

Value Alignment in Participatory Budgeting

Marc Serramia
City, University of London
London, United Kingdom

marc.serramia-amoros@city.ac.uk

Juan A. Rodriguez-Aguilar
Artificial Intelligence Research Institute (IIIA-CSIC)
Barcelona, Spain
jar@iiia.csic.es

Maite Lopez-Sanchez
University of Barcelona
Barcelona, Spain
maite_lopez@ub.edu

Stefano Moretti
LAMSADE, Université Paris-Dauphine
Paris, France
stefano.moretti@lamsade.dauphine.fr

ABSTRACT

Participatory budgeting empowers citizens to take an active role in shaping their government's policies by influencing the allocation of a limited budget. In this process, citizens file various proposals and then collectively decide which ones should receive funding through a voting system. While participatory budgets have garnered significant attention in research and practice, one aspect so far overlooked is the ethical dimension of the proposals. Thus, beyond just gauging citizen preferences, we propose also to consider how these initiatives align with the government's core values. Specifically, we apply optimisation techniques to solve a multi-criteria decision problem that considers both citizen support and value alignment when choosing the proposals to fund. We illustrate our method in two real case studies and analyse how we can combine both criteria in an egalitarian way that does not necessarily compromise the will of citizens and may encourage governments to broaden the objectives and increase the allocated budget.

KEYWORDS

Participatory budgets; Value alignment; Ethics

ACM Reference Format:

Marc Serramia, Maite Lopez-Sanchez, Juan A. Rodriguez-Aguilar, and Stefano Moretti. 2024. Value Alignment in Participatory Budgeting. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024)*, Auckland, New Zealand, May 6 – 10, 2024, IFAAMAS, 9 pages.

1 INTRODUCTION

In democratic countries, the prevailing model is representative democracy, wherein citizens limit their incidence to political life to the election of their representatives. Alternatively, participatory democracy pursues that citizens meaningfully contribute to political decision-making[39]. In this context, participatory budgeting constitutes an innovative mechanism that has lately gathered a lot of attention. This is so because it empowers citizens to take an active role in shaping their government's policies. Through this process, a

government allocates a budget to be spent on those proposals that garner –through voting– the most citizen support.

Since their onset in the 80s in Brazil [43], governments have widely adopted these processes¹. However, as reported by the UN [34], both participation normally ranges between 1% and 15% of the voting population and allocated budget represents between 1% and 10% of the total executed budget. Moreover, although research in social choice has shown interest in participatory budgeting [1, 3, 13, 20], so far it has overlooked the ethical dimension of the proposals. Against this background, in this paper we propose to consider, in addition to citizen support, the alignment of the proposals with the moral (strategic) values of the government. This serves a dual purpose: it helps compensate for potential legitimacy limitations in citizen participation and instils confidence in governments to expand the scope of participatory budgeting to encompass a wider range of decisions. Overall, we argue it poses a suitable/balanced combination of representative and participatory democratic models.

Our proposal goes in line with the European Union technical report highlighting the importance of value alignment in policy-making both as a means to make decisions and also to explain them [27]. For example, a proposal (or policy) that improves public transportation infrastructures promotes (i.e., it is aligned with) the value of environmentalism, and this can also be used to justify an investment. Obviously, different political parties have different preferences over values, but as they have been elected by the population as their representatives, we argue it is legitimate to also consider (together with the citizens' direct support) their value preferences when choosing the proposals to fund. Importantly, the two real-world participatory budgets we have studied so far show that incorporating value alignment into the decision-making process does not necessarily compromise the will of the people. Instead, it enhances the synergy between representative and participatory democracy, ensuring that the voice of the citizens resonates with the overarching goals of the government. Hopefully, our proposed method will encourage the governments to increase the budgets allocated to participatory processes.

Within value alignment literature [12, 37], value inference constitutes an open problem. Previous work in this area range from the proposal of a general pipeline [23], to contextual value detection



This work is licensed under a Creative Commons Attribution International 4.0 License.

¹According to the participatory budgeting atlas [10], there were 10081 participatory budget processes in the world (all continents) in 2019. The pandemic reduced the number down to 4032 in 2020. Most processes are set by local governments and large cities. Some examples of cities that have adopted these processes are New York[7], Santiago [9], Kakogawa[18], Madrid [28], Barcelona [2], Paris [35], or Cape Town[32].

[24] or value system aggregation [22]. In our particular case, we need to elicit the political or strategic values of the representatives in the government. For that, we envision three alternative ways. First, it could be done directly, through direct surveys, as done in the World Value Survey [52]. Second, we could use different indirect value elicitation mechanisms, such as using related indicators (e.g., Montes et Sierra [30] suggest the usage of the Gini index to model inequality) or by inferring them by applying Natural Language Processing techniques in text (as done by Liscio et al. [24]). In this paper, we opt for a third option to define the governmental value system (i.e., the values and their relative preferences). In particular, we resort to the government action plans (be it a municipal action plan or a national action plan) to manually infer the values from the chapters within, and then, map their associated expenditure to value preferences.

This paper advances the state of the art in participatory budgets by providing:

- A formal definition of a new approach to budgeting, considering value alignment as the main driver for the decision.
- A hybrid approach that combines both maximising citizen satisfaction and value-alignment.
- Two new participatory budgeting solutions inspired by classic bargaining game solutions.
- An experimental evaluation to show the effects of our hybrid approach to real-world participatory budget data.

As for the paper's structure, Section 2 introduces participatory budgets and their solutions. In Section 3 we introduce a first approach to solving participatory budgets, which aims at maximising citizen satisfaction. Section 4 is devoted to introducing the value-aligned approach. We describe the mixture of the two aforementioned approaches in Section 5. Section 6 then details how to solve our proposed mixed approach, and in Section 7 we test it with real data. Finally, Section 8 provides conclusions and future work.

2 PARTICIPATORY BUDGETS

Within a participatory budgeting process, citizens collectively decide which proposals should receive funding through a voting system. The social choice literature has explored a range of such voting methods[20]: some permit unlimited votes (e.g., approval voting), others limit the number of votes (e.g., k-approval voting or knapsack voting[13]), and still others use ordered ballots (where participants express their preferences by ranking proposals). Moreover, some researchers include fairness criteria when allocating budgets.

In this paper though, we encompass most contemporary (real-world) participatory budgeting processes by abstracting from the actual voting process and by simply evaluating proposals in terms of the votes they gather and their cost. Thus, from a set of proposals P , we formalise $vot : P \rightarrow \mathbb{N}$ such that $vot(p)$ counts the number of votes received by a proposal $p \in P$. Similarly, we define a cost function $cost : P \rightarrow \mathbb{R}^+$, such that $cost(p)$ is the cost of implementing $p \in P$. Moreover, we abuse notation to compute the overall cost of a set of proposals $S \subseteq P$ as $cost(S) = \sum_{p \in S} cost(p)$.

When deciding which proposals to fund, we must also consider the relations they may have. Hence, proposals may be exclusive if they cannot be executed simultaneously (e.g. pedestrianising a street and improving its traffic lanes) or may have a generalisation

relationship (e.g. pedestrianising the city centre is more general than pedestrianising a single street) that implies redundancy. On the one hand, we define proposal exclusivity as an irreflexive, symmetric, intransitive binary relation \mathcal{R}_x , such that $(p_1, p_2) \in \mathcal{R}_x$ means p_1 and p_2 are exclusive. On the other hand, generalisation is defined as an irreflexive, antisymmetric, and transitive binary relation \mathcal{R}_g , such that $(p_1, p_2) \in \mathcal{R}_g$ means p_1 is generalised by p_2 .

