Participatory Budgeting with Donations: The Case of Selective Voters

Philip Lazos, Evangelos Markakis[a,b] and Georgios Papasotiropoulos[c,*]

[a]Athens University of Economics and Business, Greece
[b]Input Output Global (IOG), Greece
[c]University of Warsaw, Poland

Abstract. Participatory budgeting allows citizens to decide how to allocate public funds among projects. Motivated by recent real-world applications in both municipal and blockchain environments, we propose and study a framework where voters can donate additional private funds to enhance their own satisfaction, using cumulative ballots to express preferences. We introduce the first mechanisms for this setting and evaluate them primarily based on the satisfaction of axioms, while also exploring their algorithmic and strategic aspects.

1 Introduction

Participatory budgeting (PB) empowers citizens to decide how a budget is allocated among projects benefiting the public good. Participants vote on project options, and their preferences are aggregated to fund projects while staying within budget. Our work centers on scenarios where the budget is financed by both public funds and contributions from voters wishing to support specific projects.

PB procedures have gained widespread real-world applicability. According to Dias and Júlio [10], over 7000 implementations of PB had taken place worldwide by 2018. Whether at the level of a country, municipality, neighborhood, or even smaller communities, PB is being employed widely to determine budget allocations. To motivate our study, we spotlight the elections of 2019[1] and 2021[2] in the Polish city of Gdynia, where an individual partially financed a cultural performance with personal funds, while district councils used external funds to finance projects like children's games, workshops, pedestrian infrastructure, and community center equipment. These projects were ultimately implemented through a combination of public budget allocations and contributions from individual donors (private or public), enabling funding for projects that would have been impossible solely through the available public budget.

Moving away from traditional PB elections, there has also been a notable surge of interest in PB within blockchain governance systems. For instance, Project Catalyst pools ADA cryptocurrency transaction fees and allocates funds based on the votes of stakeholders. This process repeats a few times per year, having funded more than 1.6k projects (out of more than 7k proposals), with a total cost exceeding $79 million and based on over 2.5 million votes. Similarly, Gitcoin has directed over $60 million to public goods, with 270k supporters backing more than 3.5k projects funded by both public pools and contributions of individuals. In these setups, stakeholders are able to vote on fund allocation, with their votes being weighted by their stake. Our study is directly motivated by these two applications (which are affiliated with two prominent cryptocurrencies, Ethereum and Cardano), and stems from ongoing discussions within blockchain communities on improving PB procedures already implemented in practice. These discussions are also relevant to other blockchain-related entities that use forms of PB, such as DAOs [18].

Our paper fits within the line of work on models for PB with donations. The closest works to ours are by Chen et al. [7], who initiated the study of participatory budgeting models where voters can pledge donations to support projects, and by Wang et al. [19], who applied a similar approach but focused on approval ballots. Both propose rules for their models and evaluate them mainly through specific axioms—a method we also adopt. Crucially, while these studies assume voters are motivated by the community's benefit, we focus on *selective voters*, whose donations are driven by personal satisfaction.

Motivated by the focus on selfish behavior, we also investigate strategic aspects related to donations, which have not been addressed before for such a setting. Aziz and Ganguly [1] examined similar questions but in a setting where the entire budget comes solely from the agents themselves (with no public funds available) and voter utility depends on the total money spent on her approved projects. For a similar framework, an efficient rule with strong incentive and fairness properties was suggested by Brandl et al. [5]. However, our setting fundamentally differs by incorporating a common public budget to be distributed. A shift from the charitable funding perspective to one where agents care about where their money goes—similar to our approach but still without a shared budget—is explored by Aziz et al. [2]. They proposed quasi-linear utilities to capture voter satisfaction, which depends on both their pledged donations and the selected projects, focusing primarily on algorithmic approaches and presenting strong negative results. In a conceptually similar study to ours, Boehmer et al. [4] examine how to assess the performance of losing projects in PB, including measures like reducing their costs; such a reduction could result from donations made by their supporters.

In our work, we examine a framework similar to the one studied by Chen et al. [7] and Wang et al. [19]. The mechanisms proposed in their study are based on the principle that voters make donations in order to reduce public spending. As we will discuss extensively in Section 3, this principle suggests that the voters' donations may ultimately be used towards projects they don't necessarily support, since

* Corresponding Author. Email: gpapasotiropoulos@gmail.com.
[1] bo.gdynia.pl/wp-content/uploads/2021/05/2019-09-05_Raport-podsumowujacy-BO-2019_wersja-zaktualizowana-1-1.pdf
[2] bo.gdynia.pl/wp-content/uploads/2021/12/2021-10-06_Raport-podsumowujacy-BO-2021_aktualizacja.pdf

their preferred projects might already be funded through the public budget or others' donations. This approach may not sufficiently motivate participants to contribute, especially those who aren't driven by altruism or a desire to enhance their public image. In summary, while the approach of Chen et al. [7] is compatible with voters aiming to benefit society, it may not suit selective voters who would like to donate so as to enhance their own satisfaction.

We propose and examine a PB scenario with donations where voters use cumulative ballots [8]. Cumulative balloting extends approval and ordinal ballots, allowing participants to allocate a (virtual) coin among the options. This approach aligns with our motivation from Project Catalyst and Gitcoin, as voters there can use their actual money to indicate preferences over different projects. Cumulative participatory budgeting, despite its appeal and practical use in cities like Strasbourg, Toulouse, and Gdansk, has received limited research attention [16]. For more on the use of cumulative ballots in PB, see the work of Skowron et al. [17]. Moving away from participatory budgeting, there is a broader literature on cumulative ballots in voting environments [3, 6, 9, 11, 13, 14, 15].

Contributions. First, in Section 2, we introduce and analyze an election framework that serves a dual purpose:

- It naturally captures, as a special case, the scenario of allowing donations under the classic PB setting. Therefore, our study applies to traditional PB processes where organizers permit pledging. In this regard, we complement prior work on PB with donations by focusing on scenarios where voters are interested in donating exclusively to their preferred projects.
- It mirrors voting procedures in prominent real-world blockchain systems, where a voter's stake influences her voting power. This aligns with digital governance concepts, making our study directly applicable to cryptocurrency and DAO environments, where our rules are well-suited. Moreover, our work contributes to understanding strategic considerations of participants in these systems.

