

Fairness in Long-Term Participatory Budgeting

Martin Lackner¹, Jan Maly¹ and Simon Rey^{2*}

¹DBAI, TU Wien, Vienna

²Institute for Logic, Language and Computation, University of Amsterdam, Amsterdam

{lackner, jmaly}@dbai.tuwien.ac.at, s.j.rey@uva.nl

Abstract

Participatory Budgeting (PB) processes are usually designed to span several years, with referenda for new budget allocations taking place regularly. This paper presents a first formal framework for long-term PB, based on a sequence of budgeting problems as main input. We introduce a theory of fairness for this setting, focusing on three main concepts that apply to types (groups) of voters: (i) achieving equal welfare for all types, (ii) minimizing inequality of welfare (as measured by the Gini coefficient), and (iii) achieving equal welfare in the long run. For different notions of welfare, we investigate under which conditions these criteria can be satisfied, and analyze the computational complexity of verifying whether they hold.

1 Introduction

Participatory Budgeting (PB) is a democratic tool in which citizens are asked their opinion on how to spend a public budget [Cabannes, 2004; Shah, 2007]. This process is now applied, which was invented in Brazil, is now used in many cities all around the world [Dias *et al.*, 2019]. The way it is precisely organized differs from place to place but generally the same two-stage structure is adopted [Shah, 2007]: first citizens propose projects and then they vote on these projects. Based on these votes, a set of projects is chosen that can be implemented with the available budget. Importantly, PB is usually planned to run for several years. For instance, a participatory budgeting process in Paris spanned 6 years (from 2014 to 2020) [City of Paris, 2020], and New York runs an ongoing program since 2011 [New York City Council, 2020]. The general idea of PB is to establish it as a regular, ongoing process for sustained citizen participation.

Even though PB has received substantial attention in recent years through the lens of (computational) social choice [Aziz and Shah, 2020], its formalizations generally consider PB as a one-shot process. This assumption significantly limits the scope of an analysis. In particular, it disregards the possibility of achieving fair outcomes *over time*, although a fair solution may be impossible to obtain in individual PB instances. The

main purpose of our work is to close this gap: we introduce *perpetual participatory budgeting*, a formal framework that encompasses key characteristics of long-term PB.

The long-term perspective of perpetual PB leads to conceptual challenges but brings notable advantages. To highlight the potential of this approach, we introduce and study notions of fairness in this setting and analyze to which extent strong fairness guarantees can be achieved in long-term processes. We are mainly concerned with fairness towards *types* of voters. A type is a pre-defined subset of voters, for example all voters in a certain district or socio-demographic groups (e.g., age, education, income). Furthermore, to be able to speak about fairness, we have to specify how we measure the welfare of types. We consider three main forms of welfare: The first is *satisfaction*, which intuitively corresponds to the agreement between a voter’s ballot and the chosen projects, weighted by cost. The satisfaction of a type is the average satisfaction of its voters. The second welfare notion is *relative satisfaction*, which is similar to satisfaction but measures the satisfaction relative to the voter’s maximally achievable satisfaction. The third is the *share* of a type, which is the money spent on satisfying this type. It is natural to require that a type’s share is proportional to its size.

In a first step, we define a very strong fairness criterion by requiring that all types achieve the same welfare. This is not only unachievable for obvious reasons in single-round PB, but we can show that there are arbitrary long perpetual PB instances where equal welfare remains unachievable. However, while equal welfare is often unattainable, different sets of projects can lead to vastly different distributions of welfare. Thus, as a second fairness criterion, we use the Gini index, a well-known inequality measure for income, to measure inequality with respect to welfare. This measure can be used, e.g., to analyze real-world budget allocations. In this paper, we take a computational approach: given a perpetual PB instance, we can use the Gini index as an optimization goal and search for the least unequal budget allocation. We show that testing for the optimality of a solution is already co-NP-complete, even in simple settings.

As a third fairness criterion we require that all types have the same welfare in the long run, i.e., they are asymptotically equal. Our main result for this fairness criterion is that it is always possible to achieve equal relative satisfaction in the limit if there are only two types. In contrast, equal shares and

*Contact Author

equal-satisfaction are impossible to achieve even in the long run, in particular due to inhomogeneous types. It remains an interesting open problem whether equal relative satisfaction can be guaranteed in the limit for any number of types.

To sum up, our paper contains two main contributions: (i) the framework of perpetual participatory budgeting and (ii) the analytic and computational study of three fairness criteria in this framework. These (strong) fairness criteria cannot be guaranteed in a single round of PB and thus necessitate our perpetual setting.

Related Work. The standard PB setting from the perspective of computational social choice has been extensively described [Lu and Boutilier, 2011; Talmon and Faliszewski, 2019; Aziz and Shah, 2020], and has then been extended in several directions [Jain *et al.*, 2020; Rey *et al.*, 2020; Shapiro and Talmon, 2017; Benadè *et al.*, 2020; Baumeister *et al.*, 2020; Peters *et al.*, 2020; Skowron *et al.*, 2020; Laruelle, 2021]. The main focus of prior works were normative properties: e.g., monotonicity [Talmon and Faliszewski, 2019], strategy-proofness [Goel *et al.*, 2019; Freeman *et al.*, 2019], the core property [Fain *et al.*, 2016], and proportionality [Aziz *et al.*, 2018]. A recent study of district fairness [Hershkowitz *et al.*, 2021] investigates whether a city-wide PB can guarantee each district the social welfare they would have had by running a district-wide PB. While we consider districts (called *types*), our notions of fairness differ.

