Fairness in Long-Term Participatory Budgeting

Extended Abstract

Martin Lackner
DBAI, TU Wien
Vienna, Austria
lackner@dbai.tuwien.ac.at

Jan Maly
DBAI, TU Wien
Vienna, Austria
jmaly@dbai.tuwien.ac.at

Simon Rey
ILLC, University of Amsterdam
Amsterdam, The Netherlands
s.j.rey@uva.nl

ABSTRACT

Participatory Budgeting processes are usually designed to span several years, with referenda for new budget allocations taking place regularly. This paper presents the first formalization of long-term PB. We introduce a theory of fairness for this setting, investigate under which conditions our fairness criteria can be satisfied, and analyze the computational complexity of verifying them.

KEYWORDS

Participatory Budgeting; Computational Social Choice; Long-term fairness

ACM Reference Format:
Martin Lackner, Jan Maly, and Simon Rey. 2021. Fairness in Long-Term Participatory Budgeting: Extended Abstract. In *Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), Online, May 3–7, 2021*, IFAAMAS, 3 pages.

1 INTRODUCTION

Participatory Budgeting (PB) is a democratic tool in which citizens are asked their opinion on how to use a public budget [6, 23]. This process, which was invented in Brazil, is now used in many cities all around the world [8]. PB is usually planned to run for several years. For instance, a participatory budgeting process in Paris spanned 6 years [7], and New York runs an ongoing program since 2011 [20]. 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 [2–4, 9, 10, 14–16, 18, 19, 21, 22, 24–26], 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. We intend to close this gap by introducing *perpetual participatory budgeting*, a formal model which captures key characteristics of long-term PB. This model is inspired by related work on voting [17] and utility aggregation [11, 12].

The long-term viewpoint of perpetual PB leads to conceptual challenges but brings notable advantages. In this paper, we focus on notions of fairness in this setting and analyze to which extent stronger fairness guarantees can be achieved in long-term processes. In particular, we are concerned with fairness towards *types* of voters, where a type is a pre-defined subset of voters, for example all voters in a certain district or socio-demographic groups.

Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), U. Endriss, A. Nowé, F. Dignum, A. Lomuscio (eds.), May 3–7, 2021, Online.

2 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 projects occurring throughout the process. Their cost is given by the *cost function* $c : \mathfrak{P} \to \mathbb{N}$. 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.[1] Let $\mathcal{T}$ be the set of types, *the type function* $T : \mathcal{N} \to \mathcal{T}$ 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,*
- *$A_j : \mathcal{N} \to 2^{\mathcal{P}_j}$ 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 make the assumption that every project is approved by at least one agent and that every agent approves of at least one project.

The outcome of a budgeting problem $I_j = \langle \mathcal{P}_j, b_j, A_j \rangle$ is a budget allocation $\pi_j \subseteq \mathcal{P}_j$ which is *feasible* if $c(\pi_j) \leq b_j$. It is *exhaustive* if it is feasible and there is no $p \in \mathcal{P}_j \setminus \pi$ such that $c(\pi \cup \{p\}) \leq b_j$.

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 $\boldsymbol{I} = (I_1, \ldots, I_k)$. A feasible solution is then a vector $\boldsymbol{\pi} = (\pi_1, \ldots, \pi_k)$ where for every round $j \in \{1, \ldots, k\}$, $\pi_j \subseteq \mathcal{P}_j$ is feasible for I_j.

3 A FAIRNESS THEORY

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.

A welfare measure F is a function taking as inputs a k-PPB instance $\boldsymbol{I}$, a solution $\boldsymbol{\pi}$, a type $t \in \mathcal{T}$ and a round $j \in \{1, \ldots, k\}$, and returning the "welfare score" $F(\boldsymbol{I}, \boldsymbol{\pi}, t, j) \in \mathbb{R}$ of the solution $\boldsymbol{\pi}$ for type t at rounds 1 to j given the instance $\boldsymbol{I}$.

We consider three fairness criteria that tell us whether a distribution of welfare is fair. Our first criteria might be the simplest: the goal is to equalize the welfare measures between all types.

Definition 2 (Equal-F). *For a welfare measure F, a solution $\boldsymbol{\pi}$ for the k-PPB instance $\boldsymbol{I}$ satisfies* equal-F *if for every two types $t, t' \in \mathcal{T}$ and every round $j \in \{1, \ldots, k\}$, we have $F(\boldsymbol{I}, \boldsymbol{\pi}, t, j) = F(\boldsymbol{I}, \boldsymbol{\pi}, t', j)$.*

[1] It is also possible that each voter has her own type. Therefore, fairness towards individual voters can be considered a special case of fairness towards types.

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 (i.e. feasible) ballots. The column "Complex." lists the computation complexity of checking whether there exists a solution satisfying equal-F and respectively whether a solution is Gini-optimal. This analysis is not relevant for convergence to Equal-F as we deal with infinite sequences

	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	✗	✗	?
Convergence to Equal-F	✓	✓ (ex. ballots)	✗		✓ (knap. ballots)	?		✓	✗	
F-Gini optimality	✓	✓	✓	?	✓	✓	?	✓	✓	co-NP-c

Equalizing the welfare scores is, however, often too strong of a requirement. A weaker criteria is to try to converge toward equality in the long run.