Overall, we need to distribute a given budget $b \in \mathbb{R}^+$, thus, the set of chosen proposals must be in budget and should not include exclusivity or generalisation relationships. Formally:

DEF. 1. *A feasible solution of a participatory budgeting process is any subset $Sol \subseteq P$ satisfying: i) $cost(Sol) \leq b$; ii) $\forall p_1, p_2 \in Sol, (p_1, p_2) \notin \mathcal{R}_x$; iii) $\forall p_1, p_2 \in Sol, (p_1, p_2) \notin \mathcal{R}_g$.*

From the set of all feasible solutions, we aim at finding those that maximise some criteria. Subsequent sections consider citizen satisfaction and value alignment.

3 CITIZEN SATISFACTION

In most real-world participatory budgeting processes, citizens initially cast their votes for preferred proposals so that proposals can be subsequently selected by following a descending order of gathered votes ($vot(p)$) as long as the given budget b is not exceeded [11]. This process can be naturally seen as a proxy for maximising citizen satisfaction. However, considering proposals in this order does not guarantee that the selected proposals actually maximize the total number of accumulated votes [39].

Alternatively, we consider the problem of finding the combination of proposals that maximize citizen satisfaction. The solution to this problem can then be defined as follows:

DEF. 2. *Given a participatory budget with proposals P , the maximum citizen satisfaction solution $P_{sat}^* \subseteq P$ is such that:*

$$P_{sat}^* = \underset{P' \subseteq P}{\operatorname{argmax}} \sum_{p_i \in P'} vot(p_i) \quad (1)$$

subject to:

$$\sum_{p_i \in P_{sat}^*} cost(p_i) \leq b \quad (2)$$

$$\forall p_i, p_j \in P_{sat}^*, (p_i, p_j) \notin \mathcal{R}_x \quad (3)$$

$$\forall p_i, p_j \in P_{sat}^*, (p_i, p_j) \notin \mathcal{R}_g \quad (4)$$

where $vot(p_i)$ stands for the number of votes of proposal $p_i \in P$, $cost(p_i)$ represents its cost, b is the allocated budget, and \mathcal{R}_x and \mathcal{R}_g stand for exclusive and generalisation relations.

To find this solution, we encode a binary integer program (BIP) which can be solved with state-of-the-art solvers such as CPLEX[16] or Gurobi[14]. This encoding assigns a binary variable $x_i \in \{0, 1\}$ to each proposal $p_i \in P$ to denote whether the proposal is selected (1) or not (0) and uses the target function:

$$\operatorname{Maximise} \sum_{p_i \in P} x_i \cdot vot(p_i) \quad (5)$$

subject to the constraints relating to Eq. 2, 3, 4 respectively:

$$\sum_{p_i \in P} x_i \cdot cost(p_i) \leq b \quad (6)$$

$$x_i + x_j \leq 1 \quad \forall (p_i, p_j) \in \mathcal{R}_x \quad (7)$$

$$x_i + x_j \leq 1 \quad \forall (p_i, p_j) \in \mathcal{R}_g \quad (8)$$

4 VALUE ALIGNMENT

As previously argued in the introduction, we propose to consider the ethical dimension of the proposals to choose in participatory budgeting processes. Within ethics, moral values express the moral objectives *worth striving for* [49, p.72]. Examples of human values include fairness, respect, freedom, security, and prosperity [4]. Additionally, both Sociology and Psychology have long studied human values² as well as their relative importance across individuals and societies. As governments do also have value (strategic) preferences, we can include them in the proposal decision process by borrowing the concept of value system from [41].

DEF. 3. A value system is a structure $\langle V, \geq \rangle$, where V is a set of values and \geq is a ranking over them (that is a total preorder or, in other words, a total, reflexive, and transitive binary relation).

It is worth noticing that proposals can demote or promote values. For example, a proposal to pedestrianise a city centre promotes the value of environmentalism but demotes the value of mobility. Furthermore, a proposal $p \in P$ can relate to the values in a value system $\langle V, \geq \rangle$ with different degrees. Thus, a proposal to pedestrianise the whole city centre will promote environmentalism more than the one just considering a single street. Next, we define the proposal promotion function that captures this relationship:

DEF. 4. Given a set of proposals P and a value system $\langle V, \geq \rangle$, the promotion function $\text{prom} : P \times V \rightarrow [-1, 1]$ defines the degree of promotion or demotion between each proposal and each value. Negative degrees $[-1, 0)$ represent demotion, while positive ones $(0, 1]$ represent promotion, and 0 denotes the lack of relation.

In general, a participatory budget solution contains several proposals promoting and/or demoting (with different degrees) several values, which in turn are more or less preferred than other values. Hence, in order to take all these aspects into account, we define the value alignment score by considering three dimensions: the value preferences, the proposals' relation to values, and the multiple proposals in a solution. We first look at the value preferences and consider a value relevance function in line with [42]

DEF. 5. Given a value system $\langle V, \geq \rangle$, a value relevance function is a function $r : V \rightarrow \mathbb{R}^+$ satisfying that $\forall v, v' \in V \quad v \geq v' \Leftrightarrow r(v) \geq r(v')$.

Thus, if v is more preferred than v' , then its relevance $r(v)$ must be greater than $r(v')$ (and if v and v' are indifferently preferred, then $r(v) = r(v')$). Note that the value relevance function r is very general, as it just requires that the more preferred a value is, the higher its relevance. Therefore, decision-makers may opt for alternative computations. For simplicity, here we compute the relevance of a value v as the sum of relevances of its less preferred values (with the least preferred values having a relevance of 1). Hence, given a value system $\langle V, \geq \rangle$, we define the relevance of a value $v \in V$ as:

$$r(v) = 1 + \sum_{v' > v} r(v') \quad (9)$$

²Different value models have been proposed and compared [4], rendering Schwartz's value model [38] as the most comprehensive one [15].

Second, as each proposal might be related to more than one value, we assess each proposal's value alignment by combining both the degree of promotion/demotion to each value and the value's relevance.

DEF. 6. Given a proposal $p \in P$, a value system $\langle V, \geq \rangle$, a relevance function r , and a promotion function prom , the value alignment of p is:

$$\text{al}(p) = \sum_{v \in V} r(v) \cdot \text{prom}(p, v) \quad (10)$$

Finally, a feasible solution will usually contain more than one proposal. Hence, the overall value alignment of a set of proposals is the sum of the individual value alignment of its proposals³:

DEF. 7. Given a feasible solution $\text{Sol} \subseteq P$ and a relevance function r , the value alignment of Sol is:

$$\text{al}(\text{Sol}) = \sum_{p \in \text{Sol}} \text{al}(p) = \sum_{p \in \text{Sol}} \sum_{v \in V} r(v) \cdot \text{prom}(p, v) \quad (11)$$

Similar to the citizen satisfaction problem introduced in Section 3, here we consider the problem of finding the subset of proposals that maximise value alignment.

DEF. 8. Given a participatory budget with proposals P , the maximum value alignment solution $P_{\text{al}}^* \subseteq P$ is such that:

$$P_{\text{al}}^* = \underset{P' \subseteq P}{\text{argmax}} \sum_{p_i \in P'} \text{al}(p_i) \quad (12)$$

Subject to the constraints in Eqs. 2, 3, and 4.