Then, in Section 4, we propose two rules that make meaningful use of the submitted willingness to donate by each voter; each with its own strengths and weaknesses. To demonstrate their effectiveness, our study begins by incorporating axioms that (i) ensure the alignment of the examined rules with the interests of selective voters, (ii) build on previously proposed axioms to show that allowing donations does not harm the electorate (under various interpretations), and (iii) align with established natural axioms that were not known to be satisfiable in the classic PB setting but are made possible within our framework. We also establish that only one of the rules runs in polynomial time, assuming P≠NP. Furthermore, in Section 6, we investigate strategic aspects of the proposed rules, specifically focusing on when and how voters may act strategically. We show that while various forms of manipulation of the outcome under the suggested mechanisms are theoretically possible, there are instances in which malicious actions are computationally infeasible.

2 Preliminaries

In Section 2.1 we outline the specifics of the model we examine, which generalizes classic PB. As a result, the election rules we propose apply not only to settings that involve monetary components (such as those tailored to blockchain governance that motivated our work; see Section 1) but, importantly, also to classic PB scenarios with donation allowance (such as those studied in [7, 19], and referred to as "PB scenarios" below). In Section 2.2, we discuss and define the axioms by which our rules will primarily be evaluated.

2.1 Formal Model

The input of our problem consists of a set of m candidate projects $P = \{p_1, p_2, \ldots, p_m\}$, a set of n voters $V = \{v_1, v_2, \ldots, v_n\}$ and a limit $L \in \mathbb{R}_{\geq 0}$ on the available *public budget*. Each project $p_j \in P$ has an *implementation cost* $c_j \in \mathbb{R}_{>0}$. Moreover, each voter $v_i \in V$ comes with a total *power* $s_i \in \mathbb{R}_{>0}$, which represents the stake with which the voter enters the system. In blockchain environments (such as those associated with Ethereum and Cardano which, as outlined earlier, motivate our work) an agent's stake in the system is the total number of currency tokens she holds and determines precisely her voting power in decision-making processes. As such, the parameter s_i is inherently linked to a monetary value. Having specified s_i, a voter v_i chooses before the election to split it in any way she prefers into a *voting weight* $w_i \in \mathbb{R}_{>0}$ and a *contribution parameter* $d_i \in \mathbb{R}_{\geq 0}$ that models the (maximum) amount of money she is willing to donate. Hence, it should hold that $w_i + d_i = s_i$. Since d_i is an upper bound that v_i declares for her potential contribution, she may ultimately be asked to contribute less or even 0.

Clearly, this setup closely aligns with digital democratic systems, particularly in the context of digital finance. It is important to also note that simply by setting $w_i = 1$ and $d_i = 0$, for each voter v_i, we uncover the classic PB model as being studied in the computational social choice literature. This shows the direct connection between our model and traditional PB frameworks. Moreover, all our positive results, along with the negative ones that hold for voters of unit (or pairwise equal) weights, directly apply there.

The (cumulative) *ballot* of voter v_i is defined as a function $u_i : P \to \mathbb{R}_{\geq 0}$ such that $\sum_{p_j \in P} u_i(p_j) = 1$. Intuitively, the value of $u_i(p_j)$ determines the fraction of the weight owned by voter v_i that she would like to assign to p_j to indicate the level of support towards it. At the same time, u_i is viewed as specifying the utility of v_i for each project. Ultimately, the ballot of v_i is scaled by the weight w_i that she possesses. Therefore, the *support* of voter v_i towards project p_j is given by $\sigma_i(p_j) = u_i(p_j) \cdot w_i$. This support will be used to determine which projects will be granted funding. Evidently, $\sum_{p_j \in P} \sigma_i(p_j) = w_i$. In contrast to classic cumulative voting, the total support voters can distribute among the projects might differ between voters. In traditional cumulative voting, each voter splits a fixed number of points among the candidates. In our model, the total amount to be distributed (which is w_i for voter v_i) varies depending on the weight each voter has chosen to participate with, which, in turn, is determined by her stake in the system (and her donation).

A voter $v_i \in V$ *supports* a project $p_j \in P$ if $u_i(p_j) > 0$ (equivalently if $\sigma_i(p_j) > 0$). For a project p_j we denote by $A(p_j)$ the set of voters who support it. Moreover, $U(p_j)$ is the *total support* that the voters in V allocate to project p_j, i.e., $U(p_j) = \sum_{i \in [n]} \sigma_i(p_j)$. We allow the extension of these notations to bundles of projects, by taking project-wise summation. The donations that the voters may be asked to make (and which are guaranteed to not exceed d_i for each voter v_i) are affecting a voter's acquired utility only implicitly via the set of accepted projects. If a set T of projects is selected for implementation, the final utility of voter v_i is precisely equal to $\sum_{p_j \in T} u_i(p_j)$. This is independent of her weight and contribution reflecting the idea that a donation represents a monetary amount the voter is willingly and freely giving away, in analogy to [7, 19]. We denote by $\mathbf{P}$ the set of projects P together with their costs $c = (c_j)_{j \in [m]}$ and by $\mathbf{V}$ the set of voters together with the tuple (w, d, u) which corresponds to the tuple of vectors that are associated with the voters' preferences, namely $w = (w_i)_{i \in [n]}, d = (d_i)_{i \in [n]}$, and $u = (u_i)_{i \in [n]}$. A *generalized budgeting scenario*, or simply a *scenario*, is a tuple $S = (\mathbf{P}, \mathbf{V}, L)$. We refer to scenarios of pairwise equal voting weights as *PB scenarios*.