None of the aforementioned works consider participatory budgeting as a repeating, ongoing process. We are only aware of one exception in which the outcome of the previous year is taken into account to compute a budget allocation [Shapiro and Talmon, 2017] (only for tie-breaking however). We take previous rounds of PB into account in a more comprehensive fashion. Finally, let us mention that our “perpetual” perspective has been also considered in classical voting [Lackner, 2020] and utility aggregation [Freeman *et al.*, 2017; Freeman *et al.*, 2018].

2 Motivating Example

Let us begin with an example that demonstrates the advantages of taking a long-term perspective for PB.

Example 1. *Imagine a city with five inhabitants (1 to 5). There are two districts: 1, 2 and 3 live in the first district (type t_1) and 4 and 5 in the second one (type t_2). A PB process will be run over the next three years with a vote occurring every year. The following table indicates—per year—the proposed projects, their cost and the agents’ approval ballots. A check mark (✓) indicates that the agent approves of the project. The budget limit for every year is 10.*

Projects	Year 1				Year 2				Year 3			
	p_1	p_2	p_3	p_4	p_5	p_6	p_7	p_8	p_9	p_{10}	p_{11}	p_{12}
Cost	6	2	2	4	5	5	3	2	7	7	$4/3$	$29/3$
t_1	1 ✓	2 ✓	3 ✓	4 ✓		6 ✓	7 ✓	8 ✓	9 ✓		11 ✓	
t_2		4 ✓		5 ✓	6 ✓		7 ✓	8 ✓		10 ✓	11 ✓	12 ✓

We assume an online process: when the budgeting decision has to be made for a given year, only the past is known and agents have no knowledge about future years.

Now, suppose that the municipality wants to select budget allocations that maximize either the total number of approvals¹ or the total number of approvals weighted by the cost² (the most commonly used methods in real-life [Aziz and Shah, 2020]). In both cases, all projects that are boxed in the table would be selected (if ties are broken accordingly).

It can be argued that these outcomes are not particularly fair with respects to the two districts. Agents in district t_1 are favored by the outcomes. For example, 1, 2 and 3 approve on average roughly 90% of the selected projects, while 4 and 5 only roughly 40%. As will be shown later in the paper (Examples 2 and 4), it is actually possible to reach a much more equal treatment of the two districts in the last year, by taking into account what happened during the first two years. In this paper, we introduce concepts that make precise in which sense the modified solutions are fairer than the original one, and we discuss whether fair solutions are guaranteed to exist.

3 Perpetual Participatory Budgeting

In essence, our framework consists of a sequence of budgeting problems over several rounds. Let \mathfrak{P} be the set of all the projects occurring throughout the process. Their cost is given by the *cost function* $c : \mathfrak{P} \rightarrow \mathbb{N}$. To simplify the notation, we will write $c(P)$ instead of $\sum_{p \in P} c(p)$ for any $P \subseteq \mathfrak{P}$. Moreover, let \mathcal{N} be the *set of agents* taking part in the process; we assume this set to remain the same in all rounds. Every agent belongs to a type that can represent the district she lives in or any other characteristics. Observe that each agent can have her own type. All our fairness notions and results extend to this special case. Let \mathcal{T} be the set of types, the *type function* $T : \mathcal{N} \rightarrow \mathcal{T}$ indicates for every agent $i \in \mathcal{N}$ her type $T(i)$. For simplicity, we will sometimes consider a type $t \in \mathcal{T}$ as the set of agents having this type $\{i \in \mathcal{N} \mid T(i) = t\}$. In that respect, $|t|$ denotes the number of agents having type $t \in \mathcal{T}$.

Definition 1 (Budgeting problem). *A budgeting problem for round j is defined by the tuple $I_j = \langle \mathcal{P}_j, b_j, A_j \rangle$ where:*

- $\mathcal{P}_j \subseteq \mathfrak{P}$ is the set of available projects,
- $b_j \in \mathbb{N}_{>0}$ is the available budget,
- $A_j : \mathcal{N} \rightarrow 2^{\mathcal{P}_j}$ is the approval function giving for every $i \in \mathcal{N}$ the set of projects $A_j(i) \subseteq \mathcal{P}_j$ she approves of.

We also make the assumption that every project is approved by at least one agent and that every agent approves of at least one project; projects without approvals as well as agents with empty ballots can be removed in a pre-processing stage.

The outcome of a budgeting problem $I_j = \langle \mathcal{P}_j, b_j, A_j \rangle$ is a budget allocation $\pi_j \subseteq \mathcal{P}_j$. It is *feasible* if $c(\pi_j) \leq b_j$ and $\mathcal{A}(I_j)$ is the set of all feasible budget allocations for I_j . It is *exhaustive* if it is feasible and there is no project $p \in \mathcal{P}_j \setminus \pi_j$ such that $c(\pi \cup \{p\}) \leq b_j$. We also speak about feasible and exhaustive ballots using the same definition. Feasible ballots are usually referred to as knapsack ballots.

¹The sum of approvals of the selected projects.

²The sum over the selected projects of their cost times approval.