As before, we can solve this problem using optimisation techniques. Thus, we encode it as a BIP by assigning again a binary variable $x_i \in \{0, 1\}$ to each proposal $p_i \in P$ and by considering the following target function:

$$\text{Maximise} \sum_{p_i \in P} x_i \cdot \text{al}(p_i) \quad (13)$$

subject to the constraints in Eqs. 6, 7, and 8.

5 LEVERAGING VALUE ALIGNMENT AND CITIZEN SATISFACTION

In order to create synergies between representative and participatory democracy, we now focus on combining value alignment and citizen satisfaction into the decision-making process of participatory budgeting. Thus, we aim to ensure that the voice (votes) of the citizens resonates with the strategic values of the government. We formalise this as a bi-objective problem, the so-called value-aligned participatory budgeting problem (VAPBP), and solve it utilizing optimisation techniques in Section 6.

The value-aligned participatory budgeting problem (VAPBP) aims at maximising both citizen satisfaction and value alignment in participatory budgeting processes. Thus, by considering the number of gathered votes as a proxy for citizen satisfaction, it can be formalised as a bi-objective problem:

³We assume that the proposals in P are as all-encompassing as possible, i.e., that there are no two proposals in P that could be combined into one meaningfully. This is because dividing proposals into smaller ones could affect their value alignment. It is worth mentioning though, that good democratic practices recommend to check proposals for possible combinations before moving into the voting phase.

DEF. 9 (VAPBP). Given a set of proposals P with exclusivity relations in \mathcal{R}_x and generalisation relations in \mathcal{R}_g , an available budget b , a citizen satisfaction function $\text{vot}(p)$, and an ethical alignment function $\text{al}(p)$, the value-aligned participatory budgeting problem (VAPBP) is that of finding the subset of proposals $P^* \subseteq P$ such that:

$$P^* = \underset{P' \subseteq P}{\operatorname{argmax}} \left(\sum_{p_i \in P'} \text{vot}(p_i), \sum_{p_i \in P'} \text{al}(p_i) \right)$$

subject to the constraints in Eqs. 2, 3, and 4.

A feasible solution for a VAPBP must satisfy the conditions in Def. 1. Henceforth, we refer to the set $SS \subseteq \mathcal{P}(P)$ containing all the feasible solutions of a VAPBP as its *solution space*.

The VAPBP is a bi-objective **combinatorial** optimisation problem. In multi-objective optimization, the goodness of a solution is determined by **dominance**. Informally, a solution x_1 dominates another solution x_2 if: (i) solution x_1 is no worse than x_2 in all objectives, and (ii) solution x_1 is strictly better than x_2 in at least one objective. In our case, we can formalise dominance as follows.

DEF. 10 (VAPBP SOLUTION DOMINANCE). A solution $Sol \in SS$ dominates another solution $Sol' \in SS$ if:

- $\text{vot}(Sol) \geq \text{vot}(Sol')$ and $\text{al}(Sol) \geq \text{al}(Sol')$.
- $\text{vot}(Sol) > \text{vot}(Sol')$ or $\text{al}(Sol) > \text{al}(Sol')$.

Hence, our aim is to find the feasible solutions in the solution space that are not dominated by others. We say that these solutions satisfy *strong Pareto optimality*.

DEF. 11 (STRONG PARETO OPTIMALITY). A solution $Sol \in SS$ is strong Pareto optimal if there is no other solution $\nexists Sol' \in SS$ dominating it.

Notice that all strong Pareto-optimal solutions form what is usually called the Pareto front, which we formalise hereunder.

DEF. 12 (PARETO FRONT OF THE VAPBP). Given a VAPBP with proposals P , exclusivity relations \mathcal{R}_x , generalisation relations \mathcal{R}_g , budget b , citizen satisfaction utility vot , and alignment utility al , the Pareto front of the problem is $PF \subseteq SS$, such that $\forall P^* \in PF, \nexists P' \in SS$ with P' dominating P^* .

Regarding participatory budgeting, ensuring that the selected solutions are Pareto-optimal (they lay within the Pareto front) is crucial. Neither the citizens nor the government are interested in spending money on a sub-optimal selection of proposals. Without strong Pareto-optimality, there might be another set of proposals that is better concerning one of the two criteria (either citizen satisfaction or ethical alignment) while being at least as good for the other criterion.

Hence, any solution within the Pareto front (PF) is a potential candidate to be selected in a participatory budgeting process. However, we are interested in those that provide a good compromise between both criteria.

We can regard a VAPBP as a bargaining game [31]. In short, a bargaining game consists of two or more parties bargaining over some goods (e.g., a buyer and seller that must agree on the price of some good they want to trade). As to participatory budgeting, a VAPBP is a bargaining game between the citizens and the government, who must decide on the proposals to approve. Citizens

want some proposals to be approved (and cast their votes to favour them), while the government (representing non-participating citizens) defines their preferences over proposals through alignment. The bargaining between these two parties will lead to a solution (a selection of proposals) between both parties' preferences.

The literature has explored several rules for finding solutions in bargaining games (e.g., the survey by Thomson et al. [48]). However, these rules typically assume that the feasible set is an infinite, convex, and compact subset of the Euclidean space. As to participatory budgeting, our feasible set, the solution space SS , is finite. Hence, while we can redefine the rules in [48], their properties may differ from those discussed there. Here we focus on two of the classic bargaining game rules: the Nash product and Kalai-Smordinsky rules.

First, the Nash product is commonly considered in bargaining games and other areas like multi-agent resource allocation [5]. Given a feasible solution for the VAPBP, the Nash product obtains the value of the solution as the product of citizen satisfaction and ethical alignment. The Nash rule would select as solutions the feasible solutions with maximum Nash product. We formally define the Nash product rule as follows.

DEF. 13 (NASH PRODUCT RULE FOR THE VAPBP). Consider a VAPBP with a solutions space SS . Let $\text{prod}(P')$ be the product of the citizen satisfaction and alignment utilities for a set of proposals $P' \in SS$:

$$\text{prod}(sol) = \sum_{p_i \in P'} \text{vot}(p_i) \cdot \sum_{p_i \in P'} \text{al}(p_i)$$

The Nash value-aligned participatory budget rule selects the following feasible solution:

$$P_{Nash}^* = \underset{P' \in SS}{\operatorname{argmax}} \text{prod}(P')$$

The Nash product rule satisfies Pareto-optimality. Note that given two sets of proposals, if one of them is better in terms of either (or both) criteria, its Nash product will be greater⁴. Besides Pareto-optimality, Mariotti et al. [29] study other properties of the Nash product rule which we inherit. In particular, their axiomatisation⁵ considers the following additional properties: Covariance with scale transformations, Symmetry, and Independence of irrelevant alternatives. Without entering into formal details, we describe what these properties mean in participatory budgeting terms. First, covariance with scale transformations ensures that the solution remains the same independently of the linear transformations applied to criteria utilities (e.g., the solution will be maintained independently of measuring citizen satisfaction and ethical alignment in $[0,1]$ or $[0,100]$ and multiply all utilities by 100). Second, Symmetry ensures that if the citizens and government utilities are exchanged, then the solution will remain the same. Third, Independence of irrelevant alternatives guarantees that the solution is maintained if we add “irrelevant” proposals (e.g., a proposal that is more expensive and worse in terms of citizen satisfaction and alignment than any other proposal).

⁴Assuming there is at least one proposal with non-zero utilities for all criteria considered (citizen satisfaction and ethical alignment in our case). This is a common assumption in bargaining games.

