Participatory Budgeting with Donations: The Case of Selective Voters

Philip Lazos, Evangelos Markakis, Georgios Papasotiropoulos

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 over 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[3]. 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[4]. 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 lowering their costs.

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 their preferred projects might already be funded through the public budget or others' donations. Notably, 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

[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
[3] projectcatalyst.io
[4] gitcoin.co

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,[5] 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, 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 by Chen et al. [7] and Wang et al. [19]). In Section 2.2, we discuss and formally 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 $P = \{p_1, p_2, \dots, p_m\}$, a set of n voters $V = \{v_1, v_2, \dots, v_n\}$ and a limit L 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:

[5] en.wikipedia.org/wiki/List_of_participatory_budgeting_votes

- 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 [7] and [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:[6]

- 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. While Support Monotonicity is a clearly desirable axiom, no rules in the classic PB setting are known to satisfy it; our setting, however, allows for its satisfaction. 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.

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*

[6] 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.

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 $U(F(S')) \geq U(F(S))$.*

3 Prelude to our 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 Supplementary Material.

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 *Sequential 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 Sequential and the Pareto rule fail Preference-Donation Alignment, even for PB scenarios.*

Variations of these two rules have also been considered by Chen et al. [7] and Wang et al. [19]. For all, results analogous to Observation 3 can be established, indicating that the existing rules do not immediately align with the aspirations of selective voters.

4 Our 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. We begin by presenting a detailed description of the methods we propose. Their pseudocodes are given as Algorithms 1 and 2 below.

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 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).

$B^* \leftarrow$ 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

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 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 easily 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. 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, taking into account 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 one referring to the satisfiability of a different axiom. The parts corresponding to Donation No-Harm, Support-Increase Monotonicity and Donation-Support Monotonicity, are deferred to the Supplementary Material.

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 $w_2 = 0.2 + \epsilon$, $u_2 = (0, 0.2/0.2{+}\epsilon, \epsilon/0.2{+}\epsilon)$. Moreover $s_1 = 3.1$ and $s_2 = 0.2 + 2\epsilon$.

$L = 2$		Project 1	Project 2	Project 3
	parameters	$c_1 = 2$	$c_2 = 2$	$c_3 = 2 + \varepsilon$
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, there are scenarios where DA-Pareto would need to solve an NP-hard problem in order to return a solution. Namely, the first step of the rule, i.e., the computation of B^*, essentially involves solving a knapsack instance.[7]

[7] Our reduction does not rule out pseudopolynomial time algorithms for the first step. But even if such an algorithm exists for computing B^*, it would not be immediately useful since the DA-Pareto mechanism (in its naive implementation) involves an exponential for-loop as a second step.

Theorem 5. *DA-Pareto fails to be polynomially computable, assuming P≠NP, even for PB scenarios. DA-Greedy is polynomial.*

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 one referring to the satisfiability of a different axiom, and missing parts are deferred to the Supplementary Material.

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 $w_1 = 3 - \epsilon$ and $w_2 = 1$, therefore voters' ballots could be expressed as $u_1(p_j) = \sigma_1(p_j)/3-\epsilon$ 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	$d_1 = 5$	$2 - \epsilon$	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 $2 - \epsilon$. □

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.

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.

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	✘	✘

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 as follows:

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).*[8] *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 Supplementary Material.

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 ${}^{2}/{}_{3.1}$. □

On the upside, we will 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 [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. Our negative results also extend to these cases. Searching for a mechanism satisfying all of the proposed axioms, or for a polynomial-time mechanism that provides similar guarantees to DA-Pareto, are the obvious open questions. Questions around bribery [12] also form an area for future investigation. Proportionality considerations to PB with donations are undoubtedly important. Finally, another critical direction is the experimental evaluation of our rules using data either from traditional PB settings or from the blockchain domain.

[8] Another direction could involve overstating donations. Though, this does not align well with our interpretation of d_i, which we treat as a firm upper bound on what a voter is willing to give away.

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.

Participatory Budgeting with Donations: The Case of Selective Voters

Technical Appendix

A Omitted Proofs from Section 3

Proof of Theorem 1.

Consider the instance of 2 voters and 2 projects presented in the following table, where the entry corresponding to voter v_i and project j depicts $\sigma_i(p_j)$. In this instance it holds that $u_1 = (1, 0)$ and $u_2 = (0, 1)$. We pick a value of α that is any finite number greater or equal to 1.