A perpetual participatory budgeting instance of length $k \in \mathbb{N}_{>0} \cup \{\infty\}$ (or k -PPB instance) is a sequence of k budgeting problems $\mathbf{I} = (I_1, \dots, I_k)$. A vector $\pi = (\pi_1, \dots, \pi_k)$ where $\pi_j \subseteq \mathcal{P}_j$ for every round $j \in \{1, \dots, k\}$ will be called a solution for \mathbf{I} . It is said to be feasible (resp. exhaustive) for \mathbf{I} if every $\pi_j \in \pi$ is feasible (resp. exhaustive) for I_j .³

4 A Fairness Theory for PPB

Solutions can benefit some types while disadvantaging others. To be able to reason about the quality of solutions, we will introduce several *fairness criteria*. In order to discuss whether a solution is fair or unfair, we first need a way to measure the welfare of types.

Definition 2 (Welfare Measure). *A welfare measure F is a function taking as inputs a k -PPB instance \mathbf{I} , a solution π , a type $t \in \mathcal{T}$ and a round $j \in \{1, \dots, k\}$ and returning the welfare score $F(\mathbf{I}, \pi, t, j) \in \mathbb{R}$ for type t of the solution π for the first j rounds of \mathbf{I} .*

Let us begin with fairness criteria. Specific welfare measures are introduced in a second time.

4.1 Fairness Criteria

The foundation of our fairness theory is that the fairest solution should be so that all types are treated exactly the same and thus enjoy the same level of welfare. This requirement is our first fairness criteria.

Definition 3 (Equal- F). *For a welfare measure F , a solution π for the k -PPB instance \mathbf{I} satisfies equal- F at round $j \in \{1, \dots, k\}$ if for every two types $t, t' \in \mathcal{T}$, we have:*

$$F(\mathbf{I}, \pi, t, j) = F(\mathbf{I}, \pi, t', j).$$

Moreover, a solution π satisfies equal- F if it is equal- F at round j for all rounds $j \in \{1, \dots, k\}$.

As Equal- F can be too strong of a requirement, we introduce two relaxations in the following.

A first approach when perfect fairness cannot be achieved, is to try to optimize for it. This idea is particularly relevant when the long-term perspective is adopted as subsequent rounds can compensate for unfairness in previous rounds. We pursue this approach by introducing the Gini coefficient [Gini, 1912] of a solution—a well-known measure of inequality given a multi-set of values—that can be used as a minimization objective. In the following, we use the standard formulation [Blackorby and Donaldson, 1978].

Definition 4 (F -Gini). *Let $\vec{v} = (v_1, \dots, v_k) \in \mathbb{R}^k$ be a vector ordered in non-increasing order, i.e., such that $v_i \geq v_j$ for all $1 \leq i \leq j \leq k$. The Gini coefficient of \vec{v} is given by:*

$$\text{gini}(\vec{v}) = 1 - \frac{\sum_{i=1}^k (2i-1)v_i}{k \sum_{i=1}^k v_i}.$$

For a welfare measure F , the F -Gini coefficient of a solution π for the k -PPB instance \mathbf{I} at round $j \in \{1, \dots, k\}$ is then:

$$\text{gini}_F(\mathbf{I}, \pi, j) = \text{gini}(\vec{F}(\mathbf{I}, \pi, j)),$$

³One could weaken the feasibility requirement of solutions by allowing unused budget to be used in later rounds. For our results, it is not relevant which definition we take.

where $\vec{F}(\mathbf{I}, \pi, j)$ is a vector containing $F(\mathbf{I}, \pi, t, j)$ for all types $t \in \mathcal{T}$, ordered in non-increasing order.

A solution π is F -Gini-optimal at round j with respect to a set S of solutions for \mathbf{I} , if there is no solution $\pi' \in S \setminus \{\pi\}$ with $\text{gini}_F(\mathbf{I}, \pi', j) < \text{gini}_F(\mathbf{I}, \pi, j)$.

It can be checked that F -Gini-optimality is indeed a relaxation of equal- F in the sense that for all welfare measure F , a solution π satisfies equal- F if and only if its F -Gini coefficient reaches 0 (the minimum of the F -Gini coefficient).

Another approach we follow is to require perfect fairness but only in the long run. For that we introduce *convergence to equal- F* which formalizes the idea of asymptotically equalizing the welfare of the different types.

Definition 5 (Convergence to equal- F). *For a welfare measure F , a solution π for the ∞ -PPB instance \mathbf{I} converges to equal- F if for every two types $t, t' \in \mathcal{T}$:*

$$\frac{F(\mathbf{I}, \pi, t, k)}{F(\mathbf{I}, \pi, t', k)} \xrightarrow{k \rightarrow +\infty} 1.$$

We will study the computational complexity of the following problems related to these fairness criteria. Note that this computational analysis does not apply for convergence to equal- F as we deal with infinite sequences there.

EQUAL- F	
Input:	A k -PPB instance $\mathbf{I} = (I_1, \dots, I_k)$ and a solution $\pi = (\pi_1, \dots, \pi_{k-1})$.
Question:	Is there a non-empty and feasible budget allocation π_k for I_k such that $(\pi_1, \dots, \pi_{k-1}, \pi_k)$ provides equal- F at round k ?
F -GINI-OPTIMALITY	
Input:	A k -PPB instance $\mathbf{I} = (I_1, \dots, I_k)$ and a solution $\pi = (\pi_1, \dots, \pi_k)$.
Question:	Is π Gini-optimal at round k w.r.t. all non-empty, feasible solutions?