⁵Notice that in their axiomatisation they consider weak Pareto optimality because that is enough to fully axiomatise the Nash rule, but it also satisfies strong Pareto optimality as discussed above.

Second, say that we represent each set of proposals in SS in the two-dimensional Euclidean space where the x-axis represents citizen satisfaction and the y-axis represents ethical alignment. Then, the Kalai-Smordinsky (abbreviated KS) rule would select the set of proposals that is closer to the line drawn between the status quo point $(0,0)$ and the utopia point (represented by the maximum value of citizen satisfaction and the maximum alignment achieved by any solution), hence favouring solutions closer to the utopia point. Formally, a solution $P' \in SS$, would be represented by point $(\sum_{p \in P'} \text{vot}(p), \sum_{p \in P'} \text{al}(p))$. We obtain the utopia point as $(\text{vot}_{\max}, \text{al}_{\max})$, where $\text{vot}_{\max} = \arg\max_{p' \in SS} \sum_{p \in P'} \text{vot}(p)$, and $\text{al}_{\max} = \arg\max_{p' \in SS} \sum_{p \in P'} \text{al}(p)$. Therefore, we obtain the distance between a feasible solution $P' \in SS$ and the line going through the status quo and ideal points as:

$$\text{dis}(P') = \frac{|\text{al}_{\max} \sum_{p \in P'} \text{vot}(p) - \text{vot}_{\max} \sum_{p \in P'} \text{al}(p)|}{\sqrt{\text{vot}_{\max}^2 + \text{al}_{\max}^2}}$$

Thus, we can formalise the KS rule as follows.:

DEF. 14 (KALAI-SMORDINSKY RULE FOR THE VAPBP). *Given a VAPBP with solution space SS , the Kalai-Smordinsky (abbreviated KS) rule selects the set of proposals $P_{KS}^* = \arg\min_{P' \in SS} \text{dis}(P')$.*

As for the properties of the KS rule, Nagahisa et al. [17] provide an axiomatisation for cases with a finite feasible set. They identify the following properties: Continuity, Independence, Monotonicity, Symmetry, and Invariance. Despite some changes in properties' names, Nagahisa's axiomatisation is fairly similar to Mariotti's on the classic Nash solution, as described above. Thus, Symmetry is the same for both authors. Nagahisa et al. [17] call Invariance what Mariotti calls Covariance with scale transformations. Continuity for the KS solution replaces Weak Pareto Optimality for the Nash solution. Continuity ensures the stability of a solution so that a small change in its set of proposals will not lead to a large change in the solution.

Independence of irrelevant alternatives for a Nash solution is traded for two other properties: a weaker version of it (referred to as "Independence" by Nagahisa et al. [17]) and, most notably, Monotonicity. Informally, Monotonicity ensures that solutions improve if the set of *available* proposals improves. For instance, this is the case when adding low-cost and highly-preferred proposals (by either citizens or the government, or both) to the set of available proposals. We regard Monotonicity as a desirable property for a participatory budgeting process whose proposals change dynamically. For instance, if some proposal is expanded during the voting phase or if new proposals need to be considered due to a sudden new necessity of the city/region.

Importantly, note that while the KS solution for an infinite feasible set satisfies strong Pareto-optimality [48], Nagahisa et al. [17] show that this is not the case when the feasible set is finite. This means that if there are two solutions, P' and P'' , with different citizen satisfaction but equal alignment, then the KS solution might be the one with less citizen satisfaction. Since, as discussed above, strong Pareto-optimality is important for participatory budgeting, we tweak the KS solution to ensure it satisfies this property. To do so, given a VAPBP, instead of applying the KS rule to the whole

solution space SS , we apply it to the Pareto front of the problem. We call this the Pareto Kalai-Smordinsky (abbreviated Pareto KS) rule, which we formalise as follows.

DEF. 15 (PARETO KALAI-SMORDINSKY RULE FOR THE VAPBP). *Given a VAPBP with solution space SS and Pareto front PF , the Pareto Kalai-Smordinsky rule selects the set of proposals $P_{PKS}^* = \arg\min_{P' \in PF} \text{dis}(P')$.*

Our Pareto KS rule satisfies strong Pareto-optimality trivially and satisfies the properties of the classic KS rule if we restrict them to the Pareto front. Nonetheless, we can only guarantee Pareto-optimality from the point of view of the whole solution space.

As a final comment, note that when the solution space is finite, the proposed rules may not select a unique set of proposals. Different solutions may have the same Nash product value. Or there might be two solutions whose distance to the line used by the Pareto KS rule is equal. Although unusual, those cases would benefit from defining some criterion to select a single solution (e.g., the one with greater citizen satisfaction).

Having formally defined the rules to select solutions out of the VAPBP Pareto front, next section details how to compute solutions.

6 SOLVING THE VAPBP

Note that the VAPBP is an instance of the well-known multi-objective knapsack problem (MOKP)[26] with a single knapsack (whose capacity corresponds to the allocated budget).

A classic multi-objective optimisation method is the weighted sum method, which scalarises a set of objectives into a single objective by adding each objective pre-multiplied by a user-supplied weight. The weighted sum method allows solving a MOKP using off-the-shelf LP libraries. The disadvantages of this method are: (i) it is difficult to set the weight vectors to obtain a Pareto-optimal solution in a desired region in the objective space, and (ii) it cannot find certain Pareto-optimal solutions in the case of a non-convex objective space. This second disadvantage is important since our solution space is not convex (as shown in the experiments in Sec. 7). As discussed in the previous section, the importance of the Pareto-optimal property means we only consider Pareto-optimal solutions (the Nash and Pareto KS solutions). Since these will always be on the Pareto front, instead of the weighted sum optimisation approach, we advocate for computing the whole Pareto front and selecting our solutions out of it. As discussed by Visée et al. [51], this is a complex task as the so-called non-supported efficient solutions (the ones which "break the convexity of the set") are not straightforward to compute. Fortunately, there are several exact methods (complete algorithms) in the literature to find the whole Pareto front of the MOKP as surveyed in [26]. For example, researchers have defined highly efficient, exact algorithms, as discussed in [53] or [47].

From the Pareto front, which in our case is a finite set, we can assess the Nash product of the points on the front to select the one maximising it as the Nash solution. On the other hand, we can also employ the Pareto front to compute the Pareto KS solution by applying the rule described in Def. 15.

7 EXPERIMENTAL EVALUATION

Here we illustrate how our proposal succeeds in leveraging value alignment and citizen satisfaction in the real-world participatory budgeting processes of Barcelona[6] and Warsaw[33]. We chose them due to data availability and extensive documentation.

We provide a github repository [40] with the supplementary material and code of our experiments and the data in an extended version of Pabulib format [45] that includes value relevance.

7.1 Barcelona

The participatory budget of Barcelona 2020-2023[6] involved about 40,000 people and initially distributed an overall budget of 30M€ among its 10 city districts (e.g., Les Corts was allocated 2M€, Sants 3.6M€, or Nou Barris 3.6M€). Subsequently, the process continued with 5 different phases that mostly consisted of: citizens making, debating about, and prioritising proposals; city technicians specifying how much they would cost and their categories (i.e., areas such as education or ecology); and citizens voting for the proposals. The voting method was knapsack voting [13] and votes were used to rank proposals for each district. Proposals were finally approved from top to bottom until the district’s budget was spent (skipping those proposals whose cost exceeded the remaining budget).

