyset$ whenever \mathcal{I}' is obtained from \mathcal{I} by splitting a project $p \in \mathcal{W}_R(\mathcal{I})$ into a set of projects P' with $c(P') = c(p)$ and changing every A_i having p to $(A_i \setminus \{p\}) \cup P'$.*

The next axiom requires that, if we replace a set of winning projects that are approved by exactly the same set of voters by a single project having the same total cost, then the single project wins.

Definition 7 (Merging Monotonicity). *A PB rule R is said to satisfy merging monotonicity iff, for any instance \mathcal{I} , $p \in \mathcal{W}_R(\mathcal{I}')$ whenever \mathcal{I}' is obtained from \mathcal{I} as follows: for any $S \in \mathcal{R}(\mathcal{I})$ and $P' \subseteq S$ such that $A_i \cap P' \in \{P', \emptyset\}$ for every voter i , merge the projects in P' into a single project p with cost $c(P')$ and change every A_i with P' to $(A_i \setminus P') \cup \{p\}$.*

The next axiom requires that no winning project should be dropped if it becomes less expensive.

Definition 8 (Discount Monotonicity). *A PB rule R is said to satisfy discount monotonicity iff, for any instance \mathcal{I} , $p \in \mathcal{W}_R(\mathcal{I}')$ whenever \mathcal{I}' is obtained from \mathcal{I} by reducing the cost of $p \in \mathcal{W}_R(\mathcal{I})$ to $c(p) - 1$.*

The following axiom requires that no winning project should be dropped if the budget increases.

Definition 9 (Limit Monotonicity). *A PB rule R is said to satisfy limit monotonicity iff, for any instance \mathcal{I} such that no project in P costs $b + 1$, $\mathcal{W}_R(\mathcal{I}) \subseteq \mathcal{W}_R(\mathcal{I}')$ whenever \mathcal{I}' is obtained from \mathcal{I} by increasing the budget to $b + 1$.*

The next axiom requires that all sets in $R(\mathcal{I})$ are maximal [Aziz and Shah, 2021].

Definition 10 (Strong Exhaustiveness). *A PB rule R is said to satisfy strong exhaustiveness iff, for any instance \mathcal{I} ,*

$$\forall S \in R(\mathcal{I}) \quad \forall p \in P \setminus S \quad c(S) + c(p) > b$$

Next, we introduce a new axiom, weak exhaustiveness, which requires that any winning non-maximal set can be made maximal without compromising on the win.

Definition 11 (Weak Exhaustiveness). *A PB rule R is said to satisfy weak exhaustiveness iff, for any instance \mathcal{I} ,*

$$\forall S \in R(\mathcal{I}), \forall p \in P \setminus S, c(S) + c(p) \leq b \implies S \cup \{p\} \in R(\mathcal{I})$$

Theorem 6. *The MPB rule satisfies: (a) splitting monotonicity (b) merging monotonicity (c) weak exhaustiveness.*

Theorem 7. *The MPB rule does not satisfy: (a) discount monotonicity (b) limit monotonicity (c) strong exhaustiveness.*

Proofs of the above theorems are given in the long version [Sreedurga et al., 2022]. We now examine two important axioms from multi-winner voting literature. The first axiom is an analogue of unanimity, which requires that a project approved by all voters needs to win [Faliszewski et al., 2017].

Definition 12 (Narrow-top Criterion). *A PB rule R is said to satisfy narrow-top criterion iff, for any instance \mathcal{I} , $p \in \mathcal{W}_R(\mathcal{I})$ whenever $p \in A_i$ for every $i \in N$.*

Proposition 2. *The MPB rule does not satisfy narrow-top criterion.*

A proof using a simple counter example is given in [Sreedurga et al., 2022]. The next axiom tries to capture diversity in the voting rules [Aziz et al., 2018a; Brandt et al., 2016].

