Participatory Budgeting with Donations and Diversity Constraints

Jiehua Chen, Martin Lackner, Jan Maly
TU Wien, Vienna, Austria
jiehua.chen@tuwien.ac.at, {lackner, jmaly}@dbai.tuwien.ac.at

Abstract

Participatory budgeting (PB) is a democratic process where citizens jointly decide on how to allocate public funds to indivisible projects. In this work, we focus on PB processes where citizens may provide additional money to projects they want to see funded. We introduce a formal framework for this kind of PB with donations. Our framework also allows for diversity constraints, meaning that each project belongs to one or more types, and there are lower and upper bounds on the number of projects of the same type that can be funded. We propose three general classes of methods for aggregating the citizens' preferences in the presence of donations and analyze their axiomatic properties. Furthermore, we investigate the computational complexity of determining the outcome of a PB process with donations and of finding a citizen's optimal donation strategy.

1 Introduction

Participatory budgeting (PB) (Cabannes 2004; Shah 2007) is a democratic tool that enables voters to directly decide about budget spending. The general procedure of PB is that voters are presented with a number of projects (e.g., building a library or a park) and are asked to vote on these projects. Then, a PB aggregation rule is used to select a subset of projects—a so-called *bundle*—to be funded. This bundle has to be feasible, which typically means that the total cost must not exceed the available budget, and sometimes further adhere to fairness constraints.

In this paper, we propose a new addition to the participatory budgeting process, namely adding donations. In our model, voters can pledge donations to projects they support. If such a project is funded, the donations are levied and only the remaining cost is covered by the public budget. Consequently, projects with donations can be funded with a reduced impact on the public budget. At first glance, allowing donations in PB referenda brings major advantages: As the total available budget increases, a larger overall satisfaction is achievable. In addition, voters with an intense preference for a project can support this project financially and thus increase the chance of it being funded. However, allowing donations also comes with a significant risk, namely that it may allow wealthier voters to donate more money than voters who are less well off and thus exert an unfairly large influence on the PB process. Therefore, one of the main goals of the paper is to determine whether it is possible to include donations in the PB process in a way that avoids this risk while also achieving the advantages.

Our chosen model is based on PB with cardinal preferences (see also Peters, Pierczynski, and Skowron (2020)), i.e., voters have numbers associated with projects that reflect their preferences. Cardinal preferences capture, e.g., settings with approval ballots (only 0 and 1 are used), settings where voters can distribute points to projects (where usually the sum of points is fixed and the same for all), and settings where these numbers accurately correspond to the utility of voters. Further, our model allows for diversity constraints (see also (Benabbou, Chakraborty, and Zick 2019)): Each project belongs to one or more types (based on classifications such as "youth and education" or "transport and mobility") and for each type there is a minimum and maximum number of projects to be funded. This added generality allows us to model many interesting variations of PB, for example city-wide referenda where districts have their own "project quota". Note that it is straightforward to extend our model to include constraints with a minimum/maximum amount of budget spent (see also Hershkowitz et al.).

As mentioned before, the largest concern of allowing donations in PB is that voters who can afford to donate money have additional power to influence the outcome. This may undermine the acceptance of the process among the voters who cannot afford to donate. Therefore, we consider the following desideratum as crucial for aggregation rules in PB with donations.

D1 *Donation-no-harm*. Allowing donations *should not* make any voter less satisfied (independent of whether the voter donated herself).

A natural approach for handling donations in PB is to reduce the cost of a project by the amount of donations pledged to the project and then to apply a normal PB aggregation rule. In this paper, we exemplarily examine eight standard PB aggregation rules, four based on global optimization and the remaining four on greedy optimization (see Section 2 for the formal definitions; see also Aziz and Shah 2020). We show that this approach, however, violates D1 for all eight standard PB rules. Consequently, we propose two further natural approaches for adapting a PB rule R to deal with donations. The first, *Sequential-R*, is to first run R without considering donations and then run R again with the money saved due to donations. If possible, the process is repeated. The second approach, *Pareto-R*, also first applies R ignoring all donations. Based on the winning bundle A (i.e., a subset of winning projects), it selects a bundle with maximum social welfare among the bundles that Pareto-dominate A, taking donations into account. Both Sequential-R and Pareto-R satisfy D1 for all considered PB aggregation rules.

In order to gain a better understanding of the advantages and disadvantages of different ways of adapting PB rules for donations, we consider three more desiderata. All three capture the idea that donating more money should not have unintended, harmful consequences.

D2 *Donation-project-monotonicity*. Increasing the donation from any voter to a *winning* project *should not* lead to this project not winning anymore.

D3 *Donation-welfare-monotonicity*. Increasing the donation from any voter to a project *should not* lead to a decrease of the social welfare (for a given welfare definition).

D4 *Donation-voter-monotonicity*. Donating to a project *should not* make the voter who donates less satisfied than not donating to this project (keeping her donations to other projects unchanged).

We find that Pareto-R has especially nice axiomatic properties since it satisfies D1–3 (for social welfare notions compatible to the one used in the rule). However, it fails to satisfy D4. In fact, this is essentially unavoidable since we show that—under natural assumptions—D4 is impossible to satisfy. All results hold independently of diversity constraints.

In addition to the axiomatic analysis, we study the computational complexity questions that arise in our framework. We focus on two computational problems.

The first problem, called R-WINNER, is to decide whether a given bundle is a winner under R. We show that winner determination for all *global rules* (i.e., based on global optimization) is either strongly or weakly coNP-hard. In the latter case, the problem can be solved in pseudo-polynomial time for a constant number of types. The hardness results also hold for their sequential and Pareto variants. On the other hand, winner determination for all *greedy rules* (i.e., based on greedy optimization) is polynomial-time solvable. The tractability result also holds for their sequential variant. Therefore, the sequential variant of the greedy rules can be seen as a reasonable alternative to their intractable Pareto variant, as it satisfies D1 and D2 and offers better computational properties.

As no reasonable PB rule can satisfy D4, voters generally need to carefully consider which projects they want to donate to. Therefore, we study a second problem, called R-DONATION, which is to decide whether a given voter can effectively spend a given amount of money so as to achieve a higher utility (than with an initial donation). While it is straightforward that R-DONATION is naturally contained in Σ_2^{P}, a complexity class from the second level of the polynomial hierarchy (Papadimitriou 1994), the presence of diversity constraints makes the problem indeed Σ_2^{P}-complete for all global rules and NP-complete for all greedy rules, even under severe restrictions to the input instances. We also show a somewhat unexpected result that even if no diversity constraints are imposed, the problem is at least beyond NP for global rules, while it remains (weakly) NP-hard for some greedy rules and becomes polynomial-time solvable for the remaining ones.