Alternatively, as we aim at considering both votes and value-alignment in the participatory budget, we need to infer the values involved. As previously discussed in the introduction, different methods –such as surveys or indirect indicators– can be applied. Here we resort to the municipal action plan –where the municipal expenditure is detailed for different areas/chapters– and assume that the more money spent on an area, the more relevant the moral value related to that area. For instance, we assume that a municipality that invests the most money in schools is a government that values education the most. The first two columns in Table 1 show Barcelona’s spending in categories related to values⁶. Value relevance $r(v)$ in the third column is then computed as the value spending over the total spending, which is 2065.6 million euros (M€) if just considering categories related to values.

Next, in order to avoid speculation, we simply define promotion as $prom(p, v) = 1$ if proposal p is related to value v (i.e., categorised in the related area) or $prom(p, v) = 0$ otherwise. Then, from $r(v)$ and $prom(p, v)$ we can apply Eq. 10 to compute the alignment $al(p)$ of all proposals. Table 2 lists some example proposals. In particular, the alignment for p_6 is the sum of the relevance of all the values it promotes (Equality and Ecology, in this case).

Finally, considering the citizen satisfaction function $vot(p)$ and the alignment function $al(p)$ as described above, we are ready to formalise each VAPBP according to Def. 9 and solve it by applying social choice functions over the Pareto front as described in Section 5. As for the algorithm to build the Pareto front we use the one by Tamby et al.⁷ [47]. In particular, we employed the implementation provided in Jump the package by Lubin et al. [25]. Figure 1 shows

Value	Spending (M€)	Relevance $r(v)$
Equality	392.4	0.18997
Social Welfare	403.2	0.19520
Housing	163.5	0.07915
Healthcare	23.8	0.01152
Education	161.8	0.07833
Culture	164.2	0.07949
Sport	28	0.01356
Ecology	246.7	0.11943
Security	373.8	0.18096
Economic growth	108.2	0.05238

Table 1: Barcelona’s spending [8] in areas/related values.

# p	Proposal desc.	Promoted Values	al	$cost$	vot
p_1	Improve E.I. Park	Ecology	0.11943	0.85	3351
p_2	Reform spt camp	Sport	0.01356	1.6	2890
p_3	Pedestrianise o.S.	S.Welfare	0.1952	0.2	2483
p_6	Renew veget. & children’s area	Ecology, Equality	0.3094	0.535	2132
p_7	Superblock Sants	Ecology	0.11943	0.55	2004
p_8	Cit. Usage F.G.C.	S.Welfare	0.19520	0.3	1927
p_9	Led screens spt	Sport	0.01356	0.07	1824
p_{12}	Pedestrianise C.s.	S.Welfare	0.1952	0.4	1626
p_{17}	Changing rooms	Sport	0.01356	0.275	1083
p_{18}	Redevelop Co.Sq.	S.Welfare	0.19520	1.3	790
p_{19}	Reform Casa. St.	S.Welfare	0.19520	0.588	518

Table 2: The related values, alignment $al(p)$, cost (in M€), and votes of a subset of proposals¹⁰ from Barcelona’s Sants district. Proposals are numbered in decreasing order of $vot(p)$.

the results for a sample of 3 out of the 10 districts in Barcelona⁸. For all plots, the x-axis represents the citizen satisfaction achieved by each solution and the y-axis represents its value alignment. The blue dots represent the solutions laying in the Pareto front (PF, see Def. 12) of the VAPBP for each district. The red cross (×) represents the Nash solution (P_{Nash}^* , see Def. 13), the pink triangle the Pareto KS solution (P_{PKS}^* , see Def. 15), the green square the solution maximising citizen satisfaction only (P_{sat}^* , see Def. 2), and the black star is the utopia point (i.e., the maximum possible for each criterion). If solutions coincide we paint them as an orange cross (×) and note it in the legend.

We can see in the plots that considering only one criterion may hinder the other one, whereas our proposed solutions present a good compromise between the two. For example, in the case of Les Corts (Figure 1a) maximising citizen satisfaction produces a solution P_{sat}^* with 92,9% of the maximum possible value alignment (i.e., the green square gets 1.63 out of 1.76), whereas maximising value alignment produces a solution P_{al}^* (see Def. 8) that only accounts for around 79.5% of the maximum citizen satisfaction possible (i.e.,

⁶For the sake of simplicity, in a few cases, values were defined as the addition of a couple of areas (e.g., the “Economic growth” value came from “R+D+i” and “commerce, tourism, SMEs” areas, which had spending of 1.6M€ and 106.6M€ respectively).

⁷This is an exact algorithm that is an extension of the classic epsilon-constraint method. Surprisingly, while epsilon-constraint is not exact we have seen it produces the same solutions in almost all cases.

⁸For space constraints we only show the 3 most interesting plots, nonetheless we refer the reader to our supplementary material and code [40] where the rest of the plots can be found.




Figure 1: Results for a sample of 3 out of the 10 processes in Barcelona (see all 10 in supplementary material [40]).

the blue dot on the left gets 6548 out of 8241). This difference would be more noticeable if citizens votes did not favour the government’s values. In contrast, our proposed solutions succeed in finding a suitable balance between representative and participatory democratic models, as they both get closer to the utopia point. In particular, both P_{Nash}^* and P_{PKS}^* (orange cross) achieve 96,8% of maximum citizen satisfaction (7977 out of 8241) and 97,2% of maximum value-alignment (1,71 out of 1.76). Interestingly, the Nash and the Pareto KS solutions coincide in most districts of Barcelona. Conversely, they are not coincident in Sants (Figure 1b and Table 2⁹).

In the case of Sants, the maximum citizen satisfaction solution (P_{sat}^* , green square) selects proposals $\langle p_1, p_3, p_6 - p_{11}, p_{13}, p_{14}, p_{16} \rangle$ (P_{sat}^* just has 86,6% of al). Alternatively, the Nash solution (P_{Nash}^* , red cross) selects proposals $\langle p_1, p_3, p_6, p_8 - p_{14}, p_{16} \rangle$ thus exchanging p_7 for p_{12} (as this has higher alignment) (P_{Nash}^* reaches 98,2% of sat and 90,53% of al). Finally, the Pareto KS solution (P_{PKS}^* , pink triangle) selects proposals $\langle p_3, p_6, p_8 - p_{14}, p_{16}, p_{17}, p_{19} \rangle$ by selecting p_{17} and p_{19} instead of the most voted p_1 to accumulate higher alignment and thus, to get closer to the diagonal between (0, 0) and the ideal solution (black star) (P_{PKS}^* gets 90,1% of sat and 95,2% of al). On the other end of the spectrum, the solution P_{al}^* that maximises alignment is $\langle p_3, p_6 - p_8, p_{10} - p_{14}, p_{16}, p_{19} \rangle$, which includes p_7 instead of p_9 and p_{17} (P_{al}^* though just gets a 85,9% of sat). Additionally, in order to measure the similarity between solutions, we apply the overlap Szymkiewicz–Simpson coefficient[50] by computing the size of their intersection divided by the smaller of the size of the two solutions (i.e., $o(P_i, P_j) = |P_i \cap P_j| / \min(|P_i|, |P_j|)$). Then, we have that $o(P_{al}^*, P_{sat}^*) = o(P_{PKS}^*, P_{sat}^*) = 9/11 = 0.82$, whereas $o(P_{Nash}^*, P_{sat}^*) = 10/11 = 0.91$, so that we can conclude that P_{Nash}^* is more similar to P_{sat}^* than P_{PKS}^* or P_{al}^* ¹⁰.