4.2 Welfare Measures

The first welfare measures we define are based on the satisfaction of an agent. Even though agents can have personal utility functions to express their satisfaction for a given outcome (e.g. [Peters *et al.*, 2020]), this information is usually private, i.e., unknown to the decision maker. We thus need to approximate the satisfaction of an agent. We use a standard definition for satisfaction [Talmon and Faliszewski, 2019].

Definition 6 (Satisfaction). *Let $\mathbf{I} = (I_1, \dots, I_k)$ be a k -PPB instance and $\pi = (\pi_1, \dots, \pi_k)$ a solution for \mathbf{I} . For round $j \in \{1, \dots, k\}$, whose budgeting problem is $\langle \mathcal{P}_j, b_j, A_j \rangle$, we define the marginal satisfaction of agent $i \in \mathcal{N}$ as:*

$$\text{sat}_j^m(\mathbf{I}, \pi_j, i) = c(\pi_j \cap A_j(i)).$$

Moreover, the marginal satisfaction and the satisfaction of a type $t \in \mathcal{T}$ for round $j \in \{1, \dots, k\}$ are defined by:

$$\text{sat}_j^m(\mathbf{I}, \pi_j, t) = \frac{1}{|t|} \sum_{i \in t} \text{sat}_j^m(\mathbf{I}, \pi_j, i)$$

$$\text{sat}_j(\mathbf{I}, \pi, t) = \sum_{1 \leq j^* \leq j} \text{sat}_{j^*}^m(\mathbf{I}, \pi_{j^*}, t).$$

	Satisfaction				Relative Satisfaction			Share		
	2 agents	3 agents	> 3 agents	Complex.	2 types	> 2 types	Complex.	2 agents	> 2 agents	Complex.
equal- F	✗	✗	✗	NP-c	✗	✗	NP-c	✗	✗	NP-c
conv. to equal- F	✓	✓ (ex. ballots)	✗		✓ (knap. ballots)	?		✓	✗	
F -Gini optimality	✓	✓	✓	co-NP-c	✓	✓	co-NP-c	✓	✓	co-NP-c

Table 1: Summary of the results. The columns specifying a number of agents/types are for existence guarantees: a ✓ indicates that for all instances with the specified number of agents/types, there exists a solution satisfying the fairness criteria; and the ✗ the opposite. The tags “ex. ballots” and “knap. ballots” indicates that the result only holds with exhaustive or knapsack ballots. The column “Complex.” indicates the computational complexity of the problems stated in Section 4.1 where NP-c stands for NP complete and co-NP-c for co-NP complete.

Example 2. Let us illustrate satisfaction on Example 1. It can be checked that by the end of year 2, the satisfaction of type t_1 is $17 + 2/3$ while that of type t_2 is only 8. Hence, selecting only p_{12} in the last year would lead to a solution such that both types would have a satisfaction of $17 + 2/3$.

One potential drawback of satisfaction is its strong dependence on voters’ approval sets. For example, if agent 1 approves a proper subset of agent 2’s approved projects ($A_j(1) \subset A_j(2)$) and all their approved projects are funded ($A_j(2) \subseteq \pi_j$), then agent 2 is more satisfied than agent 1. However, it can be argued that the welfare of both agents should be equal as all projects they wanted to be funded have actually been funded; neither agent 1 or 2 can be made happier (subject to the available information). To take this into account, we define *relative satisfaction* which measures how close the satisfaction of an agent is to his best-case scenario.

Definition 7 (Relative satisfaction). Let $\mathbf{I} = (I_1, \dots, I_k)$ be a k -PPB instance and $\pi = (\pi_1, \dots, \pi_k)$ a solution for \mathbf{I} . For round $j \in \{1, \dots, k\}$ corresponding to the budgeting problem $\langle \mathcal{P}_j, b_j, A_j \rangle$, we define the marginal relative satisfaction of agent $i \in \mathcal{N}$ as:

$$rsat_j^m(\mathbf{I}, \pi_j, i) = \frac{c(\pi_j \cap A_j(i))}{\max\{c(A) \mid A \subseteq A_j(i) \text{ and } c(A) \leq b_j\}}.$$

Moreover, marginal relative satisfaction and the relative satisfaction of a type $t \in \mathcal{T}$ for the round $j \in \{1, \dots, k\}$ are defined as follows:

$$rsat_j^m(\mathbf{I}, \pi_j, t) = \frac{1}{|t|} \sum_{i \in t} rsat_j^m(\mathbf{I}, \pi_j, i)$$

$$rsat_j(\mathbf{I}, \pi, t) = \sum_{1 \leq j^* \leq j} rsat_{j^*}^m(\mathbf{I}, \pi_{j^*}, t).$$

Example 3. In Example 1, the relative satisfaction scores by the end of year are $23/12$ for type t_1 and $73/60$ for type t_2 . One can verify that there is no budget allocation for the third year that would lead to equal-relative satisfaction.

Satisfaction and relative satisfaction are two concepts which relate to the idea of utilitarianism. Indeed, only the impact of the selected solution on the agents or types is taken into consideration and not the way the resources were spent. Although utilitarian welfare is attractive, other notions can be considered in participatory budgeting. The most important alternative might be distributive welfare which aims at spending an equal amount of resources on each agent or type. To account for distributive welfare, we introduce another welfare measure called the *share* of a type.