To sum up, our work provides a first axiomatic and computational analysis of PB with donations and diversity constraints, in the form of both upper and lower bounds. We discuss features and pitfalls of this idea, propose methods to handle donations, and analyze their computational demands. Table 1 summarizes our findings. Due to space limits, most proofs are deferred to (Chen, Lackner, and Maly 2021).

Related work. Participatory budgeting has received substantial attention through the lens of (computational) social choice in recent years, see e.g., (Fain, Goel, and Munagala 2016; Aziz, Lee, and Talmon 2018; Freeman et al. 2019; Goel et al. 2019; Laruelle 2021); we refer to the survey by Aziz and Shah (2020) for a detailed overview of this line of research. However, donations have not been considered in the indivisible PB model that we are concerned with in this work. The allocation of donations has been studied in a model related to divisible participatory budgeting albeit without external budget (**?**).

In contrast, diversity constrains have been studied in PB in the form of an upper bound on the amount of money spent on each type (Jain et al. 2021) . However, to the best of our knowledge, our work is the first to consider diversity constrains with both upper and lower bounds. Additionally, PB with project interactions (Jain, Sornat, and Talmon 2020) is another approach using project types to guarantee diverse outcomes, albeit by changing the utility functions of the voters instead of the set of feasible outcomes. Finally, diversity constraints have been studied in multi-winner voting (Bredereck et al. 2018; Celis, Huang, and Vishnoi 2018; Yang and Wang 2018; Bei et al. 2020), which can be considered a special case of PB where projects have unit costs.

2 Preliminaries

Given a non-negative integer z, let $[z]$ denote the set $\{1, 2, \dots, z\}$. The input of our PB problem consists of a set of m projects $C = [m]$, a set of n voters $V = [n]$ and a set of types $T = [t]$ along with the following extra information: Each project $j \in C$ has a *cost* $\mathsf{c}_j \in \mathbb{N}_0$, and a *type vector* $\boldsymbol{\tau}_j \in \{0, 1\}^t$, where $\boldsymbol{\tau}_j[z] = 1$ means that project j has type z; and $\boldsymbol{\tau}_j[z] = 0$ otherwise. Each voter $i \in V$ has (i) a *satisfaction function* $\mathsf{sat}_i \colon C \to \mathbb{N}_0$, which models how much she would like a project to be funded, and (ii) a *contribution vector* $\boldsymbol{b}_i \in \mathbb{N}_0^m$ such that for each project $j \in C$ the value $\boldsymbol{b}_i[j]$ indicates how much money she is willing to donate if project j should be selected.

In the following, we call the tuple $\mathcal{P} = (t, (\mathsf{c}_j)_{j\in[m]}, (\boldsymbol{\tau}_j)_{j\in[m]}, (\mathsf{sat}_i)_{i\in[n]}, (\boldsymbol{b}_i)_{i\in[n]})$ a *PB profile* and call every subset of projects a *bundle*. A *PB instance* $I = (\mathcal{P}, B, \boldsymbol{\ell}, \boldsymbol{u})$ contains, in addition to the PB profile $\mathcal{P}$, a set of constraints that a winning bundle has to satisfy. These are determined by the budget $B \in \mathbb{N}_0$ and the

PB Rule R	D1	D2	D3	D4	R-WINNER w/o diversity	R-WINNER w. diversity	R-DONATION w/o diversity	R-DONATION w. diversity
Apply $\mathsf{R}^\star_\diamond$	×	✓	✓	×	w. coNP-c / coNP-c	coNP-c	w. NP-h + w. coNP-h / $\mathrm{P}^{\mathrm{NP}}_{\|\|}$-h	Σ^{P}_2-c
Sequential-$\mathsf{R}^\star_\diamond$	✓	✓	×	×	w. coNP-h / coNP-c	coNP-h	w. NP-h + w. coNP-h / $\mathrm{P}^{\mathrm{NP}}_{\|\|}$-h	Σ^{P}_2-c
Pareto-$\mathsf{R}^\star_\diamond$	✓	✓	✓	×	coNP-h	coNP-h	NP-h + coNP-h / $\mathrm{P}^{\mathrm{NP}}_{\|\|}$-h	Σ^{P}_2-c
Apply $\mathsf{G}^\star_\diamond$	×	✓	×	×	in P	in P	w. NP-h / in P	NP-c
Sequential-$\mathsf{G}^\star_\diamond$	✓	✓	×	×	in P	in P	w. NP-h / NP-c	NP-c
Pareto-$\mathsf{G}^\star_\diamond$	✓	✓	✓	×	coNP-c	coNP-c	NP-h	NP-h

Table 1: Desiderata for the PB rules and complexity results. $\mathsf{R}^\star_\diamond$ and $\mathsf{G}^\star_\diamond$ stand for rules based on global optimization and greedy optimization, respectively (c.f. Section 2). "w." stands for "weakly".

diversity constraints specified by two vectors $\boldsymbol{\ell} \in \mathbb{N}_0^t$ and $\boldsymbol{u} \in \mathbb{N}_0^t$ representing the lower and upper bound on the number of projects funded per type. Throughout the paper, we assume that $\mathcal{P}$ denotes a PB profile of the form $(t, (\mathsf{c}_j)_{j\in[m]}, (\boldsymbol{\tau}_j)_{j\in[m]}, (\mathsf{sat}_i)_{i\in[n]}, (\boldsymbol{b}_i)_{i\in[n]})$ and I denotes a PB instance of the form $(\mathcal{P}, B, \boldsymbol{\ell}, \boldsymbol{u})$.

We say that a bundle $A \subseteq C$ is *feasible* for I if both the budget and diversity constraints are fulfilled, i.e., if:

Budget constraint: $\sum_{j\in A} \max(0, \mathsf{c}_j - \sum_{i\in V} \boldsymbol{b}_i[j]) \leq B$.

Diversity constr.: $\boldsymbol{\ell}[z] \leq \sum_{j\in A} \boldsymbol{\tau}_j[z] \leq \boldsymbol{u}[z], \quad \forall z \in T$.

We write $\mathbb{C}(I)$ to denote the set of all feasible bundles for I. We say that A is *exhaustive* if adding any additional project to A will violate the budget or diversity constraints.