$L = 1$		Project 1	Project 2
	parameters	$c_1 = 1$	$c_2 = 2$
v_1	$w_1 = 1 - \varepsilon$	$1 - \varepsilon$	0
v_2	$w_2 = \alpha$	0	α

We focus first on the scenario where both voters contribute 0, i.e., $s_i = w_i$, for $i = 1, 2$. Then, the only feasible solution is $\{p_1\}$. Now, consider the scenario where $d_2 = 1$, i.e. $s_2 = \alpha + 1$; leaving all other parameters unchanged. Then $\{p_2\}$ is feasible as well, while $\{p_1, p_2\}$ is again not feasible. The support towards $\{p_1\}$ by the electorate is $1 - \varepsilon$, while $\{p_2\}$ has a total support of α. The latter solution not only maximizes total support for any chosen $\alpha \geq 1$ but is also returned by any approximation algorithm with a guarantee of α against the maximal total support. Choosing $\{p_1\}$, the only feasible option in the initial scenario, gives voter v_1 a utility of 1, whereas $\{p_2\}$, the solution in the second scenario where v_2 donates, provides a utility of 0 to v_1. This reduction in satisfaction of a voter due to donations proves that any α-approximate with respect to the total support mechanism fails to satisfy Donation No-Harm.

Similarly, if we consider the same instance as the original one but this time with $s_1 = 2 - \epsilon$ and $d_1 = 1$, once again, the optimal and any α-approximate solution would be $\{p_2\}$. With a similar reasoning as before, we can prove that v_1 essentially funds a project not supported by her. Thus, Preference-Donation Alignment is not satisfied either.

Proof of Theorem 2.

The statement holds due to a simple reduction from (the decision version of) knapsack. Consider a knapsack instance, with a given capacity, and where each item comes with a cost and a utility parameter. Let t be the target value, so that the decision question is to determine if there exists a feasible solution under a cost constraint with a total value of at least t. The construction uses a single voter, say v_1 with no donation, i.e. $d_1 = 0$. We set the public budget L to match the knapsack's capacity. We also create one project p_i for each item in the knapsack instance and set $u_1(p_i)$ and c_i to match the utility and cost respectively of the corresponding item in the knapsack instance. Say that the voter has $w_1 = 1$, therefore the support towards projects equals voter's utility. Finally, we also use the same target value t for the objective of our problem. Notice that the existence of only a single voter implies that the optimal solution of the examined problem maximizes the voter's utility, within the public budget.

To prove the forward direction of the reduction, it suffices to observe that if there is a way to select fitting items that meet the utility target t of the decision version of knapsack, choosing the corresponding projects will be feasible, given the public budget of our problem. Since the voter proposes a zero donation, she, trivially, does not contribute to projects she does not support, so Preference-Donation Alignment holds. Donation No-Harm is also always true with one voter, as increasing the donation cannot harm others, hence, the two considered axioms are satisfied. This establishes the forward direction; the reverse is identical.

Proof of Observation 3.

First, consider an instance with a single voter, namely v_1, with public budget $L = 1$, and two projects p_1 and p_2, as depicted in the following table, where the entry corresponding to v_1 and project j depicts $u_1(p_j) = \sigma_1(p_j)$.

$L = 1$		Project 1	Project 2
	parameters	$c_1 = 1$	$c_2 = 1$
v_1	$w_1 = 1, d_1 = 1$	1	0

In this instance, the Sequential method begins by selecting and funding the optimal set of projects that can be afforded within the public budget, initially choosing p_1 since its cost does not exceed L and $\sigma_1(p_1) > \sigma_1(p_2)$. Next, the rule incorporates the donation from voter v_1, since $\sigma_1(p_1) > 0$. Increasing the available budget by 1 allows p_2 to be funded within the updated budget constraint. As a result, the final selected set is $\{p_1, p_2\}$. In this scenario, p_1 is funded directly from the public budget, while p_2 got funded because of the donation by v_1, even though v_1 did not support p_2.

To examine how the Pareto rule operates in a scenario involving selective agents, let's consider an instance with two voters, namely v_1 and v_2, and the projects p_1, p_2 and p_3, as depicted in the following table, where the entry corresponding to v_i and project j depicts $u_i(p_j) = \sigma_i(p_j)$.

$L = 1$		Project 1	Project 2	Project 3
	parameters	$c_1 = 1$	$c_2 = 2$	$c_3 = 3$
v_1	$w_1 = 1, d_1 = 1$	0	0	1
v_2	$w_2 = 1, d_2 = 0$	0	1	0

Initially, the Pareto rule identifies $\{p_1\}$ as the only affordable bundle that can be funded within the public budget. It then constructs a set of feasible solutions that includes $\{p_1\}$ and all feasible bundles that dominate it. It holds that $\{p_1, p_2\}$ exceeds the budget limit when incorporating donations from voters, and the same applies to any bundle including p_3. Therefore the only feasible bundle that we could consider is $\{p_2\}$. Indeed, $\{p_2\}$ becomes feasible if v_1 contributes and, moreover, $\{p_2\}$ dominates $\{p_1\}$ because it gets more support from v_2 and no less from v_1 compared to $\{p_1\}$. Consequently, Pareto has to select among $\{p_1\}$ and $\{p_2\}$, and chooses $\{p_2\}$, since it receives the maximum support. Therefore, v_1 has contributed towards the election of a project she does not support.