Considering all 10 districts, the Nash solution selected a set of proposals with an average of 98.3% of the maximum possible citizen satisfaction (standard deviation $\sigma = 1.1\%$, with a minimum of 96.8%) and an average of 96.6% of the maximum possible value alignment (standard deviation $\sigma = 3.33\%$, with a minimum of 90.53%). While

the Pareto KS selected a set of proposals with an average of 96.73% of sat ($\sigma = 3.43\%$, $min=90.08\%$) and an average of 97.27% of al ($\sigma = 2.22\%$, $min=94.27\%$). Moreover, the average similarity with P_{sat}^* is $o(P_{Nash}^*, P_{sat}^*) = 0.916$ ($\sigma = 0.051$) and $o(P_{PKS}^*, P_{sat}^*) = 0.896$ ($\sigma = 0.067$), so they are fairly similar.

Overall, we consider that the obtained solutions P_{Nash}^* and P_{PKS}^* are suitably balanced with respect to both criteria, and this balance compensates for potential legitimacy limitations in citizen participation and encourages governments to increase the budget allocation.

7.2 Warsaw

We consider the data from Warsaw’s 2021¹¹ participatory budget [46]. The budget was distributed among the 18 districts of the city plus a city-wide chapter (to include city-wide proposals). The process in Warsaw is similar to that of Barcelona (see [36]).

We proceed as for Barcelona to elicit values, to compute the corresponding value-alignment, and solve the VAPBP problem for each district. Figure 2 depicts the results for 6 districts considering citizen satisfaction (x-axis) and value alignment (y-axis). Again, blue dots signal the Pareto front. As before, we can see in the plots that considering only one criterion may highly hinder the other one, whereas our proposed Nash and Pareto KS solutions present a good compromise between the two. For example, in the case of the Warsaw city-wide participatory budget (Figure 2e) maximising value alignment produces a solution P_{al}^* that only accounts for around 71.9% of the maximum citizen satisfaction possible (sat), whereas maximising citizen satisfaction produces a solution P_{sat}^* with 81,8% of the maximum possible value alignment (al). On the other hand, our proposed solutions P_{Nash}^* and P_{PKS}^* achieve 95.1% or more in terms of al and sat . Considering all 18 districts (and city-wide participatory budget), the Nash solution P_{Nash}^* selected a set of proposals with an average of 97.32% of sat (standard deviation $\sigma = 1.31\%$, with a minimum of 95.32%) and an average of 96.53% of al ($\sigma = 2.13\%$, $min = 92.61\%$). The Pareto KS solution P_{PKS}^* selected a set of proposals with an average of 96.74% of sat ($\sigma = 1.19\%$, $min = 94.15\%$) and an average of 96.7% of al ($\sigma = 1.09\%$, $min = 94.08\%$). As for average similarity with P_{sat}^* , we obtained $o(P_{Nash}^*, P_{sat}^*) = 0.941$ ($\sigma = 0.019$) and $o(P_{PKS}^*, P_{sat}^*) = 0.931$ ($\sigma = 0.033$), so again, they are similar to the solution with highest citizen satisfaction.

⁹See the complete set of proposals for Sants in the supplementary material [40].

¹⁰The proposals that were actually chosen in Sants by ranking by vot were $\langle p_1 - p_4, p_9 \rangle$, as the first 4 proposals consumed most of the available budget and p_9 just costs 70,000€. Its overlap coefficient with P_{sat}^* is $3/5=0.6$ (so it is not very similar). Obviously, this solution is not plotted in Fig.1b because it does not belong to the Pareto front: it just gets 60,4% (12961) of sat and only 24% of al (0,46118).

¹¹As of writing, the most recent year for which government spending is available [44].

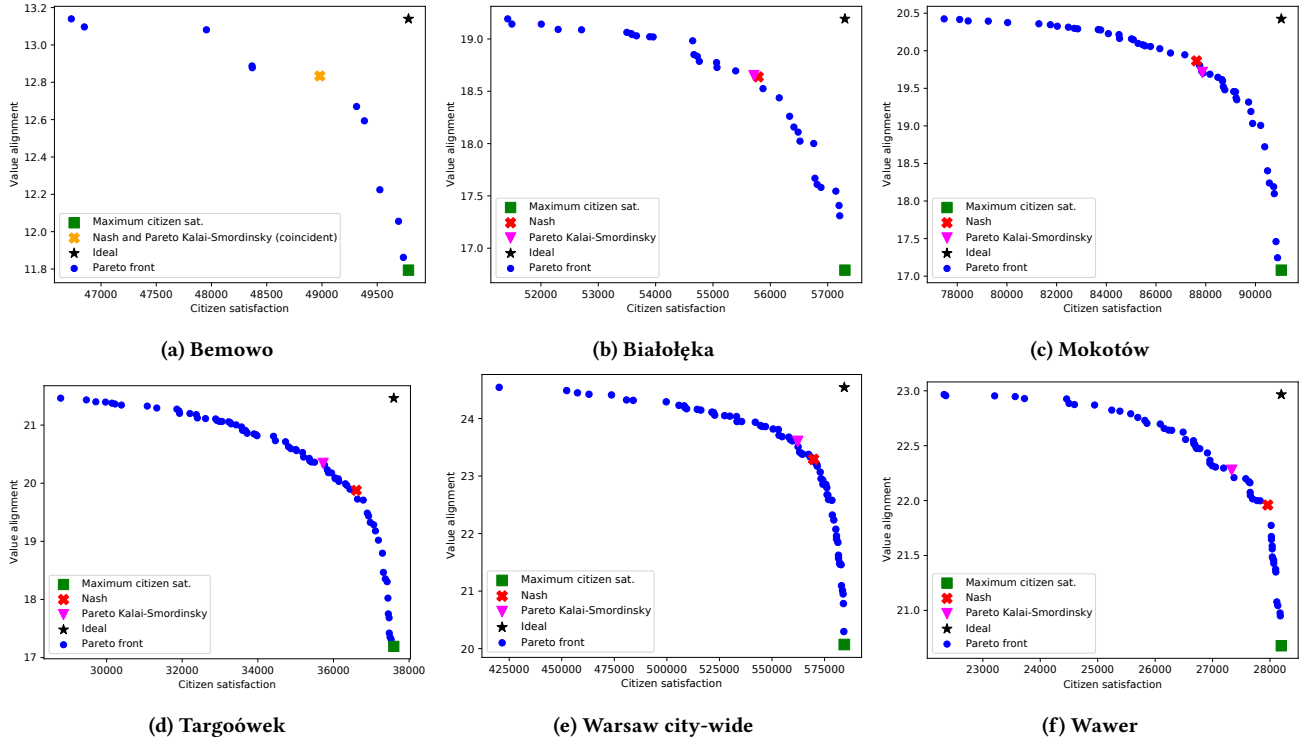


Figure 2: Results for a sample of 6 out of the 19 processes in Warsaw (see all 19 in supplementary material [40]).

7.3 Experiments conclusions

We used real-world data to illustrate how our proposed rules manage to find solutions in line with citizen preferences and government’s values. We can conclude that, on the one hand, if there is a reasonable coincidence between the citizens’ interests and the government’s strategic values, then our approach manages to successfully balance representative and participatory democratic models, that is, to consider the government’s strategic values without compromising the will of the people. In fact, we accomplish almost optimal results in terms of value alignment and citizen satisfaction. On the other hand, our proposal would also become useful to objectively signalling fundamental discrepancies between the citizens’ will and the government representing them. For those cases, the proposed social choice functions would still compute a suitable trade-off that would avoid potential bias.