Finally, we introduce some additional notions and notations. We say that voter i's contribution vector $\boldsymbol{b}_i$ is *satisfaction consistent* if (i) for all $j \in C$ with $\mathsf{sat}_i(j) = 0$ it holds that $\boldsymbol{b}_i[j] = 0$, and (ii) for all $j, j' \in C$ with $\boldsymbol{b}_i[j] > \boldsymbol{b}_i[j']$ it holds that $\mathsf{sat}_i(j) > \mathsf{sat}_i(j')$. Further, we say that a contribution vector $\boldsymbol{b}'$ is a *j-variant* of a contribution vector $\boldsymbol{b}$ if for each $j' \in C$ with $j' \neq j$ it holds that $\boldsymbol{b}'[j'] = \boldsymbol{b}[j']$ (they only differ for project j). Given a contribution vector $\boldsymbol{b}'_v$ for a voter v we use $I - \boldsymbol{b}_v + \boldsymbol{b}'_v$ to denote the PB instance where the contribution vector of v is replaced with $\boldsymbol{b}'_v$. For a PB instance I with profile $\mathcal{P}$, let $\mathcal{P}^0$ and I^0 denote the profile and instance derived from $\mathcal{P}$ and I where all donations are zero.

3 Aggregation Rules

We consider aggregation rules that select a *winning bundle*, which is feasible and maximizes the satisfaction of the voters. To achieve this, we aggregate the satisfactions in two steps. First, we aggregate the voters' utility towards a bundle of projects, and then we aggregate the utilities of all voters. We consider two options for each step.

We call a function which lifts satisfaction functions for single projects to bundles *utility function*, denoted as μ. We consider two standard functions suited for cardinal preferences: summing the satisfaction of each project (additive) or choosing the highest satisfaction of all projects (maximum) in a bundle A.

$$\mu_i^+(I, A) \coloneqq \sum_{j\in A} \mathsf{sat}_i(j), \quad \mu_i^{\max}(I, A) \coloneqq \max_{j\in A} \mathsf{sat}_i(j).$$

We omit I from the function and write $\mu_i(A)$ instead if the corresponding PB instance is clear from the context.

Next, given a PB instance I and a bundle $A \subset C$, a *scoring function* score computes a number indicating the overall utilities of the voters towards A. We consider two types of scoring functions: the *sum* scoring functions return the sum of utilities of the voters towards a given bundle; the *min* scoring functions return the minimum satisfaction. For each $\star \in \{\max, +\}$, we define

$$\mathsf{score}^\star_\Sigma(I, A) \coloneqq \sum_{i\in V} \mu_i^\star(A), \ \mathsf{score}^\star_{\min}(I, A) \coloneqq \min_{i\in V} \mu_i^\star(A).$$

We look at eight aggregation rules based on either global or greedy optimization. Let $\star \in \{\max, +\}$ and $\diamond \in \{\Sigma, \min\}$.

Aggregation rules based on global optimization. Rule $\mathsf{R}^\star_\diamond$ selects a feasible bundle A with maximum $\mathsf{score}^\star_\diamond(I, A)$. As convention, in case of multiple feasible bundles have maximum score, we select one according to an arbitrary but fixed tie-breaking rule. Define $\mathcal{R} \coloneqq \{\mathsf{R}^+_{\min}, \mathsf{R}^{\max}_{\min}, \mathsf{R}^+_\Sigma, \mathsf{R}^{\max}_\Sigma\}$.

Aggregation rules based on greedy approaches. The greedy rule $\mathsf{G}^\star_\diamond$ proceeds iteratively by adding in each step a project p to the winning bundle C' that maximizes $\mathsf{score}^\star_\diamond(C' \cup \{p\})$ among the projects for which $C' \cup \{p\}$ is feasible. In case of more than one project maximizing the score in an iteration, we select one according to an arbitrary, fixed tie-breaking order. Observe that the greedy approach selects exhaustive bundles as long as the diversity constraints contain no lower bounds. Indeed, the greedy approach may not even produce a feasible bundle in the presence of lower bounds. Therefore, we will always assume that no lower bounds are specified when talking about the greedy approach. Define $\mathcal{G} \coloneqq \{\mathsf{G}^+_{\min}, \mathsf{G}^{\max}_{\min}, \mathsf{G}^+_\Sigma, \mathsf{G}^{\max}_\Sigma\}$.

Aggregation rules with donations. As the eight aggregation rules maximize a function over the set of feasible bundles, they can simply handle donations via the definition of budget constraints. Note that, using this approach, allowing donation is equivalent to reducing the cost of the respective project. However, as we will see, this simple way of handling donations has some undesirable consequences. Therefore we also propose to consider two other natural variants for handling donations. Let $R \in \mathcal{R} \cup \mathcal{G}$.

Sequential-R first applies R on I^0, i.e., the instance without donations. If afterwards some budget is left (due to donations), R is applied again with the remaining budget but

Algorithm 1: Sequential-$R(I)$

$C \leftarrow [m]$;
while *C changes in the previous iteration* **do**
  $A_0 \leftarrow R(\mathcal{P}^0(C), B, \boldsymbol{\ell}, \boldsymbol{u})$;
  $C \leftarrow C \setminus A_0$;
  $\boldsymbol{\ell} \leftarrow \boldsymbol{\ell} - \sum_{j \in A_0} \boldsymbol{\tau}_j$; $\boldsymbol{u} \leftarrow \boldsymbol{u} - \sum_{j \in A_0} \boldsymbol{\tau}_j$;
  $B \leftarrow B - \sum\limits_{j \in A_0} \max(0,\ \mathsf{c}_j - \sum\limits_{j \in A_0, i \in V} \boldsymbol{b}_i[j])$;
return $([m] \setminus C) \cup R(\mathcal{P}(C), B, \boldsymbol{\ell}, \boldsymbol{u})$

still without donations; this step is repeated as long as new projects are added. In a last step, R is applied directly, thus guaranteeing an exhaustive bundle.[1]

The second method is *Pareto-R*. Let $A_0 = R(I^0)$. Now, consider the set of bundles $\mathcal{X}$ consisting of A_0 and all bundles $A^* \in \mathbb{C}(I)$ that μ-dominate A_0. Here, a bundle $X \subseteq C$ *μ-dominates* another bundle $Y \subseteq C$ if for each voter $i \in V$ it holds that $\mu_i(X) \geq \mu_i(Y)$ and there exists a voter $i \in V$ with $\mu_i(X) > \mu_i(Y)$. Pareto-R chooses a bundle $A \in \mathcal{X}$ with maximum $\mathsf{score}(A)$. Observe that Pareto-R can be applied to the greedy rules, but the resulting rule violates the idea of greediness and avoiding global optimization. Indeed, as we will see, Pareto-R negates the main advantage of the greedy rules that the winner determination can be done in polynomial time (Theorem 11).