An *aggregation method* or *election rule* is a procedure F that given a generalized budgeting scenario S, selects a bundle of projects $B \subseteq P$ to be implemented, an m-dimensional vector β such that $\beta_j \in \mathbb{R}_{\geq 0}$ indicates how much from the public budget will be spent towards the implementation of project p_j and a mapping δ such that $\delta_i(p_j) \in \mathbb{R}_{\geq 0}$ indicates the amount of money that v_i is being asked to contribute towards p_j. A *solution* $F(S) = (B, \beta, \delta)$ is feasible for a scenario $S = (\mathbf{P}, \mathbf{V}, L)$ if it simultaneously satisfies the following:

- No voter should be asked to spend more than the amount of money she declared that she is willing to contribute, i.e., $\sum_{p_j \in B} \delta_i(p_j) \leq d_i, \forall v_i \in V$.
- The public budget spent for all funded projects should not exceed the public budget limit, i.e., $\sum_{p_j \in B} \beta_j \leq L$.
- The total amount of money contributed towards any project $p_j \in P$ from both the public budget and the voters' contributions is equal to c_j if $p_j \in B$, and 0 otherwise.

For a project p_j, we denote by $D_j(F(S))$ the donors of p_j, i.e., the voters of a scenario S selected by the aggregation method F to donate to p_j. Hence, it includes every v_i for which $\delta_i(p_j) > 0$ under the solution $F(S)$. For notational convenience, we will sometimes use $F(S)$ to denote the bundle B (instead of the tuple (B, β, δ)).

2.2 Axioms

We will now present intuitions and formal definitions of the metrics of evaluation of our methods, namely, (i) *Donation No-Harm*, (ii) *Preference-Donation Alignment*, (iii) *Support (Redistribution/Increase) Monotonicity* and (iv) *Donation-Support Monotonicity*.

► The axiom of *Donation No-Harm* ensures that allowing donations will not make any voter less satisfied, regardless of whether the voter donated herself or not. This ensures that wealthy voters cannot influence the election in a way that decreases the satisfaction of voters who rely on public budget spending for projects they like, making it a principle of democratic character. It was the principal axiom in the works by Chen et al. [7] and Wang et al. [19], where the authors primarily aimed to show that allowing donations should not result in greater participant dissatisfaction than in a framework without donations. For further motivation we refer to the aforementioned works; for us its role is primarily to position our work within the existing landscape of PB rules with donations.

Axiom 1: Donation No-Harm. *An aggregation method F is said to satisfy Donation No-Harm if in any two scenarios S and S' where the contribution parameter d_i equals 0 for every voter v_i under S, while being positive for at least one voter in S' (with all other parameters being equal), it holds that $u_i(F(S)) \leq u_i(F(S'))$, for every v_i.*

► An axiom that distinguishes our work from previous literature is the axiom of *Preference-Donation Alignment*. At its core, this axiom asserts that a voter should not be compelled to contribute to projects she does not support. Voters who are conscious of where their funds are allocated would not willingly participate in PB elections where their contributions might go towards projects they oppose. Therefore, since a solution includes the information about which voter donates to which projects, our goal is to ensure that each voter's contribution is allocated only to projects she supports.

Axiom 2: Preference-Donation Alignment. *An aggregation method F is said to satisfy Preference-Donation Alignment if for every scenario S and every project p_j selected for implementation under $F(S)$ it holds $D_j(F(S)) \subseteq A(p_j)$; meaning that only voters supporting a certain project might be asked to pay for it.*

This axiom is particularly relevant in certain scenarios, especially those motivated by the applications driving our work. Specifically, consider situations where the voting rule may not be easily understood by all participants, or where participants seek simple, clear assurances of the rule's quality without the need to verify its underlying reasoning themselves. Then, a rule satisfying Preference-Donation Alignment can be persuasive to voters, potentially leading to broader acceptance. Various closely related axioms—like allowing for donations to a project you do not support of, but only if it results in the election of projects you favor—can be defined and analyzed. However, these may trade off the simplicity of validation, as voters might still need to understand the mechanism's specifics to feel confident about how their donation was used. Therefore, while Preference-Donation Alignment is not the only axiom that aligns with the goals of a selective voter, it is a natural and well-suited starting point.

► The axiom of *Support Monotonicity* is related to the support that a voter v_i assigns to a project. Recall for a project p_j, this equals to the ballot $u_i(p_j)$ multiplied by the voter's weight w_i. It ensures that increasing a voters' support for a winning project (without increasing the support of any other project) does not diminish its chances of being selected. More precisely, in our context, v_i can increase the support σ_i towards p_j in two ways:[3]

- by *reallocating* u_i among projects, such that only p_j gains increased support while the support to every other project either decreases or remains unchanged,
- by *augmenting* s_i, and consequently w_i (without altering d_i) to enhance the overall voting power of v_i and then increase the support exclusively towards p_j while keeping the rest unchanged, as explained in the example that follows.

This distinction follows directly from our model but contrasts with the type of cumulative voting systems for PB discussed by Skowron et al. [17], where only the first option applies. This is because in classic PB systems it holds $w_i = 1$ for every voter v_i, and augmentation of s_i is not feasible. For better understanding of how the support for a project can increase without raising the support for others, we present an illustrative example.

Example 1. *Consider a voter with a voting weight of 2 who submits the following ballot on 4 projects:* $(1/10, 4/10, 2/10, 3/10)$. *This results in the following support vector:* $(0.2, 0.8, 0.4, 0.6)$. *Now, let us consider increasing the support for the first project to* 0.4. *Two indicative support vectors that are in line with this increment are the following:* $(0.4, 0.7, 0.4, 0.5)$ *and* $(0.4, 0.8, 0.4, 0.6)$. *The first is made possible by a redistribution of the ballot. Namely, if the voting weight is kept to* 2, *the voter could submit the ballot* $(4/20, 7/20, 4/20, 5/20)$. *For the second vector, if the weight increases to* 2.2, *e.g., through an exogenous increase of her stake, then the voter could submit* $(4/22, 8/22, 4/22, 6/22)$. *This yields an increase in the support for the first project, leaving the support towards the remaining unchanged.*

We distinguish between two variants of the axiom based on how a voter v_i can increase the support towards a specific project p_j: *Support-Redistribution Monotonicity* corresponds to the case when the support increases due to redistributing her ballot (the first option discussed above), and *Support-Increase Monotonicity* corresponds to the case when the support rises because of an increase in voting weight (via power) as the voter acquires larger stake.