Definition 13 (Clone Independence). *A PB rule R is said to satisfy clone independence iff, for any instance \mathcal{I} , $R(\mathcal{I})$ does not change when a group of voters, all having exactly the same approval vote A_v , is replaced by a single voter with approval vote A_v .*

Proposition 3. *The MPB rule satisfies clone independence.*

The above proposition follows from the fact that redundant votes do not affect the value of maxmin objective. Clearly, this axiom represents diversity, but it is rather narrow. Faliszewski et al. [2017] identified that the fairness notion of diversity does not have a clear axiomatic representation in the literature. We address this gap by introducing a new axiom, called *Maximal Coverage*, to capture diversity in PB as well as in multi-winner voting settings.

Let us define *covered voters* as the voters with at least one approved project funded and a *redundant project* as a project whose removal from an outcome does not change the set of covered voters. To achieve diversity, we need to cover as

many voters as possible. Our new axiom ensures that a redundant project is funded only when no more voters can be covered by doing otherwise. For example, while allocating time for plenary talks, the organizers might want to explore as many novel ideas as possible, without eliminating any key area. They might allocate a time slot to a proposed plenary talk on a not-so-popular topic (hoping to prop up a new area), by dropping 1 out of 4 proposed talks in a popular area. This *covers* or reaches out to the audience from the new area and broadens the reach of the conference.

Definition 14 (Maximal Coverage). *A PB rule R is said to satisfy maximal coverage iff, for any instance \mathcal{I} , $S \in R(\mathcal{I})$, $p \in S$ such that $\{j_1 : p \in A_{j_1}\} \subseteq \{j_2 : (S \setminus \{p\}) \cap A_{j_2} \neq \emptyset\}$, and $i \in N$,*

$$\mathcal{W}_R(\mathcal{I}) \cap A_i = \emptyset \implies c(a) > b - c(S \setminus \{p\}) \quad \forall a \in A_i \quad (8)$$

We consider an example to understand this better. Suppose a budget of \$2.25B is to be allocated to projects in two counties, X and Y, with populations of 10000 and 6000, respectively. County X proposes projects $\{X_1, X_2, X_3\}$ costing $\{\$0.5B, \$1B, \$1B\}$, and County Y proposes projects $\{Y_1, Y_2, Y_3\}$ costing $\{\$0.7B, \$0.7B, \$0.8B\}$ respectively. Assume each voter approves all and only the projects of her county. A utilitarian rule selects the set $\{\{X_2, X_3\}\}$. Let p be X_2 . If $i = 2$ and $a = Y_1$, then the first part of Eq. (8) satisfies but $c(Y_1) \leq b - c(X_3)$. Hence it does not satisfy maximal coverage. If we apply the MPB rule, we get $\mathcal{W}_M(\mathcal{I}) = \{X_1, X_2, X_3, Y_1, Y_2\}$ and $M(\mathcal{I}) = \{\{X_1, X_2, Y_1\}, \{X_1, X_2, Y_2\}, \{X_1, X_3, Y_1\}, \{X_1, X_3, Y_2\}\}$. It is notable that Eq. (8) holds.

Proposition 4. *The MPB rule satisfies maximal coverage.*

Proof. Let S, p, i, a be as stated in Definition 14. Assume $\exists S$ such that the first part of Eq. (8) holds. Since $\mathcal{W}_M(\mathcal{I}) \cap A_i = \emptyset$, $u_i(S) = 0$. Hence the minimum utility from any set in $M(\mathcal{I})$ is 0. But, since $\mathcal{W}_M(\mathcal{I}) \cap A_i = \emptyset$ and $a \in A_i$, $\{a\} \notin M(\mathcal{I})$. This is possible only if $c(a) > b$ and, thus, the second part of Eq. (8) holds. \square