Example 1. *Consider the following PB instance I with 5 projects $p_1, \ldots, p_5$, two voters, budget $B = 5$, and without diversity constraints:*

	$\mathsf{c}(\cdot)$	sat_1	sat_2	b_1	b_2
p_1	3	5	5	0	1
p_2	3	9	0	0	0
p_3	2	1	2	0	0
p_4	3	3	3	0	0
p_5	1	1	1	0	0

We consider rule R^+_Σ, and its sequential and Pareto variants. One can verify that the winner under R^+_Σ is $A_1 = \{p_1, p_2\}$ as it maximizes score^+_Σ ($= 19$). Without donations, the winner is $A_0 = \{p_1, p_3\}$. Hence, Sequential-R^+_Σ starts by selecting A_0. Then it runs R^+_Σ on the instance created by removing p_1 and p_3 from $\mathcal{P}^0$ with budget of 1 (the cost for p_1 is reduced by 1 due to voter 2's donation). Now, $\{p_5\}$ is the winner (for R^+_Σ), leaving 0 budget for the next round. Hence the final winning bundle is $A_2 = \{p_1, p_3, p_5\}$.

Pareto-R^+_Σ maximizes score^+_Σ among the projects which μ^+-dominate A_0. While A_1 has a higher score than A_0 it does not μ^+-dominate A_0 since voter 2 is worse off ($\mu^+_2(A_1) = 5 < 7 = \mu^+_2(A_0)$). Indeed, $A_3 = \{p_1, p_4\}$ and A_2 are the only feasible bundles which μ^+-dominate A_0. Among those A_3 has the highest score and is hence the winner under Pareto-R^+_Σ.

[1] This step is necessary as some remaining projects may be unaffordable without donations.

4 Axioms Regarding Donations

In this section, we axiomatically analyze the different methods for handling donations. We start with a property that we view as crucial for the acceptability of donations in PB, namely that no voter may end up less satisfied with the outcome than in a process without donations.

Definition 1 ($\hat{\mu}$-donation-no-harm)**.** An aggregation rule R satisfies *$\hat{\mu}$-donation-no-harm* if for each PB instance I and each voter x it holds that $\hat{\mu}_x(R(I)) \geq \hat{\mu}_x(R(I^0))$.

If a rule R is based on a utility function μ, then that R *satisfies donation-no-harm* means it satisfies μ-donation-no-harm. We will use the same shorthand for the other axioms. It turns out that the naive approach of just reducing the costs of the projects does not guarantee this crucial property.

Theorem 1. *All $R \in \mathcal{R} \cup \mathcal{G}$ fail donation-no-harm, even if there are no diversity constraints and only satisfaction consistent donations are allowed.*

Proof sketch. To show the statement, consider the following PB instance I with three voters $1, 2, 3$ and three projects p_1, p_2, p_3. The budget B is 5, and there are neither donations nor diversity constraints. The costs of the projects, and the preferences of the voters are as follows:

	$\mathsf{c}(\cdot)$	sat_1	sat_2	sat_3
p_1	2	6	2	2
p_2	4	1	4	4
p_3	3	0	5	3

There are two feasible and exhaustive bundles $A_1 = \{p_1, p_3\}$ and $A_2 = \{p_2\}$. One can verify that A_1 is the winner for every $R \in \mathcal{R} \cup \mathcal{G}$. Now, if voter 3 donates one unit to p_2 then $A_3 = \{p_1, p_2\}$ becomes feasible. Indeed, A_3 is the unique winner under every $R \in \mathcal{R} \cup \mathcal{G}$. One can verify that this is a worse result for voter 2 under μ^+ and $\mu^{\max}$. ☐

It is straightforward to see that Pareto-R and Sequential-R satisfy donation-no-harm.

Proposition 2. *For each R, Pareto-R satisfies $\hat{\mu}$-donation-no-harm, where $\hat{\mu}$ is used for the $\hat{\mu}$-domination, and Sequential-R satisfies $\hat{\mu}$-donation-no-harm for all* monotonic *utility function $\hat{\mu}$ (i.e., for all $A \subseteq B$ and for all voters i it holds that $\hat{\mu}_i(A) \leq \hat{\mu}_i(B)$).*

While donation-no-harm guarantees that there is no incentive for any voter to reject the inclusion of donations in the PB process, it would also be desirable to incentivize voters to donate. Clearly, it is in general not possible to guarantee every voter that donating money will increase his satisfaction. However, we would like to ensure that donating money has no unintended, harmful consequences. First, we look at unintended consequences for projects:

Definition 2 (donation-project-monotonicity)**.** An aggregation rule R satisfies *donation-project-monotonicity* if for each PB instance I, each voter x, and each donation $\boldsymbol{b}'_x$ which is a j-variant of $\boldsymbol{b}_x$ with $\boldsymbol{b}_x[j] < \boldsymbol{b}'_x[j]$ it holds that if $j \in R(I)$, then $j \in R(I - \boldsymbol{b}_x + \boldsymbol{b}'_x)$.

Increasing the donation for a project j only makes new bundles available that all include j and has no effect on the other bundles. Therefore, the following holds:

Proposition 3. *R, Sequential-R, and Pareto-R satisfy donation-project-monotonicity for all* $R \in \mathcal{R} \cup \mathcal{G}$.

Next, we consider the overall satisfaction of the voters.

Definition 3 (score-donation-welfare-monotonicity)**.** An aggregation rule R satisfies score-*donation-welfare-monotonicity* if for each PB instance I, each voter x, and each contribution vector $\boldsymbol{b}'_x$ which is a j-variant of $\boldsymbol{b}_x$ with $\boldsymbol{b}_x[j] < \boldsymbol{b}'_x[j]$ it holds that $\mathsf{score}(I, R(I)) \leq \mathsf{score}(I', R(I'))$, where $I' = I - \boldsymbol{b}_x + \boldsymbol{b}'_x$.

As before, we omit score, if it is clear from the context. All rules based on global optimization and their Pareto variants satisfy donation-welfare-monotonicity, as increasing the donation only increases the set of feasible bundles.

Proposition 4. R *and Pareto-*R *satisfy donation-welfare-monotonicity for all* $\mathsf{R} \in \mathcal{R}$.

This leaves Sequential-R and the greedy rules, which do not satisfy donation-welfare-monotonicity.

Proposition 5. *For all* $\mathsf{R} \in \mathcal{R}$ *and* $\mathsf{G} \in \mathcal{G}$, *Sequential-*R, G, *and Sequential-*G *fail donation-welfare-monotonicity, even if there are no diversity constraints and only satisfaction-consistent donations are allowed.*

The final property asserts that a voter should not be worse off if she decides to donate more money to a project. We consider the slightly weaker requirement that a voter should not be worse off if she donates money for a project than if she donates no money for that project.