Definition 8 (Share). Let $\mathbf{I} = (I_1, \dots, I_k)$ be a k -PPB instance with a solution $\pi = (\pi_1, \dots, \pi_k)$. For round $j \in \{1, \dots, k\}$ with budgeting problem $\langle \mathcal{P}_j, b_j, A_j \rangle$, the marginal share of agent $i \in \mathcal{N}$ is defined as:

$$share_j^m(\mathbf{I}, \pi_j, i) = \sum_{p \in \pi_j \cap A_j(i)} \frac{c(p)}{|\{i' \in \mathcal{N} \mid p \in A_j(i')\}|}$$

Moreover, the marginal share and the share of a type $t \in \mathcal{T}$ for round $j \in \{1, \dots, k\}$ are defined as:

$$share_j^m(\mathbf{I}, \pi_j, t) = \frac{1}{|t|} \sum_{i \in t} share_j^m(\mathbf{I}, \pi_j, i)$$

$$share_j(\mathbf{I}, \pi, t) = \sum_{1 \leq j^* \leq j} share_{j^*}^m(\mathbf{I}, \pi_{j^*}, t).$$

Example 4. Once again coming back to Example 1, we can show that equal-share can be achieved by the end of the last year. Indeed by the end of year 2, the shares are $5 + 1/3$ for t_1 and 2 for t_2 . Now, by selecting p_{10} and p_{11} in the third year, we reach a solution where each type has a share of $5 + 2/3$.

Observe that trying to equalize the shares of the different types requires the *average* share of each type to be equal, meaning that we require the total share of a type to be proportional to its size. In this sense, the fairness criteria equal-share can be considered a proportionality concept.

Note that relative share—in contrast to relative satisfaction—is not a very sensible property as distributive fairness should hardly depend on the ballots.

5 Realizing Fairness

We will now explore the fairness criteria and the welfare measures we have defined previously. All our results are summarized in Table 1. Note that most of the proofs are omitted due to space constraints but can be found in the appendix.

5.1 Achieving Perfect Fairness: Equal- F

We first explore the criteria that we consider to represent a situation of perfect fairness: equal- F .

Unfortunately, it is easy to check that equal- F cannot be guaranteed even for a single round (except by selecting an empty budget allocation) for all of our welfare measures. Consider the following example with two agents where no non-empty solution satisfies either equal-satisfaction, equal-relative satisfaction or equal-share in any round.

Example 5. Let I be a k -PPB instance with two agents 1 and 2 of types t_1 and t_2 respectively. Furthermore, let $b_j = 1$ for every round $j \in \{1, \dots, k\}$ and let $c(p) = 1$ for all $p \in \mathfrak{P}$.

In the first round, agent 1 approves only of project p_1 and agent 2 only of p_2 . In all following rounds, they both only approve of p_1 . Assume w.l.o.g. that p_1 is selected in the first round. Then, for every solution $\pi = (\pi_1, \dots, \pi_k)$, at round $j \in \{2, \dots, k\}$, we have $F_j(I, \pi, t_1) = 1 + F_j(I, \pi, t_2)$ for all three of our welfare measures F .

The example shows that in general, equal- F cannot be satisfied for our welfare measures. However, it could still be achieved on some specific instances. It turns out that for all three welfare measures, checking the existence of an equal- F solution is an NP-complete problem.

Proposition 1. The EQUAL-SATISFACTION and EQUAL-RELATIVE SATISFACTION problems are strongly NP-complete even if there is only one round.

Proposition 2. The EQUAL-SATISFACTION and EQUAL-SHARE problems are weakly NP-complete even if there is only one round and there are only two agents.

Proof (Sketch). We reduce from SUBSET SUM [Karp, 1972], i.e., the problem of finding a subset Z' of a set $Z \subset \mathbb{Z}$ such that $\sum Z' = 0$: Consider a PB instance with a project p_z for every $z \in Z$ such that $c(p_z) = |z|$. Further, there are two agents, v_+ and v_- approving all projects p_z with $z \geq 0$ resp. $z < 0$. We claim that this is a positive instance of EQUAL-SATISFACTION and EQUAL-SHARE if and only if Z is a positive instance of SUBSET SUM. \square

We observe that both results still hold if we additionally require exhaustiveness, i.e., if we ask whether there is an exhaustive solution that satisfies equal- F .

5.2 Optimizing for Fairness: F -Gini-optimality

Let us now turn our attention to F -Gini-optimality. Note first that—by definition—there will always be for every instance at least one solution which is F -Gini-optimal. Therefore, the main questions here concern computational problems and are not about existence guarantees.

We first show that F -GINI-OPTIMALITY is co-NP-complete for both satisfaction and share.

Proposition 3. SATISFACTION-GINI-OPTIMALITY, RELATIVE-SATISFACTION-GINI-OPTIMALITY and SHARE-GINI-OPTIMALITY are weakly co-NP-complete even if there is only one round and two agents.

We note that it is also co-NP-complete to check whether a solution is Gini-optimal among exhaustive solutions. An interesting open question is the complexity of finding a Gini-optimal solution that maximizes the overall welfare of the population.

5.3 Achieving Fairness in the Long Run: Convergence to Equal- F

Let us conclude our analysis by investigating convergence to equal- F . We first show that for two agents convergence to equal-satisfaction can always be guaranteed (under mild additional assumptions).