Although in our experiments the Nash and Pareto KS rules provide similar solutions, it is worth recalling that they hold different properties. Nash satisfies IIA, while Pareto KS satisfies monotonicity. In most cases, we would opt for Nash, but if new proposals can be added during the decision-making process (which is not usual), we would then favour the Pareto KS solution for its monotonicity.

8 CONCLUSIONS AND FUTURE WORK

Participatory budgeting has gathered a lot of attention, but so far, the literature has overlooked the ethical dimension of the proposals. We see the value-aligned participatory budgeting problem as a bi-objective problem of leveraging value alignment and citizen

satisfaction. Our proposal constitutes a balanced combination of representative and participatory democratic models that compensates for potential legitimacy limitations when citizen participation is low, as including values of the representative government indirectly counts for those who did not participate. Our experimental results show that our method provides solutions that are almost optimal for each of the two objectives. Importantly, considering the government’s value preferences in the decision-making process allows to find a consensus solution (between the government and citizens), hence encouraging governments to increase the budget allocation, which has so far been limited.

As future work, we plan to explore alternative value elicitation methods and non-symmetric social choice functions. For instance, it would be interesting to apply a non-symmetric version of the Nash product rule [3, 19]. This would be similar to the work of Laurelle et al. [21] who have used a non-symmetric Nash product in bargaining problems including a voting rule.

ACKNOWLEDGMENTS

We are especially thankful to Eugeni Blas and Ariane Ravier for their help which was essential to conceive our experimental evaluation. Research funded by projects ACISUD (PID2022-136787NB-I00), VAE (TED2021-131295B-C31), VALAWAI (HE-101070930), Fairtrans (PID2021-124361OB-C33), and THEMIS (ANR-20-CE23-0018). Maite Lopez-Sanchez belongs to the WAI research group (University of Barcelona), an associated unit to CSIC through the IIIA.