Definition 4 ($\hat{\mu}$-donation-voter-monotonicity)**.** An aggregation rule R satisfies *$\hat{\mu}$-donation-voter-monotonicity* if for each PB instance I, each voter x, and each donation $\boldsymbol{b}'_x$ which is a j-variant of $\boldsymbol{b}_x$ such that $\boldsymbol{b}'_x[j] = 0$ it holds that

$$\hat{\mu}_x(R(I)) \geq \hat{\mu}_x(R(I - \boldsymbol{b}_x + \boldsymbol{b}'_x)).$$

Unfortunately, this property is essentially impossible to satisfy by an exhaustive rule. To be more precise, an exhaustive rule which satisfies a very weak additional axiom cannot satisfy $\hat{\mu}$-donation-voter-monotonicity for nearly all utility functions $\hat{\mu}$. To show our result in its full generality, we need to introduce two rather technical axioms, weak continuity and weak responsiveness. Both are weak forms of well-known axioms from the voting (Zwicker 2016) resp. fair allocation (Bouveret et al. 2016) literature and essentially state that continuity and responsiveness have to be satisfied for some particularly clear-cut cases. First, we define weak continuity, which is satisfied by all considered rules.

Definition 5. An aggregation rule R satisfies *weak continuity* if for each PB instance I and for each project $j \in [m]$ the following holds: If $\mathsf{sat}_i(j) > 0$ for all $i \in [n]$ and there exists a feasible bundle that contains j, then there are values c_0 and k_0 such that j is a winner if one adds $k \geq k_0$ voters $i_1^*, \ldots, i_k^*$ that do not donate any money and have satisfaction functions such that for all $\ell \in [k]$ we have $\mathsf{sat}_{i_\ell^*}(j) \geq c_0$ and $\mathsf{sat}_{i_\ell^*}(j^*) = 0$ for all projects $j^* \neq j$.

Proposition 6. *R, Sequential-R and Pareto-R satisfy weak continuity for each* $R \in \mathcal{R} \cup \mathcal{G}$.

Before we define weak responsiveness, we recall that responsiveness requires that for all bundles A and projects $x \in A$ and $y \notin A$ with $\mathsf{sat}(x) < \mathsf{sat}(y)$ we have $\mu(A) < \mu((A \setminus \{x\}) \cup \{y\})$. Weak responsiveness essentially states that this has to hold at least for three satisfaction values and (only) for some bundles of size two.

Definition 6. A utility function μ is *weakly responsive* if there are values u_1, u_2 and u_3 with $u_1 > 0$ and $u_2 < u_3$ such that for all voters v and projects p_1, p_2 and p_3 with $\mathsf{sat}_v(p_1) = u_1$, $\mathsf{sat}_v(p_2) = u_2$ and $\mathsf{sat}_v(p_3) = u_3$, it holds that $\mu_v(\{p_1, p_2\}) < \mu_v(\{p_1, p_3\})$.

It is straightforward to see that μ^+ and $\mu^{\max}$ are weakly responsive, for example with $u_1 = 1, u_2 = 2$ and $u_3 = 3$, as are most natural utility functions. A notable exception is the function μ^{CC} where $\mu_v^{CC}(A) = 1$ if there is a project $x \in A$ such that $\mathsf{sat}_v(x) > 0$ and $\mu_v^{CC} = 0$ else. For dichotomous preferences, this equals the satisfaction function that is used in the well-known Chamberlin-Courant rule (Lang and Xia 2016). Indeed, all considered aggregation rules satisfies μ^{CC}-donation-voter-monotonicity if voters only donate for projects which offer them positive satisfaction values.

Theorem 7. *No exhaustive aggregation rule that satisfies weak continuity can satisfy μ-donation-voter-monotonicity for any weakly responsive utility function μ, even if we only allow satisfaction consistent contributions and there are no diversity constraints.*

We note that exhaustiveness is indeed necessary. Consider for example the following non-exhaustive rule R^*: Let $p_1, \ldots, p_m$ be an enumeration of C such that $\mathsf{score}_\Sigma^+(\{p_i\}) > \mathsf{score}_\Sigma^+(\{p_j\})$ implies $i < j$. Now, let k be the largest value such that $C_k = \{p_1, \ldots, p_k\}$ is a feasible bundle. Then, C_k is the winner under R^*. R^* satisfies weak continuity by the same argument as G_Σ^+ (see Proposition 6). Moreover, increasing the donation to any project can only increase the value k for which C_k is feasible. Therefore, R^* satisfies μ-donation-voter-monotonicity for all monotonic utility functions μ.

Finally, if we additionally assume that all voters only donate to the projects which give them the highest satisfaction, then $\mathsf{R}_{\min}^{\max}$ and $\mathsf{R}_\Sigma^{\max}$ satisfy donation-voter-monotonicity.

5 Central Computational Problems

We consider two decision problems that arise in our framework. The first one captures the complexity of applying an aggregation method.

R-WINNER

Input: A PB instance I and a bundle A.

Question: Is A a *(co-)winning bundle* under R?

The second problem is concerned with the effective use of donations from a voter's perspective. This problem is particularly crucial in light of Theorem 7, which tells us that voters need to carefully consider how to distribute their donation for nearly every natural voting rule.

R-DONATION

Input: A PB instance I for m projects C, a voter v, and an integer $\delta \in \mathbb{N}$ (voter v's personal budget).

Question: $\exists \boldsymbol{b}'_v \in \mathbb{N}_0^m$ with $\sum_{j \in C} \boldsymbol{b}'_v[j] \leq \delta$ such that $\mu_v(R(I')) > \mu_v(R(I))$, where $I' := I - \boldsymbol{b}_v + \boldsymbol{b}'_v$ and μ denotes the utility function underlying rule R?

$\mathcal{R}$-Winner

First, we upper-bound the complexity of $\mathcal{R}$-Winner.

Theorem 8. R-Winner *is in coNP for each* $\mathsf{R} \in \mathcal{R}$.

Next, we consider the rule R_{Σ}^{+}. Since it generalizes the famous knapsack algorithm, parts of the following results are straightforward.

Theorem 9. R_{Σ}^{+}-Winner *can be solved in* $O(n \cdot m + (B+1) \cdot (m+1)^{t+1} \cdot t)$ *time, and* Sequential-R_{Σ}^{+}-Winner *can be solved in* $O(n \cdot m + (B+1) \cdot (m+1)^2)$ *time if there are no diversity constraints. Even if there is only one voter and no donations are provided, both problems are coNP-hard for dichotomous preferences, and remain weakly coNP-hard when diversity constraints are not present.*

Pareto-R_{Σ}^{+}-Winner *is already coNP-hard even for dichotomous preferences and without diversity constraints.*

Hardness for rules other than R_{Σ}^{+} is also quite straightforward as they generalize commonly used multiwinner rules.