Definition 3 (Convergence to Equal-F). *For a welfare measure F, a solution $\boldsymbol{\pi}$ for the ∞-PPB instance $\boldsymbol{I}$* converges to equal-F *if for every two types $t, t' \in \mathcal{T}$:*

$$\frac{F(\boldsymbol{I}, \boldsymbol{\pi}, t, k)}{F(\boldsymbol{I}, \boldsymbol{\pi}, t', k)} \underset{k \to +\infty}{\longrightarrow} 1.$$

Another approach could be to try to optimize for fairness. A well studied way of doing that is to search for the solution with the lowest Gini coefficient [13]. In the following, we will use a formulation of the Gini coefficient from [5].

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*

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

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

$$gini_F(\boldsymbol{I}, \boldsymbol{\pi}, j) = gini(\vec{F}(\boldsymbol{I}, \boldsymbol{\pi}, j)),$$

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

A solution $\boldsymbol{\pi}$ is thus F-Gini-optimal *at round j with respect to a set S of solutions for $\boldsymbol{I}$, if no solution $\boldsymbol{\pi}' \in S \setminus \{\boldsymbol{\pi}\}$ is such that $gini_F(\boldsymbol{I}, \boldsymbol{\pi}', j) < gini_F(\boldsymbol{I}, \boldsymbol{\pi}, j)$.*

In the rest of this paper, we introduce three welfare measures: two of which are more egalitarian in nature while the other deals with distributive fairness.

4 WELFARE MEASURES

The first welfare measures we investigate are based on the satisfaction of an agent. The true satisfaction of an agent is usually unknown. We approximate it as in [26].

Definition 5 (Satisfaction and Relative satisfaction). *For a k-PPB instance $\boldsymbol{I}$ and a solution $\boldsymbol{\pi} = (\pi_1, \dots, \pi_k)$, we define the* satisfaction *of a type $t \in \mathcal{T}$ for round $j \in \{1, \dots, k\}$, whose budgeting problem is $\langle \mathcal{P}_j, b_j, A_j \rangle$, as:*

$$sat_j(\boldsymbol{I}, \boldsymbol{\pi}, t) = \frac{1}{|t|} \sum_{i \in t} \sum_{1 \leq j' \leq j} c(\pi_j \cap A_j(i)).$$

Similarly the relative satisfaction of type $t \in \mathcal{T}$ is

$$rsat_j(\boldsymbol{I}, \boldsymbol{\pi}, t) = \frac{1}{|t|} \sum_{i \in t} \sum_{1 \leq j' \leq j} \frac{c(\pi_j \cap A_j(i))}{\max\{c(A) \mid A \subseteq A_j(i) \land c(A) \leq b_j\}}.$$

Trying to achieve the same level of satisfaction for all types might require spending more resources for one type than for another (in particular if the ballots are less uniform for one type). In that sense these welfare measure lead to an egalitarian approach. An important alternative is distributive welfare: trying to spend the same amount of resources on each type. This is formalized by the *share* of a type.

Definition 6 (Share). *Let $\boldsymbol{I} = (I_1, \dots, I_k)$ be a k-PPB instance with a solution $\boldsymbol{\pi} = (\pi_1, \dots, \pi_k)$. For round $j \in \{1, \dots, k\}$ with budgeting problem $\langle \mathcal{P}_j, b_j, A_j \rangle$, the* share *of a type $t \in \mathcal{T}$ is*

$$share_j(\boldsymbol{I}, \boldsymbol{\pi}, t) = \frac{1}{|t|} \sum_{i \in t} \sum_{1 \leq j' \leq j} \sum_{p \in \pi_{j'} \cap A_{j'}(i)} \frac{c(p)}{|\{i' \in \mathcal{N} \mid p \in A_{j'}(i')\}|}.$$

The share also leads to more proportional criteria (e.g. [1]) as equalizing the share between types means requiring the total share of a type to be proportional to its size.

All three welfare measures have been studied both in terms of existence (can we always find a solution satisfying a given fairness criteria) and in terms of computational complexity (how hard is it that compute a solution satisfying a fairness criteria). All the results are summarized in Table 1.

5 CONCLUSION

We introduced the first framework to study long-term participatory budgeting. Taking the viewpoint of perpetual PB allowed us to achieve forms of fairness that cannot be obtained in single-round PB. Several research directions can be pursued within this framework. It would be interesting to look for natural PB procedures to compute solutions satisfying (or approximating) our fairness criteria. Another interesting question when studying fairness with respect to types is the price of fairness as studied in [15].

ACKNOWLEDGMENTS

The Austrian Science Fund (FWF): P31890 supported this work.