Proposition 4. Consider an ∞ -PPB instance I with two agents such that there exists a constant $B^* \in \mathbb{N}$ with $b_j \leq B^*$ for every round j . Furthermore, assume for every round and both agents that there is a project p with $c(p) \leq b_j$ that the agent approves of. Then, there is a non-empty feasible solution that converges to equal-satisfaction.

Proof. Call the agents 1 and 2 and assume they belong to types t_1 and t_2 respectively (as equal-satisfaction is trivially satisfied if there is only one type). We claim that there exists a solution π such that for every round j , we can guarantee:

$$sat_j(I, \pi, t_1) - B^* \leq sat_j(I, \pi, t_2) \leq sat_j(I, \pi, t_1) + B^*.$$

Let us prove the claim by induction. For the first round, it is clear that whichever non-empty budget allocation has been chosen, we have $0 \leq sat_1(I, \pi, t) \leq B^*$ for $t \in \{t_1, t_2\}$.

Now assume the claim holds for round $j - 1$. W.l.o.g. assume $sat_{j-1}(I, \pi, t_2) \leq sat_{j-1}(I, \pi, t_1)$. Let p be a project approved by 2 such that $c(p) \leq b_j$. Then, we set $\pi_j = \{p\}$. This implies that

$$sat_j^m(I, \pi_j, t_1) \leq sat_j^m(I, \pi_j, t_2) \leq B^*$$

From this together with the induction hypothesis and the assumption $sat_{j-1}(I, \pi, t_2) \leq sat_{j-1}(I, \pi, t_1)$ we can conclude that the claim also holds in round j .

Now, we know that $sat_j(I, \pi, t_1) + sat_j(I, \pi, t_2) \geq \sum_{j'=1}^j c(\pi_{j'})$. Together with the claim, this implies that $\lim_{j \rightarrow +\infty} (sat_j(I, \pi, t_1)) = \lim_{j \rightarrow +\infty} (sat_j(I, \pi, t_2)) = +\infty$.

Therefore, the proposition follows from the following:

$$\frac{sat_j(I, \pi, t_1) - B^*}{sat_j(I, \pi, t_1)} \leq \frac{sat_j(I, \pi, t_2)}{sat_j(I, \pi, t_1)} \leq \frac{sat_j(I, \pi, t_1) + B^*}{sat_j(I, \pi, t_1)}.$$

\square

Unfortunately, this result cannot be generalized—even for three agents—as the following example shows.

Example 6. Let I be a ∞ -PPB instance with three agents 1, 2, 3 where agent 1 has type t_1 and agents 2 and 3 have type t_2 . Assume $b_j = 1$ for every round j and $c(p) = 1$ for all projects $p \in \mathfrak{P}$. In every round, there are two projects and agent 1 approves of both, 2 approves of only one and 3 of the other one. Then, for every non-empty feasible solution π and every round j , we have $sat_j(I, \pi, t_1) = j$ and $sat_j(I, \pi, t_2) = \frac{j}{2}$. Therefore, we have

$$\lim_{j \rightarrow +\infty} \left(\frac{sat_j(I, \pi, t_2)}{sat_j(I, \pi, t_1)} \right) = \frac{1}{2}.$$

This counter-example can be avoided if we impose some restrictions on the ballots the agents may submit. Indeed, if ballots are exhaustive then, for three agents we can always find a solution that converge to equal-satisfaction.

Proposition 5. Consider an ∞ -PPB instance I with three agents where the ballot of each agent is exhaustive in every round and there exists a constant $B^* \in \mathbb{N}$ with $b_j \leq B^*$ for every round j . Then, there is a non-empty feasible solution that converges to equal-satisfaction.

However, by increasing the number of agents we again encounter an impossibility, even with these restricted ballots.

Example 7. Let \mathbf{I} be a ∞ -PPB instance. In every round j , we have $b_j = 10$, there are eight agents $1, \dots, 8$ such that $1, 2, 3$ have type t_1 and $4, 5, 6, 7, 8$ have type t_2 . Furthermore, there are six projects p_1, \dots, p_6 such that $c(p_1) = c(p_2) = c(p_3) = 5$ and $c(p_4) = c(p_5) = c(p_6) = 3$. The ballots are such that, for every round j :

$$\begin{aligned} A_j(1) &= \{p_1, p_4\} & A_j(2) &= \{p_2, p_5\} & A_j(3) &= \{p_3, p_6\} \\ A_j(4) &= \{p_1, p_2\} & A_j(5) &= \{p_1, p_3\} & A_j(6) &= \{p_2, p_3\} \\ A_j(7) &= \{p_4, p_5, p_6\} & A_j(8) &= \{p_4, p_5, p_6\} \end{aligned}$$

We leave it to reader to check that for each project the marginal satisfaction for type t_2 is higher than for type t_1 . This directly implies that there can be no non-empty solution converging to equal-satisfaction.

Results about convergence to equal-share are very similar to the ones with equal-satisfaction. By a similar argument as for Proposition 4, we can show that convergence to equal-share can be achieved for two agents. Unfortunately, we cannot go far beyond this, as the following example shows.