[3] A third method—redistributing s_i to decrease d_i and increase w_i—could render a winning project unaffordable, thus is unsatisfiable under any rule.

Axiom 3a: Support-Redistribution Monotonicity. *Consider two arbitrary scenarios S and S' such that for exactly one voter v_i and a project p_ℓ it holds $u_i(p_\ell) < u_i'(p_\ell)$, and $u_i(p_k) \geq u_i'(p_k), \forall k \neq \ell$, and with all other parameters of the scenarios being equal. An aggregation method F is said to satisfy Support-Redistribution Monotonicity if whenever $p_\ell \in F(S)$ it also holds that $p_\ell \in F(S')$.*

Axiom 3b: Support-Increase Monotonicity. *Consider two arbitrary scenarios S and S' such that for exactly one voter v_i and a project p_ℓ it holds $\sigma_i(p_\ell) < \sigma_i'(p_\ell)$, where σ_i, σ_i' correspond to the support in S and S' respectively. Suppose it also holds that $w_i < w_i'$ and that $\sigma_i(p_k) = \sigma_i'(p_k), \forall k \neq \ell$ (and all other parameters of the scenarios remain equal). An aggregation method F is said to satisfy Support-Increase Monotonicity if for every such a pair of scenarios S and S', whenever $p_\ell \in F(S)$ it also holds that $p_\ell \in F(S')$, for the considered project p_ℓ.*

► We would like the external increase of the potential donation of a voter to be unable to result in the election of a worse bundle of projects. The axiom of Donation-Support Monotonicity ensures that increasing a voter's contribution (keeping all other parameters, including voting weights, unchanged) can only benefit society: it cannot lead to the election of a bundle with lower total support.

Axiom 4: Donation-Support Monotonicity. *An aggregation method F is said to satisfy Donation-Support Monotonicity if in any two scenarios S and S' where their only difference comes from the contribution parameter of a voter v_i, that is d_i in S and $d_i' > d_i$ in S', it holds that $U(F(S')) \geq U(F(S))$.*

The rules we focus on aim to align with selfish voters' willingness to donate, and we refer to them as Donation-Alignment (DA) rules.

3 Prelude to Donation-Alignment Election Rules

One approach to designing an aggregation method F is to define and solve an optimization problem, e.g., as done in [2]. The utilitarian objective, which is among the most explored desiderata in the PB literature, forms the basis of our work, with egalitarian objectives and proportionality guarantees emerging as natural directions for future research. We focus on maximizing the voters' support, which is the most natural starting point. This approach aligns with the objectives outlined by [7], the rule currently used in Project Catalyst, and one of the most commonly applied rules in real-world PB processes. The support of each voter is considered as additive and is based on which of the projects she supports are selected. Specifically, given a scenario S and an aggregation method F, we say that a voter v_i assigns a support of $\sigma_i(F(S)) = \sum_{p_j \in B} \sigma_i(p_j) = w_i \sum_{p_j \in B} u_i(p_j)$ to the bundle B that has been selected by applying F in S. We thus can set as an optimization objective to maximize $\sum_{i \in [n]} \sigma_i(F(S))$ among all feasible solutions.

Notably, we will show that methods that are based on solving this optimization problem cannot satisfy the basic axioms suggested in Section 2.2—an absolute drawback for our purposes. Before stating this result, we highlight that solving the optimization problem could be done according to the following simple reduction to the knapsack problem, together with the application of any of the well known polynomial-time approximation schemes or parameterized algorithms for knapsack: create one item for each project, set the knapsack capacity to $L + \sum_{i \in [n]} d_i$ and set the utility that an item p_j would bring to $\sum_{i \in [n]} \sigma_i(p_j)$. Then, any (exact or approximate) solution to the created knapsack instance corresponds to a feasible solution for the initial PB instance. The coming result essentially shows that the main requirements we put forward in Section 2.2 contradict with the objective of maximizing the total electorate's support—which should not come as a surprise, given that the axioms are tailored to selective agents. Its full proof, along with any other omitted proofs or parts thereof, is deferred to the full version of our work.

Theorem 1. *Any mechanism that returns the bundle that maximizes the total voters' support up to any finite, positive multiplicative approximation factor fails Donation No-Harm and Preference-Donation Alignment. This result holds even for PB scenarios.*

Returning to the earlier reduction to knapsack, we also exhibit a reverse direction where knapsack can be straightforwardly reduced to our problem, establishing NP-hardness.

Theorem 2. *It is NP-hard to maximize the objective of the voters' support while satisfying Donation No-Harm and Preference-Donation Alignment. This result holds even for PB scenarios.*

Moving forward, we discuss two main rules as representatives among families of rules proposed by Chen et al. [7], adapted to our framework. It was established that these methods satisfy Donation No-Harm. However, they are not suitable for settings with selective agents, as will become apparent shortly.

The first rule, to be called the *Greedy rule*, employs a subroutine where, iteratively, a project p is added to the winning bundle C' towards maximizing the total voters' support for $C' \cup \{p\}$, ensuring feasibility at each step. The main component of the algorithm proceeds by initially setting $C' = \emptyset$ and applying this subroutine to the instance without considering contributions. If, by the end of the process, the voters express willingness to contribute to any project in the set of the selected ones, then the global budget is increased by the analogous donations. The subroutine repeats with this new budget. This continues until no further projects can be added.

The second rule, referred to as the *Pareto rule*, begins by selecting the optimal bundle in terms of electorate's support that is feasible without any donations, say B^*. It then creates a collection of bundles T that includes B^* as well as all bundles that are feasible when donations are considered, provided they Pareto dominate B^*—a bundle is said to *Pareto dominate* B^* if it receives at least the same support as B^* by all voters and strictly more by at least one voter. The rule outputs the bundle of maximum support among those in T.

The following result highlights a drawback of the two discussed rules, when applied in scenarios with selective voters, and it essentially motivates our study. However, we emphasize that such rules were not specifically designed to accommodate selective voters, so this observation should not be seen as particularly surprising.