REFERENCES

[1] Haris Aziz, Barton E. Lee, and Nimrod Talmon. 2018. Proportionally Representative Participatory Budgeting: Axioms and Algorithms. In *Proceedings of the 17th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*. 23–31.

[2] Haris Aziz and Nisarg Shah. 2020. Participatory Budgeting: Models and Approaches. In *Pathways between Social Science and Computational Social Science: Theories, Methods and Interpretations*, Tamás Rudas and Gábor Péli (Eds.). Springer.

[3] Dorothea Baumeister, Linus Boes, and Tessa Seeger. 2020. Irresolute Approval-based Budgeting. In *Proceedings of the 19th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*. 1774–1776.

[4] Gerdus Benadè, Swaprava Nath, Ariel D. Procaccia, and Nisarg Shah. 2020. Preference Elicitation for Participatory Budgeting. *Management Science* (2020).

[5] Charles Blackorby and David Donaldson. 1978. Measures of Relative Equality and their Meaning in Terms of Social Welfare. *Journal of Economic Theory* 18, 1 (1978), 59–80.

[6] Yves Cabannes. 2004. Participatory Budgeting: a Significant Contribution to Participatory Democracy. *Environment and Urbanization* 16, 1 (2004), 27–46.

[7] City of Paris. 2020. Paris Budget Participatif. https://budgetparticipatif.paris.fr/bp/. Last accessed on October the 8th 2020.

[8] Nelson Dias, Sahsil Enríquez, and Simone Júlio (Eds.). 2019. *The Participatory Budgeting World Atlas*. Epopee Records: Officinal Coordination.

[9] Brandon Fain, Ashish Goel, and Kamesh Munagala. 2016. The Core of the Participatory Budgeting Problem. In *Proceedings of the 12th International Workshop on Internet and Network Economics (WINE)*. 384–399.

[10] Rupert Freeman, David M. Pennock, Dominik Peters, and Jennifer Wortman Vaughan. 2019. Truthful Aggregation of Budget Proposals. In *Proceedings of the 20th ACM Conference on Electronic Commerce (ACM-EC)*. 751–752.

[11] Rupert Freeman, Seyed Majid Zahedi, and Vincent Conitzer. 2017. Fair and Efficient Social Choice in Dynamic Settings. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*. 4580–4587.

[12] Rupert Freeman, Seyed Majid Zahedi, Vincent Conitzer, and Benjamin C. Lee. 2018. Dynamic Proportional Sharing: A Game-Theoretic Approach. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 2, 1 (2018), 3:1–3:36.

[13] Corrado. Gini. 1912. *Variabilità e Mutabilità: Contributo allo Studio delle Distribuzioni e Delle Relazioni Statistiche*. P. Cuppini.

[14] Ashish Goel, Anilesh K. Krishnaswamy, Sukolsak Sakshuwong, and Tanja Aitamurto. 2019. Knapsack Voting: Voting Mechanisms for Participatory Budgeting. *ACM Transaction on Economics and Computation* (2019), 8:1–8:27.

[15] D. Ellis Hershkowitz, Anson Kahng, Dominik Peters, and Ariel D. Procaccia. 2021. District-Fair Participatory Budgeting. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*. To appear.

[16] Pallavi Jain, Krzysztof Sornat, and Nimrod Talmon. 2020. Participatory Budgeting with Project Interactions. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*. 386–392.

[17] Martin Lackner. 2020. Perpetual Voting: Fairness in Long-Term Decision Making. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*. 2103–2110.

[18] Annick Laruelle. 2021. Voting to Select Projects in Participatory Budgeting. *European Journal of Operational Research* 288, 2 (2021), 598–604.

[19] Tyler Lu and Craig Boutilier. 2011. Budgeted Social Choice: From Consensus to Personalized Decision Making. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*. 280–286.

[20] New York City Council. 2020. Participatory Budgeting. https://council.nyc.gov/pb/. Last accessed on October the 8th 2020.

[21] Dominik Peters, Grzegorz Pierczyński, and Piotr Skowron. 2020. Proportional Participatory Budgeting with Cardinal Utilities. *arXiv preprint arXiv:2008.13276* (2020).

[22] Simon Rey, Ulle Endriss, and Ronald de Haan. 2020. Designing Participatory Budgeting Mechanisms Grounded in Judgment Aggregation. In *Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning (KR)*. 692–702.

[23] Anwar Shah (Ed.). 2007. *Participatory Budgeting*. The World Bank, Washington, DC.

[24] Ehud Shapiro and Nimrod Talmon. 2017. A Participatory Democratic Budgeting Algorithm. *arXiv preprint arXiv:1709.05839* (2017).

[25] Piotr Skowron, Arkadii Slinko, Stanisław Szufa, and Nimrod Talmon. 2020. Participatory Budgeting with Cumulative Votes. *arXiv preprint arXiv:2009.02690* (2020).

[26] Nimrod Talmon and Piotr Faliszewski. 2019. A framework for approval-based budgeting methods. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*. 2181–2188.