Example 8. Consider again the same ∞ -PPB instance as in Example 7. We claim that for every project, selecting it would lead to a higher share for type t_2 than that of type t_1 . For project p_1 , we have $\text{share}_1(\mathbf{I}, \{p_1\}, t_1) = \frac{1}{3} \cdot \frac{5}{3} = \frac{5}{9}$ but $\text{share}_1(\mathbf{I}, \{p_1\}, t_2) = \frac{1}{5}(\frac{5}{3} + \frac{5}{3}) = \frac{2}{3}$. The case for p_2 and p_3 is similar. For p_4 , we have $\text{share}_1(\mathbf{I}, \{p_4\}, t_1) = \frac{1}{3} \cdot \frac{3}{3} = \frac{1}{3}$ but $\text{share}_1(\mathbf{I}, \{p_4\}, t_2) = \frac{1}{5}(\frac{3}{3} + \frac{3}{3}) = \frac{2}{5}$. The case for projects p_5 and p_6 is similar. It follows that, in this example, we cannot have convergence to equal shares.

Results are more positive when it comes to relative satisfaction. Indeed, we can guarantee convergence to equal-relative satisfaction when there are two types. Note that this result is much more general than Lemma 4 as types may contain an arbitrary number of agents. The proof is based on the following lemma stating that with two types, we can always favor one type over the other.

Lemma 6. Let \mathbf{I} be a k -PPB instance with non-empty knapsack ballots and two types t_1 and t_2 . Then, in every round $j \in \{1, \dots, k\}$ there are two feasible budget allocations π_1 and π_2 such that:

$$0 < \text{rsat}_j^m(\mathbf{I}, \pi_1, t_1) \geq \text{rsat}_j^m(\mathbf{I}, \pi_1, t_2) \quad \text{and} \\ \text{rsat}_j^m(\mathbf{I}, \pi_2, t_1) \leq \text{rsat}_j^m(\mathbf{I}, \pi_2, t_2) > 0.$$

Thanks to this lemma, using a similar line of reasoning as in Proposition 4, we can show that for two types we can find a solution that converges to equal relative satisfaction.

Theorem 7. Assume that \mathbf{I} is an ∞ -PPB-instance with non-empty knapsack ballots such that there are only two types and a B^* such that $b_j \leq B^*$ for all rounds j . Then, there is a non-empty feasible solution for \mathbf{I} that converges to equal relative satisfaction.

It is important to mention that the proofs of Lemma 6 and that of Theorem 7 are both constructive, in the sense that they show how to compute the relevant solutions. However, this construction does not guarantee the solution to be exhaustive. To achieve this, an additional ballot restriction is necessary.

Corollary 8. Consider an ∞ -PPB instance \mathbf{I} that satisfies all the conditions of Theorem 7. Then, there exists a non-empty feasible solution $\pi = (\pi_1, \pi_2, \dots)$ for \mathbf{I} that (i) converges to equal relative satisfaction and (ii) such that for each round j there is an agent i with $A_j(i) \subseteq \pi_j$. In particular, if all ballots are exhaustive, then every budget allocation in π is exhaustive.

Whether Theorem 7 and Corollary 8 can be extended to three and more types remains an important open question.

6 Conclusion

In this paper, we have introduced a model of participatory budgeting (called *perpetual participatory budgeting*) that takes into account the temporal component of a PB process. We have further defined a theory of fairness for this model by introducing several fairness criteria. We considered both egalitarian concepts based on voters' (relative) satisfaction and a form of proportionality based on shares. For corresponding axiomatic properties, we studied whether (and when) we can guarantee these to hold as well as the computational complexity of verifying them.

We can conclude that taking the long-term viewpoint allows us to approximate forms of fairness that cannot be obtained in single-round PB instances. This was already visible in our starting example in Section 2. Beyond that, we have established three strong fairness concepts that required a thorough analysis to judge their applicability. On the one hand, we have seen that they cannot be guaranteed in general. On the other hand, and we have identified special cases where some of these strong are guaranteed to hold.

Several research directions can be pursued within our proposed framework. For instance, it would be interesting to look for natural PB procedures that compute solutions satisfying (or approximating) our fairness criteria. In the light recent works (e.g., [Hershkowitz *et al.*, 2021]), a relevant question concerns the price of fairness, i.e., how much the satisfaction has to be reduced in order to achieve fairness. The fairness criteria we introduced may not be compatible with efficiency notions such as Pareto-optimality. Although the tension between fairness and efficiency is well-known (see, e.g., [Peters *et al.*, 2020]), it would be interesting to study the combination of fairness and efficiency criteria in our setting. Moreover, PB generalizes multi-winner voting in a way that makes it closer to the fair division of indivisible items [Bouveret *et al.*, 2016]. This linked was already hinted in Example 5 as applying "up-to-one project" criteria would have made the solution fair. It would then be interesting to investigate if and how fairness criteria from fair division (for instance envy-freeness or EF1) could be successfully adapted and applied to PB. Finally, by considering real-world data of repeated PB referenda one can analyze the possible gains of taking a long-term perspective as we propose here.

Acknowledgements

This work was supported by the Austrian Science Fund (FWF), project P31890.