5 Summary and Future Directions

In this paper, we have motivated the MPB rule for indivisible PB and undertaken a computational as well as an axiomatic study of the same. On the computational side, we proved MPB is strongly NP-hard. We then proved it is weakly NP-hard when the number of distinct votes is small and proposed a pseudo-polynomial time algorithm. We introduced a novel parameter, scalable limit, and proved that if it too is small, MPB can be solved in polynomial-time. We proposed an LP-rounding algorithm that gives additive approximation guarantees for restricted spaces of instances and showed empirically that it gives optimal outcomes on real-world PB datasets. We established an upper bound on the approximation ratio achievable by the family of exhaustive and strategy-proof algorithms. On the axiomatic side, we undertook a detailed analysis of MPB and introduced a new axiom, maximal coverage, to capture diversity in PB and multi-winner voting.

Giving an axiomatic characterization of MPB, studying other objectives (for example leximin), establishing the connections between maximal coverage and other axioms in the literature are a few of several interesting directions to pursue.

Acknowledgements

The authors thank Dr. Neeldhara Misra for suggesting the number of distinct votes as a parameter, and the anonymous reviewers for some helpful feedback. Gogulapati Sreedurga gratefully acknowledges the Prime Minister Research Fellowship funded by Government of India. Mayank Ratan Bhardwaj gratefully acknowledges the MHRD funding by Government of India.