Observation 3. *Both the Greedy and the Pareto rule fail Preference-Donation Alignment, even for PB scenarios.*

4 Donation-Alignment Election Rules

In this section we present two rules designed to be applicable in scenarios where voters are selective. Their evaluation with respect to the axioms from Section 2.2 appears in Section 5 and their strategic aspects are being explored in Section 6. Both rules could fit either for traditional PB applications or within blockchain-based systems. We highlight that cumulative ballots align with the platforms that motivated our study, though importantly, our rules also apply directly to formats like *approval* or *cardinal* ballots.

In order to satisfy the fact that any voter who gets satisfaction only because of the public budget will not get worse because of the appearance of donations from others (Donation No-Harm), both of our rules start by considering a solution that is affordable only by the

$B^* \leftarrow$ bundle maximizing total voters' support under L.
$T \leftarrow \{B^*\}$.
for *each possible bundle* $B \subseteq P$ **do**
  Compute the utility improvement against $B^*, \forall$ voter.
  if *the utility improves for at least one voter without decreasing for the rest and* B *is affordable by the public budget plus contributions from strictly benefiting voters* **then** Add B to the collection T.
Return argmax$\{U(B) : B \in T\}$.

Algorithm 1: DA-Pareto Mechanism

Sort P in non-increasing order of support-to-cost ratio.
Initialize the remaining public budget $R \leftarrow L$.
Initialize the set of selected projects $T \leftarrow \emptyset$.
for *each project* p_j *in sorted order* **do**
  if $c_j \leq R$ **then**
    Allocate public funds to p_j.
    $T \leftarrow T \cup \{p_j\}$.
    Update remaining public budget $R \leftarrow R - c_j$.
for *each project* $p_j \notin T$ *in sorted order* **do**
  if $c_j \leq R$+*remaining contributions of supporters of* p_j **then**
    Allocate public funds and supporting voters' contributions to cover c_j and add p_j to T.
    Update R and voters' available contributions.
Return set of selected projects T.

Algorithm 2: DA-Greedy Mechanism

public budget. After that, we only ask voters to contribute towards projects that they would like to support and that have not been selected for implementation by the public budget; essentially this idea serves the purpose of taking donations from a voter only if this will result in strictly improving her utility (in line with the Preference-Donation Alignment axiom).

Under the Pareto rule proposed by Chen et al. (2022), voters might end up paying more than in the initial solution where only the public budget was being considered (i.e., more than donating 0 for B^*), even if their utility remains unchanged, to improve another voter's utility. As observed in Section 3, this method does not align with Preference-Donation Alignment. We propose a variant, *Donation-Alignment Pareto* (*DA-Pareto*), addressing this concern. DA-Pareto (Algorithm 1) starts by selecting the optimal, in terms of voters' support, bundle B^*, which can be purchased within the available public budget. It then identifies a collection T of potentially winning bundles that includes B^* as well as all bundles that are feasible when incorporating the voters' donations, and dominate B^*—a bundle is said to dominate B^* if it receives strictly more support by at least one voter without receiving less by any, while only those voters who benefit pay more than they would for B^* (i.e., a non-zero amount). These conditions can be checked via a linear system. Among the bundles in T, the rule returns the one maximizing the electorate's support.

Our second suggestion, *Donation-Alignment Greedy* (*DA-Greedy*), operates in two phases (Algorithm 2). In the first, it allocates only the public funds to projects based on their support-to-cost ratio, sorted in non-increasing order (recall that the support of p_j equals $U(p_j) = \sum_{i \in [n]} \sigma_i(p_j)$). Projects are included in the solution iteratively until no further project can be funded by the remaining public budget. In the second phase, the mechanism evaluates each remaining project in descending order of their support-to-cost ratio. For each project p_j, it determines whether its cost can be covered by the remaining public budget augmented by the contributions from voters who support p_j. If affordable, the mechanism spends as much of the remaining public budget as possible on the considered project and covers the remaining cost through donations from supporting voters, aiming at equal contribution among them (or utilizing all available funds from certain voters). This process is repeated, using the remaining contributions, for each subsequent project, until all projects have been considered. Ties are broken arbitrarily.

Table 1: Axiomatic properties of the proposed election rules.

Axioms	Election Rules: DA-Pareto	DA-Greedy
Donation No-Harm	✔	✔
Preference-Donation Alignment	✔	✔
Support-Increase Monotonicity	✔	✔
Donation-Support Monotonicity	✔	✘
Support-Redistribution Monotonicity	✘	✘
Polynomially Computable (assuming P≠NP)	✘	✔

5 Axiomatic Results

We now discuss the properties of our rules. Our results are summarized in Table 1. DA-Pareto satisfies all but one of the axioms but does not have polynomial runtime. DA-Greedy guarantees polynomial computability but sacrifices the satisfaction of an extra axiom.

Theorem 4. *DA-Pareto satisfies Donation No-Harm, Preference-Donation Alignment, Support-Increase Monotonicity, Donation-Support Monotonicity, but fails Support-Redistribution Monotonicity.*

Proof. We split the proof in parts, each referring to the satisfiability of a different axiom. Missing parts are deferred to the full version.

Preference-Donation Alignment: This is satisfied by the definition of the rule. All the possible bundles in T considered by the Pareto rule as potential solutions do not require voters to fund projects they do not support. This follows since the bundles in T are either B^*, which would be funded by public funds, or any bundle B that dominates B^*, where each $p_j \in B$ would be funded by (perhaps some public funds and) voters who support it.

Support-Redistribution Monotonicity: Consider the instance depicted in the following table, where the entry corresponding to voter v_i and project j depicts $\sigma_i(p_j)$. Specifically, $L = 2$ and $w_1 = 3.1$, $u_1 = (1/3.1, 0.6/3.1, 1.5/3.1)$ and $w_2 = 0.2 + \varepsilon$, $u_2 = (0, 0.2/0.2+\varepsilon, \varepsilon/0.2+\varepsilon)$. Moreover $s_1 = 3.1$ and $s_2 = 0.2 + 2\varepsilon$.