B Omitted Proofs from Section 5

Remaining Part of the Proof of Theorem 4.

In the main part of this work we showed that DA-Pareto satisfies Preference-Donation Alignment but fails Support-Redistribution Monotonicity. To complete the proof of the theorem, it remains to prove that DA-Pareto satisfies additionally Donation No-Harm, Support-Increase Monotonicity and Donation-Support Monotonicity.

Donation No-Harm: Say that B^* is the optimal bundle with respect to the voters' support, when no donations are permitted. The rule outputs a bundle other than B^* only if it receives more support by some voters, under the condition that it receives no less by all the rest. So, if B is the winning bundle under DA-Pareto it holds that $\sigma_i(B) \geq \sigma_i(B^*)$, for every voter v_i or in words that B cannot be worse than B^* for any voter in terms of support. But then, $\frac{\sigma_i(B)}{w_i} \geq \frac{\sigma_i(B^*)}{w_i}$, or equivalently $u_i(B) \geq u_i(B^*)$.

Support-Increase Monotonicity: Say that the winning bundle B under DA-Pareto includes a project p and a voter increases her support towards p while leaving unaltered the support towards the rest of the projects. The reason why B won in the original instance was that it maximizes support among the affordable bundles. Since nothing changed with respect to contributions or public budget, the bundle B remains feasible and no more bundles than before are feasible. As a result, B will again be selected as the winning bundle, as $U(B)$ only increased after the change in the support of p, so it is still the best in terms of total voters' support, among the feasible bundles.

Donation-Support Monotonicity: Say that the winning bundle is B and a voter increases her donation. The reason why B won was because it maximizes support among the affordable bundles. The increase of the donation perhaps has as a result the existence of multiple more feasible solutions. In any case, B will still remain a feasible option. The winning bundle would be either B (so the axiom is being trivially satisfied), or a bundle that is superior to B in terms of voters' support since DA-Pareto selects the affordable bundle that maximizes support.

Proof of Theorem 5.

The construction used in the proof of Theorem 2 also applies here directly. The construction uses a single voter of 0 donation, implying that the second part of DA-Pareto (the for-loop) doesn't apply, and the relation to knapsack pertains to the first step of the rule: the computation of optimal in terms of total voters' support bundle which is affordable under L. The correctness of the reduction holds for the exact same reasons as the one for Theorem 2.

Remaining Part of the Proof of Theorem 6.

In the main part of this work we showed that DA-Greedy satisfies Preference-Donation Alignment but fails Donation-Support Monotonicity. To complete the proof of the theorem, it remains to prove that DA-Greedy satisfies additionally Donation No-Harm and Support-Increase Monotonicity but fails Support-Redistribution Monotonicity.

Donation No-Harm: Suppose that no donations exist in a PB scenario S. Then the outcome will coincide with the selection that is done during the first round of the method, because the second round only applies if voters donate. Now, consider a PB scenario S' that its only difference to S is that some voters are willing to donate. Then, the outcome of the first round of DA-Greedy on S' coincides with the outcome on S. The second round only adds more projects and will not delete any project selected in the first round. So, the outcome on S' is a superset of the outcome on S, which proves the satisfiability of Donation No-Harm.

Support-Increase Monotonicity: Say that a voter increases the support towards a project p that belongs to the winning bundle under DA-Greedy, and nothing else changes. Regarding the implications of this change, we observe that the project p has the same cost but receives more support from the voters. So, now, it has an increased support-to-cost ratio, and as a result it is being considered at most as late as in the execution of the algorithm in the original instance. Until considering p the run of the algorithm remains exactly the same, so, when considering p, selecting it is a feasible choice as well. Therefore, p will remain in the winning bundle after the increase of the support of a voter towards it.

Support-Redistribution Monotonicity: Consider a scenario of 3 projects, namely p_1, p_2, p_3, where $c_1 = c_2 = 1$, $c_3 = 2$ and $L = 2$. There is also a single voter v_1 of zero contribution who has $w_1 = 3$ and expresses a support of $(1, \varepsilon, 2 - \varepsilon)$, equivalently, her ballot vector is $({}^{1}\!/_{3}, {}^{\varepsilon}\!/_{3}, {}^{2-\varepsilon}\!/_{3})$. The project that maximizes the support-to-cost ratio is p_1, and after that, the only feasible option will be to also include p_2. Therefore, the winning bundle is $\{p_1, p_2\}$. Suppose that the voter decides to redistribute her ballot in a way that increases the support towards the winning project p_2, this time expressing the following support vector: $(0.5, 0.5 + \epsilon, 2 - \epsilon)$, equivalently, her ballot vector now is $({}^{0.5}\!/_{3}, {}^{0.5+\epsilon}\!/_{3}, {}^{2-\epsilon}\!/_{3})$. The project that maximizes the support-to-cost ratio is p_3 so it will be selected first. After that, there is no remaining budget to fund others and the winning bundle will be $\{p_3\}$. The increase of the support towards p_2 prohibited its election.