Theorem 10. *For each* $\mathsf{R} \in \{\mathsf{R}_{\min}^{\max}, \mathsf{R}_{\min}^{+}, \mathsf{R}_{\Sigma}^{\max}\}$ *and their sequential and Pareto variants,* R-Winner *is coNP-hard even for projects with unit costs, without diversity constraints or donations, and for dichotomous preferences.*

Finally, winner determination is polynomial-time solvable for all greedy rules and their sequential variants, but this does not extend to the Pareto variants.

Theorem 11. *For all* $\star \in \{+, \max\}$ *and* $\diamond \in \{\min, \Sigma\}$, $\mathsf{G}_{\diamond}^{\star}$-Winner *and* Sequential-$\mathsf{G}_{\diamond}^{\star}$-Winner *are polynomial-time solvable, while* Pareto-$\mathsf{G}_{\diamond}^{\star}$-Winner *is in coNP and at least as hard as* $\mathsf{R}_{\diamond}^{\star}$-Winner. Pareto-$\mathsf{G}_{\Sigma}^{+}$-Winner *without diversity constraints remains coNP-hard even for dichotomous preferences.*

$\mathcal{R}$-Donation

If diversity constraints are present, the picture is quite clear: $\mathcal{R}$-Donation is Σ_2^{P}-complete for all $\mathsf{R} \in \mathcal{R}$ and their two variants, while $\mathcal{R}$-Donation is NP-complete for all $\mathsf{G} \in \mathcal{G}$ and their sequential variants. Without diversity constraints, the complexity results vary: For all aggregation rules $\mathsf{R} \in \mathcal{R}$ except R_{Σ}^{+}, finding an optimal donation is as hard as the complexity class $\mathsf{P}_{||}^{\mathsf{NP}}$ (Papadimitriou 1994, Chapter 17.1), which includes NP. For R_{Σ}^{+}, it is both weakly NP-hard and weakly coNP-hard. This implies, under a widely believed complexity-theoretical assumption, that the problem is beyond NP. On the other hand, $\mathcal{R}$-Donation is polynomial time solvable for $\mathsf{G}_{\diamond}^{\max}$ and becomes NP-hard for its sequential variant while it remains weakly NP-hard for $\mathsf{G}_{\diamond}^{+}$ and all sequential variants. It is NP-hard for the Pareto variant of all greedy rules.

In the following, after locating the complexity upper bound, we first consider the case with diversity constraints, and then that without diversity constraints.

Theorem 12. R-Donation, Sequential-R-Donation, *and* Pareto-R-Donation *are in* Σ_2^P *for each* $\mathsf{R} \in \mathcal{R}$. G-Donation *and* Sequential-G-Donation *are in NP for each* $\mathsf{G} \in \mathcal{G}$.

With diversity constraints Using the power of diversity constraints, we prove Σ_2^{P}-hardness for all $\mathsf{R} \in \mathcal{R}$. That is, finding an effective donation is hard for the complexity class Σ_2^{P}, whenever diversity constraints are involved. All reductions are from a SAT variant, which is proved to be Σ_2^{P}-complete by Chen, Ganian, and Hamm [2020b, Claim 1] and originally used to prove that finding a diverse and stable matching is Σ_2^{P}-hard.

Theorem 13. R-Donation *is* Σ_2^P*-hard for* $\mathsf{R} \in \mathcal{R} \setminus \{\mathsf{R}_{\min}^{\max}\}$ *even for projects with unit costs, zero budget, and dichotomous preferences.* $\mathsf{R}_{\min}^{\max}$-Donation *is* Σ_2^P*-hard even for projects with unit costs, zero budget, and trichotomous preferences. The same hardness holds for the sequential and Pareto variants.*

Next, we show that, in the presence of diversity upper bounds, G-Donation is NP-hard for all greedy rules G and their sequential variants.

Theorem 14. G-Donation *and* Sequential-G-Donation *are NP-hard for all* $\mathsf{G} \in \mathcal{G}$*, even if the budget is zero.*

With no diversity constraints In this case, we were not able to show Σ_2^P-hardness. However, for most $\mathsf{R} \in \mathcal{R}$ and their two variants, we show that it is at least $\mathsf{P}_{||}^{\mathsf{NP}}$-hard. Before we present our results, we recall the following relations among the complexity classes: $(\mathrm{NP} \cup \mathrm{coNP}) \subseteq \mathsf{P}_{||}^{\mathsf{NP}} \subseteq \Sigma_2^{\mathsf{P}}$, where all inclusions are generally assumed to be strict.

Theorem 15. R-Donation *without diversity constraints becomes* $P_{||}^{NP}$*-hard for all* $\mathcal{R} \setminus \{\mathsf{R}_{\Sigma}^{+}\}$ *and its sequential and Pareto variants. Hardness for* $\mathsf{R}_{\Sigma}^{\max}$ *and its two variants hold even for dichotomous preferences.*

Proof sketch. We show the hardness result for $\mathsf{R}_{\Sigma}^{\max}$-Donation by reducing from a $\mathsf{P}_{||}^{\mathsf{NP}}$-complete problem (Spakowski 2005, Theorem 3.2.6), called Max-True-3SAT-Compare: Given two equal-sized sets X and Y of Boolean variables and two equal-sized sets $\phi_1(X)$ and $\phi_2(Y)$ of clauses over X and Y, respectively, where each clause contains at most 3 literals, and both $\phi_1(X)$ and $\phi_2(Y)$ are satisfiable, is "$\mathsf{max}\text{-}\mathbb{1}(\phi_1) \geq \mathsf{max}\text{-}\mathbb{1}(\phi_2)$" true? Herein, given a Boolean formula ϕ over a set Z of variables, $\mathsf{max}\text{-}\mathbb{1}(\phi)$ denotes the maximum number of variables set to true in a satisfying truth assignment for ϕ; if ϕ is not satisfiable, then $\mathsf{max}\text{-}\mathbb{1}(\phi)$ is undefined.