$L = 2$	parameters	Project 1 $c_1 = 2$	Project 2 $c_2 = 2$	Project 3 $c_3 = 2 + \varepsilon$
v_1	$d_1 = 0$	1	0.6	1.5
v_2	$d_2 = \varepsilon$	0	0.2	ε

It holds that B^* consists of the set that only includes Project 1, as it is the one maximizing total voters' support, between the two projects that are affordable from the global budget. However, the bundle that consists of Project 3 is affordable by the public budget increased by the donation of v_2 and it results to strictly greater satisfaction to v_2 and no worse for v_1 compared to the previously considered bundle, so it will be winning under DA-Pareto. Say then that v_1 redistributes her ballot, now declaring the ballot $(0.5/3.1, 0.6/3.1, 2/3.1)$, which results to the following support vector: $(0.5, 0.6, 2)$. Notice that this change increased the support towards the winning project, so, according to Support-Redistribution Monotonicity, the third project should remain in the winning bundle. However, after this change, B^* contains the second project, and the solution that only contains the third one doesn't dominate B^* anymore since v_2 prefers Project 2 to Project 3. □

On the negative side, DA-Pareto requires an exponential enumeration, at least in its most straightforward implementation. Furthermore, it needs to solve an NP-hard problem, since its first step essentially involves solving a knapsack instance.

Theorem 5. *The first step of DA-Pareto, computing B^*, amounts to answering an NP-hard question, even for PB scenarios. Furthermore, the second step of DA-Pareto, its for-loop, involves an exponential enumeration. DA-Greedy runs in polynomial time.*

Our result on the axiomatic properties of the greedy rule follows.

Theorem 6. *DA-Greedy satisfies Donation No-Harm, Preference-Donation Alignment, Support-Increase Monotonicity, but fails Support-Redistribution Monotonicity, Donation-Support Monotonicity.*

Proof. We split the proof in parts, each referring to the satisfiability of a different axiom. Missing parts are deferred to the full version.

Preference-Donation Alignment: Say that p_j is a project that belongs to the winning bundle under DA-Greedy, to be denoted by F. This is either funded exclusively by the global budget or voters will contribute as well. In the first case, the axiom holds trivially. Regarding the second, it simply suffices to observe that in the first round, no voter is being asked to donate, whereas in the second, only voters supporting a project may contribute, so $D_j(F(S)) \subseteq A_j(S)$.

Donation-Support Monotonicity: Consider the instance depicted in the following table, where the entry corresponding to voter v_i and project j depicts $\sigma_i(p_j)$. The voting weights of the two voters are, respectively $w_1 = 3 - \varepsilon$ and $w_2 = 1$, therefore voters' ballots could be expressed as $u_1(p_j) = \sigma_1(p_j)/3-\varepsilon$ and $u_2(p_j) = \sigma_2(p_j)$.

$L = 0$		Project 1	Project 2	Project 3
	parameters	$c_1 = 6$	$c_2 = 4$	$c_3 = 4$
v_1	$d_1 = 5$	$2 - \varepsilon$	0.5	0.5
v_2	$d_2 = 3$	0	0.5	0.5

Note that $L = 0$. The first project has a better ratio of total support-to-cost, so it will be considered first. However, it isn't affordable as its supporter can contribute at most 5 dollars, i.e. 1 less than the cost of the project. The rest of the projects are all affordable since all two voters support them and together they have a total budget of 8 which equals the cost of those two projects. Hence, the solution under DA-Greedy in the given instance receives a total support from the electorate that is equal to 2 by selecting projects 2 and 3. Suppose now that v_1 increases her potential donation d_1 from 5 to 6. Project 1 will now be affordable, since the first voter, a supporter of this project, has a total budget equal to the cost of the project. After selecting the first project, the budget of v_1 is exhausted, and given that the budget of v_2 isn't sufficient for buying any project, the solution after the increase of the budget of v_1 now receives a total support of $2 - \varepsilon$. □

In summary, among the rules we suggest, there is one (DA-Pareto) that satisfies most of the axioms set forth in Section 2.2. However, it cannot ensure polynomial running time. Conversely, the rule that guarantees polynomial computability (DA-Greedy) is slightly weaker in terms of axiom satisfaction, but still performs undoubtedly better compared to what has been known in the literature for scenarios involving selective agents, as it also exhibits sufficiently strong axiomatic properties. These findings align perfectly with Theorem 2, which shows that no polynomial-time computable rule can satisfy the desired axioms while also providing sufficient guarantees with respect to electorate's support. One of our proposed rules sacrifices computational efficiency to ensure certain support guarantees, while the other prioritizes efficiency at the expense of support.

Table 2: Strategic aspects of the proposed election rules. The negative statements regarding polynomial computability hold under P≠NP.

	DA-Pareto	DA-Greedy
manipulable by donation misreport	✔	✔
manipulable by ballot misreport	✔	✔
manipulation by donation in poly-time	✖	✔
manipulation by ballot in poly-time	✖	✖
election control in poly-time	✖	✖

6 Strategic Aspects

In this section we focus on strategic aspects of the proposed setting. To begin with, we will hereinafter assume that voters are able to misreport their preferences. Recall that the input given to an aggregation mechanism by voter v_i corresponds to the triplet (w_i, d_i, u_i). For notational simplicity, since s_i is known to the mechanism and w_i can be inferred from d_i, we treat the voter's input as the tuple (u_i, d_i), in words, her ballot vector (which is then scaled by w_i to form her support) as well as her contribution parameter. Suppose now that voter v_i, although having some true preferences (u_i, d_i), can choose to submit (b_i, q_i) instead, where it should obviously hold that the declared weight of v_i equals $s_i - q_i$. The tuple (b_i, q_i) might or might not be equal to (u_i, d_i). In the former case, we say that we are in a *truthful scenario*. In the remainder we mainly focus on the following:

> *When is it rational and computationally feasible for a voter to misreport her true preferences towards maximizing her utility from the resulting outcome?*

We begin by illustrating that the answer is not trivial in our setting.

Observation 7. *There exist instances where voters are better off donating than having large voting power, while in others, they are better off maximizing their voting power, under any reasonable mechanism.*

Underreporting the contribution, i.e., expressing $q_i < d_i$, directly increases the voting weight of voter v_i since s_i is fixed. As a result, our findings on donation misreporting do not apply to settings where all voters have equal weight—such as the classic PB model, where monetary contributions do not influence voting power. In contrast, our results on ballot misreporting hold for that model as well.