REFERENCES

- [1] Haris Aziz and Nisarg Shah. 2021. *Participatory Budgeting: Models and Approaches*. Springer International Publishing, Cham, 215–236. https://doi.org/10.1007/978-3-030-54936-7_10
- [2] Decidim Barcelona. 2020. Pressupostos participatius 2020-2023. Accessed August 2023.
- [3] Florian Brandl, Felix Brandt, Matthias Greger, Dominik Peters, Christian Stricker, and Warut Suksompong. 2022. Funding public projects: A case for the Nash product rule. *Journal of Mathematical Economics* 99 (2022), 102585.
- [4] An-Shou Cheng and Kenneth R Fleischmann. 2010. Developing a meta-inventory of human values. *Proceedings of the American Society for Information Science and Technology* 47, 1 (2010), 1–10.
- [5] Yann Chevaleyre, Paul Dunne, Endriss Ulle, Lang Jérôme, Lemaître Michel, Nicolas Maudet, Julian Padget, Steve Phelps, Juan Rodríguez-Aguilar, and Paulo Sousa. 2006. Issues in Multiagent Resource Allocation. *Informatica* 30 (01 2006), pag. 20.
- [6] Barcelona city council. 2020. Participatory budget 2020-2023 process, project data, and results. <https://www.decidim.barcelona/processes/PressupostosParticipatius> Only in catalan and spanish, accessed August 2023.
- [7] New York City Council. 2023. Participatory budgeting PBNYC. <https://council.nyc.gov/pb/>. Accessed Oct 2023.
- [8] Ajuntament de Barcelona. 2023. Public finances of the city of Barcelona in 2023. <https://warszawa.stat.gov.pl/en/information-about-voivodship/capital-of-voivodship/public-finance/>. Accessed August 2023.
- [9] Municipalidad de Santiago de Chile. 2023. Presupuestos Participativos Santiago. <https://www.munistgo.cl/presupuestos/>. Accessed Oct 2023.
- [10] Nelson Dias, Sahsil Enriquez, Rafaela Cardita, Simone Júlio, and Tatiane Serrano. 2021. *Participatory budgeting world atlas*. Epopeia and Oficina, Vila Ruiva - Portugal. Available at: <http://www.participatorybudgetingworks.org/news/view/60/the-latest-edition-of-the-pb-world-atlas-is-now-available>.
- [11] Joan Font, Graham Smith, Carol Galais, and PAU Alarcon. 2018. Cherry-picking participation: Explaining the fate of proposals from participatory processes. *European Journal of Political Research* 57, 3 (2018), 615–636.
- [12] Iason Gabriel. 2020. Artificial intelligence, values, and alignment. *Minds and machines* 30, 3 (2020), 411–437.
- [13] Ashish Goel, Anilesh K. Krishnaswamy, Sukolsak Sakshuwong, and Tanja Aitamurto. 2019. Knapsack Voting for Participatory Budgeting. *ACM Trans. Econ. Comput.* 7, 2, Article 8 (jul 2019), 27 pages. <https://doi.org/10.1145/3340230>
- [14] Gurobi Optimization. 2010. Gurobi. <http://www.gurobi.com/>. Accessed 06/2021.
- [15] Paul HP Hanel, Lukas F Fitzellachner, and Gregory R Maio. 2018. An empirical comparison of human value models. *Frontiers in Psychology* 9 (2018), 1643.
- [16] IBM. 1988. CPLEX. <https://www.ibm.com/analytics/data-science/prescriptive-analytics/cplex-optimizer>. Accessed 06/2021.
- [17] Ryo ichi Nagahisa and Makoto Tanaka. 2002. An axiomatization of the Kalai-Smorodinsky solution when the feasible sets can be finite. *Social Choice and Welfare* 19, 4 (2002), 751–761. <http://www.jstor.org/stable/41106486>
- [18] Kakogawa. 2023. Make our Kakogawa. <https://kakogawa.dycities.jp/>. Accessed Oct 2023.
- [19] Ehud Kalai. 1977. Nonsymmetric Nash solutions and replications of 2-person bargaining. *International Journal of Game Theory* 6 (1977), 129–133.
- [20] Annick Laruelle. 2021. Voting to select projects in participatory budgeting. *European Journal of Operational Research* 288, 2 (2021), 598–604. <https://doi.org/10.1016/j.ejor.2020.05.063>
- [21] Annick Laruelle and Federico Valenciano. 2007. Bargaining in committees as an extension of Nash's bargaining theory. *Journal of Economic Theory* 132, 1 (2007), 291–305.
- [22] Roger Lera-Leri, Filippo Bistaffa, Marc Serramia, Maite Lopez-Sanchez, and Juan Rodríguez-Aguilar. 2022. Towards Pluralistic Value Alignment: Aggregating Value Systems Through lp-Regression. In *Int. Conf. on Autonomous Agents and Multiagent Systems (New Zealand) (AAMAS '22)*. IFAAMAS, Richland, SC, 780–788.
- [23] Enrico Liscio, Roger Lera-Leri, Filippo Bistaffa, Roel IJ. Dobbe, Catholijn M. Jonker, Maite Lopez-Sanchez, Juan A. Rodríguez-Aguilar, and Pradeep K. Murukannaiah. 2023. Value Inference in Sociotechnical Systems. In *Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS '23)*. IFAAMAS, Richland, SC, 1774–1780.
- [24] Enrico Liscio, Michiel van der Meer, Luciano C. Siebert, Catholijn M. Jonker, Niek Mouter, and Pradeep K. Murukannaiah. 2021. Axes: Identifying and Evaluating Context-Specific Values. In *Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS '21)*. IFAAMAS, Richland, SC, 799–808.
- [25] Miles Lubin, Oscar Dowson, Joaquim Dias Garcia, Joey Huchette, Benoît Legat, and Juan Pablo Vielma. 2023. JuMP 1.0: Recent improvements to a modeling language for mathematical optimization. *Mathematical Programming Computation* 15, 3 (2023), 581–589. <https://doi.org/10.1007/s12532-023-00239-3>
- [26] Thibaut Lust and Jacques Teghem. 2012. The multiobjective multidimensional knapsack problem: a survey and a new approach. *International Transactions in Operational Research* 19, 4 (2012), 495–520.
- [27] Scharfbillig M, Smillie L, Mair D, Sienkiewicz M, Keimer J, Pinho Dos Santos R, Vinagreiro Alves H, Vecchione E, and Scheunemann L. 2021. *Values and Identities - a policymaker's guide*. Scientific analysis or review KJ-NA-30800-EN-N (online), KJ-NA-30800-EN-C (print), KJ-NB-30800-EN-Q. European Commission, Luxembourg (Luxembourg). [https://doi.org/10.2760/349527\(online\),10.2760/022780\(print\),10.2760/059689](https://doi.org/10.2760/349527(online),10.2760/022780(print),10.2760/059689)
- [28] Decide Madrid. 2022. Presupuestos participativos 2022. <https://decide.madrid.es/presupuestos>. Accessed August 2023.
- [29] Marco Mariotti. 1998. Nash bargaining theory when the number of alternatives can be finite. *Social choice and welfare* 15 (1998), 413–421.
- [30] Nieves Montes and Carles Sierra. 2021. Value-Guided Synthesis of Parametric Normative Systems. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems (Virtual Event, United Kingdom) (AAMAS '21)*. IFAAMAS, Richland, SC, 907–915.
- [31] John F. Nash. 1950. The Bargaining Problem. *Econometrica* 18, 2 (1950), 155–162. <http://www.jstor.org/stable/1907266>
- [32] City of Cape Town. 2022. The Public Participation Process for the Cape Flats Aquifer. <https://www.umvoto.com/our-work/public-participation-process/>. Accessed Oct 2023.
- [33] City of Varsaw participatory budget. 2023. Varsaw participatory budget. <https://en.um.warszawa.pl/-/participatory-budget>.
- [34] Global Campaign on Urban Governance. 2004. *72 frequently asked questions about participatory budgeting*. UN-HABITAT, Nairobi, Kenya.
- [35] Décider Paris. 2023. Budget Participatif. https://decider.paris.fr/bp/jsp/site/Portal.jsp?page_id=10. Accessed August 2023.
- [36] Participedia. 2021. The case of Varsaw participatory budget. <https://participedia.net/case/7815>.
- [37] Mahendra Prasad. 2018. *Social choice and the value alignment problem*. Chapman and Hall/CRC, New York, 291–314. <https://doi.org/10.1201/9781351251389-21>
- [38] Shalom H Schwartz. 2012. An overview of the Schwartz theory of basic values. *Online readings in Psychology and Culture* 2, 1 (2012), 2307–0919.
- [39] Marc Serramia, Maite López-Sánchez, Juan A. Rodríguez-Aguilar, and Patricio Escobar. 2019. Optimising Participatory Budget Allocation: The Decidim Use Case. In *Proceedings of the 22nd International Conference of the Catalan Association for Artificial Intelligence (Frontiers in Artificial Intelligence and Applications, Vol. 319)*. Jordi Sabater-Mir, Vicenç Torra, Isabel Aguiló, and Manuel González Hidalgo (Eds.). IOS Press, Mallorca, 193–202. <https://doi.org/10.3233/FAIA190124>
- [40] Marc Serramia, Maite Lopez-Sanchez, Juan A. Rodriguez-Aguilar, and Stefano Moretti. 2024. Code and Supplementary material. https://github.com/jariia/Participatory_budgeting.
- [41] Marc Serramia, Maite Lopez-Sanchez, Juan A. Rodriguez-Aguilar, Manel Rodriguez, Michael Wooldridge, Javier Morales, and Carlos Ansoategui. 2018. Moral Values in Norm Decision Making. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems (Stockholm, Sweden) (AAMAS '18)*. IFAAMAS, Richland, SC, 1294–1302.
- [42] Marc Serramia, Manel Rodriguez-Soto, Maite Lopez-Sanchez, Juan A. Rodriguez-Aguilar, Filippo Bistaffa, Paula Boddington, Michael Wooldridge, and Carlos Ansoategui. 2023. Encoding Ethics to Compute Value-Aligned Norms. *Minds and Machines* (2023). <https://doi.org/10.1007/s11023-023-09649-7>
- [43] Paolo Spada. 2017. Dataset on Brazilian Participatory Budgeting: 1989 to 2012. <https://doi.org/10.7910/DVN/EDSNJS>
- [44] Warsaw's statistical office. 2021. Public finances of the city of Warsaw in 2021. <https://warszawa.stat.gov.pl/en/information-about-voivodship/capital-of-voivodship/public-finance/>. Accessed August 2023.
- [45] Dariusz Stolicki, Stanisław Szufa, and Nimrod Talmon. 2020. Pabulib: A Participatory Budgeting Library. [arXiv:2012.06539 \[cs.DC\]](https://arxiv.org/abs/2012.06539)
- [46] Dariusz Stolicki, Stanisław Szufa, and Nimrod Talmon. 2021. Pabulib, Poland Warszawa 2021 data (19 datasets, 18 for districts of Warsaw and one for the whole city). <http://pabulib.org/?year=2021&city=Warszawa>. Accessed August 2023.
- [47] Satya Tamby and Daniel Vanderpooten. 2021. Enumeration of the nondominated set of multiobjective discrete optimization problems. *INFORMS Journal on Computing* 33, 1 (2021), 72–85.
- [48] William Thomson. 1994. Chapter 35 Cooperative models of bargaining. In *Handbook of Game Theory with Economic Applications*. Vol. 2. Elsevier, Amsterdam, 1237–1284. [https://doi.org/10.1016/S1574-0005\(05\)80067-0](https://doi.org/10.1016/S1574-0005(05)80067-0)
- [49] Ibo van de Poel and Lambert Royakkers. 2011. *Ethics, Technology, and Engineering: An Introduction*. Wiley-Blackwell, US.
- [50] MK Vijaymeena and K Kavitha. 2016. A survey on similarity measures in text mining. *Machine Learning and Applications: An International Journal* 3, 2 (2016), 19–28.
- [51] Marc Visée, Jacques Teghem, Marc Pirlot, and Ekunda L Ulungu. 1998. Two-phases method and branch and bound procedures to solve the bi-objective knapsack problem. *Journal of Global Optimization* 12, 2 (1998), 139–155.
- [52] WVS. 2023. World Values Survey. <https://www.worldvaluessurvey.org/wvs.jsp>, last visited on Sept. 2023..
- [53] Miguel Ángel Domínguez-Ríos, Francisco Chicano, and Enrique Alba. 2021. Effective anytime algorithm for multiobjective combinatorial optimization problems. *Information Sciences* 565 (2021), 210–228. <https://doi.org/10.1016/j.ins.2021.02.074>