The idea of the reduction is to construct, from an instance $(\phi_1(X), \phi_2(Y))$ of Max-True-3SAT-Compare with $|X| = |Y| = \hat{n}$ and $|\phi_1| = |\phi_2| = \hat{m}$, an equivalent instance of $\mathsf{R}_{\Sigma}^{\max}$-Donation with $2\hat{n}$ X*-projects*, $2\hat{n}$ Y*-projects* (each corresponding to a literal), $\hat{n}$ *auxiliary-projects*, and 3 special projects x_0, y_0, α_0, where our target voter v can only gain satisfaction from x_0. We define the costs of the X- and Y-projects such that projects corresponding to positive literals cost less than projects corresponding to negative literals. Therefore, if more positive projects are funded, then more money is left to select additional auxiliary-projects. The costs of the auxiliary-projects are small in comparison to the other projects. This way, the score of bundle A is linear in the number of positive X-projects (resp. Y-projects) in A. Moreover, we define the budget, the donation bound, and the costs of x_0 and y_0 such that any feasible bundle A with sufficiently large score satisfies the following properties: (i) A contains either x_0 or y_0. (ii) If A contains x_0 (resp. y_0), then it corresponds to a valid truth assignment of ϕ_1 (resp. ϕ_2). Besides voter v, we introduce a large number of additional voters to ensure that any winning (and feasible) bundle must achieve a sufficiently large score. Since voter v is only satisfied with x_0, the only way for her to improve her utility is to ensure that there exists a winning (and feasible) bundle which includes x_0. In order to achieve this, she must donate money to projects such that the number of positive X-projects is at least as large as the number of positive Y-projects in any feasible bundle including y_0, i.e., $\mathsf{max}\text{-}\mathbb{1}(\phi_1) \geq \mathsf{max}\text{-}\mathbb{1}(\phi_2)$.

Formally, let $(\phi_1(X), \phi_2(Y))$ be an instance of MAX-TRUE-3SAT-COMPARE with $X = \{x_1, \ldots, x_{\hat{n}}\}$ and $Y = \{y_1, \ldots, y_{\hat{n}}\}$, $\phi_1 = \{C_1, \ldots, C_{\hat{m}}\}$ and $\phi_2 = \{D_1, \ldots, D_{\hat{m}}\}$. To ease notation, define $L := \hat{n} + 3$ and $K := 2\hat{n} + 2\hat{m} + 4\hat{n}^2 + 4\hat{n}$. We create an instance of $\mathsf{R}^{\max}_{\Sigma}$-DONATION without diversity constraints as follows.

Besides the three distinguished projects x_0, y_0, α_0, we create the following projects: For each $x_i \in X$ (resp. $y_i \in Y$) create two *X-projects* x_i and $\overline{x}_i$ (resp. Y-projects y_i and $\overline{y}_i$). For each $i \in [\hat{n}]$, we create an *auxiliary-project* α_i. The costs of the projects are specified as follows:

x_i	$\overline{x}_i$	y_i	$\overline{y}_i$	α_i	x_0	y_0	α_0
$\hat{n}+1$	$\hat{n}+2$	$\hat{n}+1$	$\hat{n}+2$	i	$2\hat{n}$	$\hat{n}$	B

The voters have dichotomous preferences: If they are *satisfied* with a project, then they are satisfied with value one. Our target voter v is only satisfied with x_0. Additionally, we create the following $L \cdot (2\hat{n} + 2\hat{m} + 4\hat{n}^2 + 4\hat{n}) + \hat{n} + 5 = L \cdot K + \hat{n} + 4$ voters.

- For each $x_i \in X$ (resp. $y_i \in Y$), we create L voters x_i^j (resp. y_i^j), $j \in [L]$, which are only satisfied with the projects x_i and $\overline{x}_i$ (resp. y_i and $\overline{y}_i$).
- For each clause $C_\ell \in \phi_1$ (resp. $D_\ell \in \phi_2$), we create L voters c_ℓ^j (resp. d_ℓ^j), $j \in [L]$, each of which is only satisfied with the X-projects (resp. Y-projects) which correspond to the literals contained in C_ℓ (resp. D_ℓ), and project α_0.
- For each $x_i \in X$, we create $2 \cdot L$ voters u_i^j and $\overline{u}_i^j$, $j \in [L]$. Each u_i^j (resp. $\overline{u}_i^j$) is satisfied with the corresponding X-project x_i (resp. $\overline{x}_i$), and projects x_0 and α_0. Similarly, for each $y_i \in Y$, we introduce $2 \cdot L$ voters, w_i^j and $\overline{w}_i^j$, $j \in [L]$. Each w_i^j (resp. $\overline{w}_i^j$) is satisfied with the corresponding Y-projects y_i (resp. $\overline{y}_i$), and projects y_0 and α_0. Finally, for each $\mathsf{lit} \in X \cup \overline{X}$ and $\mathsf{lit}' \in Y \cup \overline{Y}$, we create L connector-voters who are only satisfied with projects lit, lit', and α_0.
- For each $i \in [\hat{n}]$, we create a voter a_i, who is only satisfied with the auxiliary-projects from $\{\alpha_{\hat{n}}, \alpha_{\hat{n}-1} \ldots, \alpha_{\hat{n}-i+1}\}$. We create one more voter a_0 who is only satisfied with all $\hat{n}$ auxiliary-projects.
- Finally, we create three distinguished voters v_1, v_2, and v_3 such that v_1 is only satisfied with y_0 and α_0, while v_2 and v_3 are only satisfied with α_0.

Finally, to complete the construction, define $B := \hat{n}(3\hat{n} + 6)$ and $\delta := \hat{n}$, and let no voter donates any money initially. Let I denote this PB instance. The proofs of the following and of the remaining results are available in the full version (Chen, Lackner, and Maly 2021): (1) the initial winning bundle has score $\geq L \cdot K + 3$ and the initial utility of v is zero, and (2) $(\phi_1(X), \phi_2(Y))$ and (I, δ) are equivalent, i.e., ϕ_1 admits a satisfying assignment σ_1 such that the number k_1 of X-variables set to true is greater or equal to $\mathsf{max}\text{-}\mathbb{1}(\phi_2)$ iff. there exists a donation vector for v with sum at most δ such that the bundle consisting of x_0, α_{k_1}, the projects corresponding to σ_1, and all Y-projects is a winner after the donation. ☐

We were not able to show $\mathsf{P}^{\mathsf{NP}}_{||}$-hardness for R^{+}_{Σ} and its two variants. However, we show that it is unlikely to be contained in NP or coNP.

Theorem 16. *For the case without diversity constraints,* R^{+}_{Σ}-DONATION *and* Sequential-R^{+}_{Σ}-DONATION *are both weakly NP-hard and weakly coNP-hard, while* Pareto-R^{+}_{Σ}-DONATION *becomes NP-hard and coNP-hard even for dichotomous preferences.*

For the greedy rules, we can show that, in the absence of diversity constraints, if a voter wants to ensure that a project is in the winning bundle, she cannot do better than donating all her available money to it. From this, it follows directly that $\mathsf{G}^{\max}_{\diamond}$-DONATION can be solved in polynomial time.