Observation 7 motivates the study of strategic aspects. A first negative result for a large family of rules, including those we proposed in Section 4, follows. It shows that not only does acting truthfully fail to result in a Nash Equilibrium for the voters, but also that the Price of Anarchy, defined as the ratio between the total voters' utility in the optimal (centralized, non-strategic) solution and in the worst (in terms of total voters' utility) equilibrium is unbounded.

Theorem 8. *The truthful scenario is not always a Nash Equilibrium, for any deterministic aggregation mechanism that decides for funding based on the voters' support on the projects. Moreover, the Price of Anarchy for such mechanisms tends to infinity as n grows, even for PB scenarios.*

In response, we now focus on questions around manipulation and control of elections. First and foremost, we investigate whether a voter can misreport her preferences (either through the declared ballot or donation) to increase her utility under the aggregation methods we proposed. We also examine whether such manipulation can always be done in polynomial time, since, even if a manipulation is theoretically possible, what matters is whether such actions can be efficiently determined. Moreover, we explore whether a controller, aiming to enforce a specific outcome by influencing the set of voters, can achieve this in polynomial time. Table 2 summarizes our findings. The main concepts of this section are formally defined below.

Definition 1. *We say that a rule F is manipulable by misreporting donations if there is a scenario in which a voter v_i can gain more utility from the outcome of F by claiming willingness to contribute $q_i < d_i$ (while keeping u_i unchanged). We say that a rule F is manipulable by misreporting ballots if there is a scenario in which a voter v_i can gain more utility from the outcome of F by casting a cumulative ballot $b_i \neq u_i$ (while keeping d_i unchanged).*

The following result shows that both rules are manipulable, and this manipulation can occur through both actions.

Theorem 9. *DA-Pareto and DA-Greedy are manipulable by misreporting donations. DA-Pareto and DA-Greedy are also manipulable by misreporting ballots, even for PB scenarios.*

Proof. We will prove the statements for DA-Pareto and the proof for DA-Greedy is deferred to the full version of our work. Towards proving that the rule is manipulable by misreporting donations, consider the instance that appears below, where an entry of the table corresponding to voter v_i and project j depicts $u_i(p_j)$.

$L = 0$		Project 1	Project 2
	parameters	$c_1 = 3$	$c_2 = 5$
v_1	$s_1 = 6$	0.75	0.25
v_2	$s_2 = 1$	0.1	0.9

First, say that $d_1 = 5$ and $d_2 = 0$, so both voters vote with a weight of one in the truthful scenario. Then, the only feasible bundle affordable by the public budget is the empty one and both bundles $\{p_1\}$ and $\{p_2\}$ dominate it while being affordable by the budget of v_1. Additionally, $\{p_1, p_2\}$ isn't feasible. Then $\{p_2\}$ will be selected as the winning bundle because $U(p_2) > U(p_1)$. Consider now the case where v_1 submits a non truthful contribution parameter $q_1 = 3 < d_1$. In turn, v_1 votes with a weight of 3 and $\sigma_1 = (2.25, 0.75)$. Then, only $\{p_1\}$ is a feasible solution, which, again, dominates the empty one. This solution gives to v_1 more utility than when reporting d_1 simply because $u_1 = (0.75, 0.25)$. So, the decrease of her donation resulted in a better for her outcome.

We now move to proving that DA-Pareto is also manipulable by misreporting ballots, even for PB scenarios. Consider the following instance, where $s_1 = 4.1, s_2 = 3.5, w_1 = w_2 = 3.1$ and $u_1 = (1/3.1, 0, 1/3.1, 1.1/3.1)$ and $u_2 = (0, 1.1/3.1, 2/3.1, 0)$. Say that the entry of the table corresponding to voter v_i and project j depicts $\sigma_i(p_j)$.

$L = 1$		Project 1	Project 2	Project 3	Project 4
	parameters	$c_1 = 1$	$c_2 = 1$	$c_3 = 1.4$	$c_4 = 10$
v_1	$d_1 = 1$	1	0	1	1.1
v_2	$d_2 = 0.4$	0	1.1	2	0

In this scenario, the best bundle affordable by the public budget is $\{p_2\}$. With donations, feasible solutions that dominate $\{p_2\}$ are $\{p_3\}, \{p_1, p_2\}, \{p_1, p_3\}, \{p_2, p_3\}$, with $\{p_2, p_3\}$ having the highest support and winning under DA-Pareto. The utility that v_1 gets is then equal to $1/3.1$. If v_1 instead submits $b_1 = (2/3.1, 0, 0, 1.1/3.1)$, the best bundle under the public budget is $\{p_1\}$. Feasible solutions that dominate it are $\{p_1, p_2\}, \{p_1, p_3\}$, with $\{p_1, p_3\}$ having the maximum total support, increasing the satisfaction of v_1 to $2/3.1$. □

We will now show that there are instances where it is computationally infeasible for a voter to determine whether misreporting her utilities could lead DA-Pareto or DA-Greedy to return a bundle that is more favorable to her than the outcome based on her truthful preferences, unless P=NP. We call U-MANIP the relevant computational problem as follows: *Given a specific voter (manipulator), can she misreport her ballot to achieve a utility of at least t, for a given value t, from the outcome of the examined rule*? For DA-Pareto, where computing the outcome is already NP-hard, studying this manipulation problem is only relevant in instances where winning bundles can be computed efficiently.

Theorem 10. *Under DA-Greedy, it is NP-hard to solve* U-MANIP. *The same holds for DA-Pareto, and this is even in cases where the winning bundle under the rule can be computed in polynomial time, specifically when all projects have identical costs. Both results hold even for PB scenarios.*

Unlike misreporting ballots, manipulation through donations is computationally easier. Given that there are instances (of non-zero contributions) where the outcome of DA-Pareto is already computationally intractable (Theorem 5) we focus exclusively on DA-Greedy.