References

- [Airiau *et al.*, 2019] Stéphane Airiau, Haris Aziz, Ioannis Caragiannis, Justin Kruger, Jérôme Lang, and Dominik Peters. Portioning using ordinal preferences: Fairness and efficiency. In *IJCAI 2019*, pages 11–17, 2019.
- [Aziz and Lee, 2021] Haris Aziz and Barton E Lee. Proportionally representative participatory budgeting with ordinal preferences. In *AAAI 2021*, pages 5110–5118, 2021.
- [Aziz and Shah, 2021] Haris Aziz and Nisarg Shah. Participatory budgeting: Models and approaches. In *Pathways Between Social Science and Computational Social Science*, pages 215–236. Springer, 2021.
- [Aziz and Stursberg, 2014] Haris Aziz and Paul Stursberg. A generalization of probabilistic serial to randomized social choice. In *AAAI 2014*, pages 559–565, 2014.
- [Aziz *et al.*, 2018a] Haris Aziz, Piotr Faliszewski, Bernard Grofman, Arkadii Slinko, and Nimrod Talmon. Egalitarian committee scoring rules. In *IJCAI 2018*, pages 56–62, 2018.
- [Aziz *et al.*, 2018b] Haris Aziz, Barton E Lee, and Nimrod Talmon. Proportionally representative participatory budgeting: Axioms and algorithms. In *AAMAS 2018*, pages 23–31, 2018.
- [Aziz *et al.*, 2019] Haris Aziz, Anna Bogomolnaia, and Hervé Moulin. Fair mixing: The case of dichotomous preferences. In *Proceedings of the 2019 ACM EC*, pages 753–781, 2019.
- [Benade *et al.*, 2018] Gerdus Benade, Nevo Itzhak, Nisarg Shah, and Ariel D Procaccia. Efficiency and usability of participatory budgeting methods. <https://www.cs.toronto.edu/~nisarg/papers/pb-usability.pdf>, 2018. Accessed: 2022-01-10.
- [Bogomolnaia *et al.*, 2005] Anna Bogomolnaia, Hervé Moulin, and Richard Stong. Collective choice under dichotomous preferences. *Journal of Economic Theory*, 122(2):165–184, 2005.
- [Brams *et al.*, 2007] Steven J Brams, D Marc Kilgour, and M Remzi Sanver. A minimax procedure for electing committees. *Public Choice*, 132(3-4):401–420, 2007.
- [Brandt *et al.*, 2016] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D Procaccia. *Handbook of computational social choice*. Cambridge University Press, 2016.
- [Caragiannis *et al.*, 2010] Ioannis Caragiannis, Dimitris Kalaitzis, and Evangelos Markakis. Approximation algorithms and mechanism design for minimax approval voting. In *AAAI 2010*, pages 737–742, 2010.
- [Duddy, 2015] Conal Duddy. Fair sharing under dichotomous preferences. *Mathematical Social Sciences*, 73:1–5, 2015.
- [Fairstein *et al.*, 2021] Roy Fairstein, Reshef Meir, and Kobi Gal. Proportional participatory budgeting with substitute projects. *arXiv preprint arXiv:2106.05360*, 2021.
- [Faliszewski *et al.*, 2017] Piotr Faliszewski, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. Multiwinner voting: A new challenge for social choice theory. *Trends in computational social choice*, 74:27–47, 2017.
- [Freeman *et al.*, 2021] Rupert Freeman, David M Pennock, Dominik Peters, and Jennifer Wortman Vaughan. Truthful aggregation of budget proposals. *Journal of Economic Theory*, 193:105234, 2021.
- [Ganian and Ordyniak, 2019] Robert Ganian and Sebastian Ordyniak. Solving integer linear programs by exploiting variable-constraint interactions: A survey. *Algorithms*, 12(12):248, 2019.
- [Garey and Johnson, 1979] Michael R Garey and David S Johnson. *Computers and intractability*, volume 174. San Francisco: freeman, 1979.
- [Goel *et al.*, 2019] Ashish Goel, Anilesh K Krishnaswamy, Sukolsak Sakshuwong, and Tanja Aitamurto. Knapsack voting for participatory budgeting. *ACM Transactions on Economics and Computation (TEAC)*, 7(2):1–27, 2019.
- [Laruelle, 2021] Annick Laruelle. Voting to select projects in participatory budgeting. *European Journal of Operational Research*, 288(2):598–604, 2021.
- [Lenstra Jr, 1983] Hendrik W Lenstra Jr. Integer programming with a fixed number of variables. *Mathematics of operations research*, 8(4):538–548, 1983.
- [Pierczyński *et al.*, 2021] Grzegorz Pierczyński, Piotr Skowron, and Dominik Peters. Proportional participatory budgeting with additive utilities. *Advances in Neural Information Processing Systems*, 34, 2021.
- [Procaccia and Tennenholtz, 2009] Ariel D Procaccia and Moshe Tennenholtz. Approximate mechanism design without money. In *Proceedings of the 10th ACM conference on electronic commerce*, pages 177–186, 2009.
- [Rey *et al.*, 2020] Simon Rey, Ulle Endriss, and Ronald de Haan. Designing participatory budgeting mechanisms grounded in judgment aggregation. In *KR 2020*, pages 692–702, 2020.
- [Röcke, 2014] Anja Röcke. *Framing citizen participation: Participatory budgeting in France, Germany and the United Kingdom*. Springer, 2014.
- [Rosenfeld and Talmon, 2021] Ariel Rosenfeld and Nimrod Talmon. What should we optimize in participatory budgeting?: An experimental study. *arXiv preprint arXiv:2111.07308*, 2021.
- [Sreedurga *et al.*, 2022] Gogulapati Sreedurga, Mayank Ratan Bhardwaj, and Y. Narahari. Maxmin participatory budgeting. *arXiv preprint arXiv:2204.13923*, 2022.
- [Stolicki *et al.*, 2020] Dariusz Stolicki, Stanisław Szufa, and Nimrod Talmon. Pabulib: A participatory budgeting library. *arXiv preprint arXiv:2012.06539*, 2020.
- [Talmon and Faliszewski, 2019] Nimrod Talmon and Piotr Faliszewski. A framework for approval-based budgeting methods. In *AAAI 2019*, pages 2181–2188, 2019.
- [Tang *et al.*, 2020] Zhongzheng Tang, Chenhao Wang, and Mengqi Zhang. Price of fairness in budget division for egalitarian social welfare. In *ICCOA 2020*, pages 594–607. Springer, 2020.