References

- [Aziz and Shah, 2020] Haris Aziz and Nisarg Shah. Participatory budgeting: Models and approaches. In Tamás Rudas and Gábor Péli, editors, *Pathways between Social Science and Computational Social Science: Theories, Methods and Interpretations*. Springer, 2020.
- [Aziz et al., 2018] Haris Aziz, Barton E. Lee, and Nimrod Talmon. Proportionally representative participatory budgeting: Axioms and algorithms. In *Proc. of 17th AAMAS Conference*, pages 23–31, 2018.
- [Baumeister et al., 2020] Dorothea Baumeister, Linus Boes, and Tessa Seeger. Irresolute approval-based budgeting. In *Proc. of 19th AAMAS Conference*, pages 1774–1776, 2020.
- [Benadè et al., 2020] Gerdus Benadè, Swaprava Nath, Ariel D. Procaccia, and Nisarg Shah. Preference elicitation for participatory budgeting. *Management Science*, 2020.
- [Blackorby and Donaldson, 1978] Charles Blackorby and David Donaldson. Measures of relative equality and their meaning in terms of social welfare. *Journal of Economic Theory*, 18(1):59–80, 1978.
- [Bouveret et al., 2016] Sylvain Bouveret, Yann Chevaleyre, and Nicolas Maudet. Fair allocation of indivisible goods. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 17. Cambridge University Press, 2016.
- [Cabannes, 2004] Yves Cabannes. Participatory budgeting: a significant contribution to participatory democracy. *Environment and Urbanization*, 16(1):27–46, 2004.
- [City of Paris, 2020] City of Paris. Paris Budget Participatif. <https://budgetparticipatif.paris.fr/bp/>, 2020. Last accessed on January the 18th 2020.
- [Dias et al., 2019] Nelson Dias, Sahsil Enríquez, and Simone Júlio, editors. *The Participatory Budgeting World Atlas*. Epopee Records: Officinal Coordination, 2019.
- [Fain et al., 2016] Brandon Fain, Ashish Goel, and Kamesh Munagala. The core of the participatory budgeting problem. In *Proc. of 12th WINE*, pages 384–399, 2016.
- [Freeman et al., 2017] Rupert Freeman, Seyed Majid Zahedi, and Vincent Conitzer. Fair and efficient social choice in dynamic settings. In *Proc. of 26th IJCAI*, pages 4580–4587, 2017.
- [Freeman et al., 2018] Rupert Freeman, Seyed Majid Zahedi, Vincent Conitzer, and Benjamin C. Lee. Dynamic proportional sharing: A game-theoretic approach. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 2(1):3:1–3:36, 2018.
- [Freeman et al., 2019] Rupert Freeman, David M. Pennock, Dominik Peters, and Jennifer Wortman Vaughan. Truthful aggregation of budget proposals. In *Proc. of 20th ACM-EC Conference*, pages 751–752, 2019.
- [Gini, 1912] Corrado Gini. *Variabilità e Mutabilità: Contributo allo Studio delle Distribuzioni e Delle Relazioni Statistiche*. P. Cuppini, 1912.
- [Goel et al., 2019] Ashish Goel, Anilesh K. Krishnaswamy, Sukolsak Sakshuwong, and Tanja Aitamurto. Knapsack voting: Voting mechanisms for participatory budgeting. *ACM Transaction on Economics and Computation*, pages 8:1–8:27, 2019.
- [Hershkowitz et al., 2021] D. Ellis Hershkowitz, Anson Kahng, Dominik Peters, and Ariel D. Procaccia. District-fair participatory budgeting. In *Proc. of 35th AAAI Conference*, 2021.
- [Jain et al., 2020] Pallavi Jain, Krzysztof Sornat, and Nimrod Talmon. Participatory budgeting with project interactions. In *Proc. of 29th IJCAI*, pages 386–392, 2020.
- [Karp, 1972] Richard M. Karp. Reducibility among combinatorial problems. In *Proc. of Symposium on the Complexity of Computer Computations*, pages 85–103. Springer, 1972.
- [Lackner, 2020] Martin Lackner. Perpetual voting: Fairness in long-term decision making. In *Proc. of 34th AAAI Conference*, pages 2103–2110, 2020.
- [Laruelle, 2021] Annick Laruelle. Voting to select projects in participatory budgeting. *European Journal of Operational Research*, 288(2):598–604, 2021.
- [Lu and Boutilier, 2011] Tyler Lu and Craig Boutilier. Budgeted social choice: From consensus to personalized decision making. In *Proc. of 22nd IJCAI*, pages 280–286, 2011.
- [New York City Council, 2020] New York City Council. Participatory Budgeting. <https://council.nyc.gov/pb/>, 2020. Last accessed on October the 8th 2020.
- [Peters et al., 2020] Dominik Peters, Grzegorz Pierczyński, and Piotr Skowron. Proportional participatory budgeting with cardinal utilities. *arXiv preprint arXiv:2008.13276*, 2020.
- [Rey et al., 2020] Simon Rey, Ulle Endriss, and Ronald de Haan. Designing participatory budgeting mechanisms grounded in judgment aggregation. In *Proc. of 17th KR Conference*, 2020.
- [Shah, 2007] Anwar Shah, editor. *Participatory Budgeting. Public Sector Governance and Accountability Series*. The World Bank, Washington, DC, 2007.
- [Shapiro and Talmon, 2017] Ehud Shapiro and Nimrod Talmon. A participatory democratic budgeting algorithm. *arXiv preprint arXiv:1709.05839*, 2017.
- [Skowron et al., 2020] Piotr Skowron, Arkadii Slinko, Stanisław Szufa, and Nimrod Talmon. Participatory budgeting with cumulative votes. *arXiv preprint arXiv:2009.02690*, 2020.
- [Talmon and Faliszewski, 2019] Nimrod Talmon and Piotr Faliszewski. A framework for approval-based budgeting methods. In *Proc. of 33rd AAAI Conference*, 2019.