Theorem 11. *A voter can determine the optimal contribution to maximize her utility under the DA-Greedy mechanism in polynomial time, provided that the rest of the parameters remain fixed.*

Strategic Election Control. We conclude with a brief note on control problems—a prevalent research area within computational social choice [12] that is relevant to the questions of the section. These problems involve a controller attempting to enforce a certain outcome by affecting the election components, most commonly by adding or deleting voters or candidates. Here, we focus on altering the set of voters. The definition of a variant involving addition or deletion of candidates is not straightforward in this context as the precise set of candidates must be pre-specified for voters to submit their cumulative ballots. Even in the single-winner setting and with no donations, both problems of controlling the outcome by adding or deleting voters are NP-hard for the Plurality voting rule under cumulative ballots [15]. Given that DA-Greedy and DA-Pareto would produce outcomes identical to Plurality in such scenarios, the relevant computational problems are also NP-hard under the examined rules.

7 Outlook

Our work complements the literature on PB with donations by focusing on selective voters—those interested in donating solely to enhance their own satisfaction. We introduced rules tailored to this setting and demonstrated their effectiveness by proving that they satisfy solid axiomatic guarantees. Motivated by the premise that voters are driven by self-interest rather than charitable motives, we also explored the strategic aspects of the PB framework, focusing on axiomatic and algorithmic questions related to manipulability, and also presented findings on game-theoretic issues and strategic control.

Our model is intentionally centered around frameworks already used in practice. Devising, formulating, and analyzing models under different voting formats or utilities is a valuable direction for future work. Our mechanisms can be adapted to settings with approval limits or ballots allowing approval, disapproval, and abstention. Questions around bribery [12] also form an area for future investigation. Another critical direction is the experimental evaluation of our rules using data either from traditional PB settings or from the blockchain domain. Finally, we note that introducing donations in PB may raise concerns about power asymmetries. Such concerns can be mitigated—at least to some extent—through careful model design and appropriate rule selection as certain axioms may offer a first safeguard against unfairness. Additional safety nets that could be integrated are donation caps and mechanisms inspired by quadratic voting. We believe that fairness considerations should be the next step in the context of PB with donations, and it is an interesting question how to define proportionality notions in the presence of donations.

Acknowledgements

This work has been supported by the H.F.R.I call "Basic research Financing (Horizontal support of all Sciences)" under the National Recovery and Resilience Plan "Greece 2.0" funded by the European Union – NextGenerationEU (H.F.R.I. Project Number: 15877). Part of this work was done while G. Papasotiropoulos was a research intern at Input Output Global (IOG). G. Papasotiropoulos has been supported by the European Union (ERC, PRO-DEMOCRATIC, 101076570). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

European Research Council
Established by the European Commission

References

[1] H. Aziz and A. Ganguly. Participatory funding coordination: Model, axioms and rules. In *Proceedings of the International Conference on Algorithmic Decision Theory*, pages 409–423, 2021.

[2] H. Aziz, S. Gujar, M. Padala, M. Suzuki, and J. Vollen. Coordinating monetary contributions in participatory budgeting. In *Proceedings of the International Symposium on Algorithmic Game Theory*, pages 142–160, 2023.

[3] S. Bhagat and J. A. Brickley. Cumulative voting: The value of minority shareholder voting rights. *The Journal of Law and Economics*, 27(2): 339–365, 1984.

[4] N. Boehmer, P. Faliszewski, L. Janeczko, D. Peters, G. Pierczyński, S. Schierreich, P. Skowron, and S. Szufa. Evaluation of project performance in participatory budgeting. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 2678–2686, 2024.

[5] F. Brandl, F. Brandt, M. Greger, D. Peters, C. Stricker, and W. Suksompong. Funding public projects: A case for the nash product rule. *Journal of Mathematical Economics*, 99:102585, 2022.

[6] R. Cagliero, F. Bellini, F. Marcatto, S. Novelli, A. Monteleone, and G. Mazzocchi. Prioritising cap intervention needs: An improved cumulative voting approach. *Sustainability*, 13(7):3997, 2021.

[7] J. Chen, M. Lackner, and J. Maly. Participatory budgeting with donations and diversity constraints. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 9323–9330, 2022.

[8] A. T. Cole Jr. Legal and mathematical aspects of cumulative voting. *South Carolina Law Review*, 2:225, 1949.

[9] D. A. Cooper. The potential of cumulative voting to yield fair representation. *Journal of Theoretical Politics*, 19(3):277–295, 2007.

[10] N. Dias and S. Júlio. The next thirty years of participatory budgeting in the world start today. *Hope for democracy*, 30:15–34, 2018.

[11] B. Fain, A. Goel, and K. Munagala. The core of the participatory budgeting problem. In *Proceedings of the Conference on Web and Internet Economics*, pages 384–399, 2016.

[12] P. Faliszewski and J. Rothe. Control and bribery in voting. In *Handbook of Computational Social Choice*, pages 146–168. 2016.

[13] R. Freeman, D. M. Pennock, D. Peters, and J. Wortman Vaughan. Truthful aggregation of budget proposals. In *Proceedings of the ACM Conference on Economics and Computation*, pages 751–752, 2019.

[14] M. Köppe, M. Koutecký, K. Sornat, and N. Talmon. Fine-grained liquid democracy for cumulative ballots. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems*, pages 1029–1037, 2024.

[15] R. Meir, A. Procaccia, J. Rosenschein, and A. Zohar. Complexity of strategic behavior in multi-winner elections. *Journal of Artificial Intelligence Research*, 33:149–178, 2008.

[16] S. Rey and J. Maly. The (computational) social choice take on indivisible participatory budgeting. *arXiv preprint arXiv:2303.00621*, 2023.

[17] P. Skowron, A. Slinko, S. Szufa, and N. Talmon. Participatory budgeting with cumulative votes. *arXiv preprint arXiv:2009.02690*, 2020.

[18] N. Talmon. Social choice around decentralized autonomous organizations: On the computational social choice of digital communities. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems–Blue Sky Ideas special track*, pages 1768–1773, 2023.

[19] S. Wang, C. Wang, T. Wang, and W. Jia. Approval-based participatory budgeting with donations. In *Proceedings of the International Computing and Combinatorics Conference*, pages 404–416, 2023.