Theorem 17. *For the case without diversity constraints,* $\mathsf{G}^{\max}_{\diamond}$-DONATION *for each $\diamond \in \{\Sigma, \min\}$ can be solved in polynomial time.*

Interestingly, this does not carry over to the sequential and Pareto variant of all greedy rules.

Theorem 18. *For the case without diversity constraints,* Sequential-$\mathsf{G}^{\max}_{\diamond}$-DONATION *and* Pareto-$\mathsf{G}^{\star}_{\diamond}$-DONATION *remain NP-hard for each $\star \in \{+, \max\}$ and $\diamond \in \{\min, \Sigma\}$. Hardness for Pareto-$\mathsf{G}^{\star}_{\Sigma}$ holds even for dichotomous preferences.*

Theorem 19. *For the case without diversity constraints,* $\mathsf{G}^{+}_{\diamond}$-DONATION *and* Sequential-$\mathsf{G}^{+}_{\diamond}$-DONATION *remain weakly NP-hard for each $\diamond \in \{\Sigma, \min\}$ even if there is only one voter and the budget is zero.*

6 Discussion

To briefly summarize our findings, we conclude that Pareto-$\mathsf{R}^{\star}_{\diamond}$ is axiomatically the most promising implementation of donations (among all exhaustive and weak-continuous rules) in a PB process. If a computationally efficient rule is sought, then we recommend the greedy-based Sequential-$\mathsf{G}^{\star}_{\diamond}$ variants.

For future work, one could investigate whether dropping exhaustiveness or weak continuity yields interesting PB rules that satisfy more axiomatic properties compared to the rules discussed in this paper. Secondly, some of our hardness results are not yet tight. Moreover, a parameterized complexity analysis could help to find effective algorithms (e.g., with the number of voters as parameter). Recent work (Peters, Pierczynski, and Skowron 2020; Hershkowitz et al. 2021) has considered proportionality in PB processes. Merging this line of research with our focus on donations could yield particularly fair and versatile PB voting rules.

Acknowledgments

Jiehua Chen was supported by the WWTF research grant VRG18-012. Martin Lackner and Jan Maly were supported by the FWF research grant P31890. Jan Maly was also supported by the FWF research grant J4581.

References

Aziz, H.; Lee, B. E.; and Talmon, N. 2018. Proportionally Representative Participatory Budgeting: Axioms and Algorithms. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS '18)*, 23–31.

Aziz, H.; and Shah, N. 2020. Participatory Budgeting: Models and Approaches. In Rudas, T.; and Péli, G., eds., *Pathways between Social Science and Computational Social Science: Theories, Methods and Interpretations*. Springer.

Bei, X.; Liu, S.; Poon, C. K.; and Wang, H. 2020. Candidate Selections with Proportional Fairness Constraints. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS '20)*, 150–158.

Benabbou, N.; Chakraborty, M.; and Zick, Y. 2019. Fairness and Diversity in Public Resource Allocation Problems. *IEEE Data Eng. Bull.*, 42(3): 64–75.

Bouveret, S.; Chevaleyre, Y.; Maudet, N.; and Moulin, H. 2016. Fair Allocation of Indivisible Goods. In Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D., eds., *Handbook of Computational Social Choice*, 284–310. Cambridge University Press.

Bredereck, R.; Faliszewski, P.; Igarashi, A.; Lackner, M.; and Skowron, P. 2018. Multiwinner Elections with Diversity Constraints. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI '18)*, 933–940.

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

Celis, L. E.; Huang, L.; and Vishnoi, N. K. 2018. Multiwinner Voting with Fairness Constraints. arXiv:1710.10057.

Chen, J.; Ganian, R.; and Hamm, T. 2020a. Stable Matchings with Diversity Constraints: Affirmative Action is beyond NP. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI '20)*, 146–152.

Chen, J.; Ganian, R.; and Hamm, T. 2020b. Stable Matchings with Diversity Constraints: Affirmative Action is beyond NP. Technical report, arXiv:2001.10087.

Chen, J.; Lackner, M.; and Maly, J. 2021. Participatory Budgeting with Donation and Diversity Constraints. Technical report, arXiv:2104.15075.

Fain, B.; Goel, A.; and Munagala, K. 2016. The Core of the Participatory Budgeting Problem. In *Proceedings of the 12th International Conference on Web and Internet Economics (WINE '16)*, 384–399. Springer.

Freeman, R.; Pennock, D. M.; Peters, D.; and Vaughan, J. W. 2019. Truthful Aggregation of Budget Proposals. In *Proceedings of the 2019 ACM Conference on Economics and Computation (EC '19)*, 751–752.

Goel, A.; Krishnaswamy, A. K.; Sakshuwong, S.; and Aitamurto, T. 2019. Knapsack Voting: Voting Mechanisms for Participatory Budgeting. *ACM Transaction on Economics and Computation*, 7(2): 8:1–8:27.

Hershkowitz, D. E.; Kahng, A.; Peters, D.; and Procaccia, A. D. 2021. District-Fair Participatory Budgeting. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI '21)*.

Jain, P.; Sornat, K.; and Talmon, N. 2020. Participatory Budgeting with Project Interactions. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI '20)*, 386–392.

Jain, P.; Sornat, K.; Talmon, N.; and Zehavi, M. 2021. Participatory Budgeting with Project Groups. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI '21)*, 276–282. ijcai.org.

Lang, J.; and Xia, L. 2016. Voting in Combinatorial Domains. In Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D., eds., *Handbook of Computational Social Choice*, 197–222. Cambridge University Press.

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

Papadimitriou, C. H. 1994. *Computational Complexity*. Addison Wesley.

Peters, D.; Pierczynski, G.; and Skowron, P. 2020. Proportional Participatory Budgeting with Cardinal Utilities. Technical report, arxiv:2008.13276.

Shah, A., ed. 2007. *Participatory Budgeting*. Public Sector Governance and Accountability Series. Washington, DC: The World Bank.

Spakowski, H. 2005. *Completeness for parallel access to NP and counting class separations*. Ph.D. thesis, University of Düsseldorf, Germany.

Yang, Y.; and Wang, J. 2018. Multiwinner Voting with Restricted Admissible Sets: Complexity and Strategyproofness. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI '18)*, 576–582.

Zwicker, W. S. 2016. Introduction to the Theory of Voting. In Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D., eds., *Handbook of Computational Social Choice*, 23–56. Cambridge University Press.