C Omitted Proofs from Section 6

Proof of Observation 7.

Consider the following scenario S of two projects and two voters, where an entry of the table corresponding to voter v_i and project j depicts $u_i(p_j) = b_i(p_j)$.

		Project 1	Project 2
	parameters	c_1	c_2
v_1	$s_1 = 1 - \varepsilon$	1	0
v_2	$s_2 = 1 - \varepsilon$	0	1

We fix $L = c_1 = c_2 = 1$ and $q_2 = 0$, towards, first, showing that the optimal strategy for v_1 is to keep her entire budget as voting power and donate nothing, casting $(b_1, q_1) = ((1, 0), 0)$ and having $w_1 = 1 - \epsilon$. Then, $U(p_j) = 1 - \epsilon$, for every project j. Hence, any reasonable method F that would break ties lexicographically, would only fund p_1, resulting to $u_1(F(S)) = 1$. On the other hand, if v_1 decides to contribute any strictly positive value then $F(S) = \{p_2\}$. This is because the bundle $\{p_1, p_2\}$ remains infeasible, while the support towards $\{p_2\}$ will still be $1 - \varepsilon$, in contrast to the support towards $\{p_1\}$ which will be reduced due to the fact that w_1 can now only be strictly less than $1 - \varepsilon$. Therefore, if v_1 contributes then $u_1(F(S)) = 0$.

Now consider the same instance but this time fixing $L = c_1 = c_2 = 1 - 2\varepsilon$ and $q_1 = 0$. It is now true that under any F that breaks ties lexicographically, the optimal strategy for v_2 corresponds to spending (almost) her entire budget s_2 as a donation. For that, one needs to observe that only if v_2 submits $q_2 = 1 - 2\varepsilon$ can result in $u_2(F(S)) > 0$, because F will fund p_1 first from the global budget, so p_2 can only be funded if v_2 donates c_2.

Proof of Theorem 8.

By reexamining the proof of Observation 7, specifically the first instance used there, we can easily see that the truthful scenario is not always a Nash Equilibrium as the true preferences of voter 1 could contain $d_1 > 0$ however she would prefer to declare $q_1 = 0$ in order to elect p_1.

Regarding the Price of Anarchy, consider an instance S of n voters and two projects p_1, p_2 of cost $c(p_1) = c(p_2) = n$ and say that $L = 0$. Moreover say that for each voter v_i it holds that $u_i = (1, 0)$ and $s_i = 1 + \varepsilon$. Observe that no voter can afford to buy a project alone, so $(b_i, q_i) = ((1,0),0)$, for each voter v_i, is a Nash Equilibrium. However, then, any method F would result to $B = \emptyset$, and, then $u_i(F(S)) = 0$ for each voter v_i. On the other hand. the strategy $(b_i, q_i) = ((1,0),1)$ will result to the purchase of p_1 by voters' funds under any reasonable mechanism. Then the bundle $B = \{p_1\}$ will be the winning one and $u_i(B) = 1, \forall i \in [n]$.

Remaining Part of the Proof of Theorem 9.

We begin by proving that, as DA-Pareto, DA-Greedy can also be manipulated by misreporting donations. Consider the instance that appears in the following table, where the entry corresponding to v_i and project j depicts $\sigma_i(p_j)$, when v_1 expresses her true preferences which involve $d_1 = 0.8$, $w_1 = 2$ and $u_1 = (1/2, 0, 1/2)$, and for v_2 it holds $(u_2, d_2) = (b_2, q_2)$.

$L=1$		Project 1	Project 2	Project 3
	parameters	$c_1 = 1$	$c_2 = 1$	$c_3 = 1.4$
v_1	$s_1 = 2.8$	1	0	1
v_2	$w_2 = 3.1, d_2 = 1$	0	1.1	2

We will show that by misrepresenting her donation, v_1 can cause both of her supported projects to be elected, which contrasts with the outcome under her true preferences, where DA-Greedy would elect only one. For the first round of DA-Greedy, considering only the public budget, the only affordable projects are project 1 and project 2, and among them, the method will select the second as it receives more support at the same cost as project 1. In the second round, the rule sorts the remaining projects in increasing order of support-to-cost ratio. Then, project 3 will be firstly considered as its ratio equals $3/1.4$; the ratio of project 1 equals 1. Therefore projects 2 and 3 will be bought. Since $d_1 + d_2 + L < c_1 + c_2 + c_3$, no further purchases can be made. As a result, it holds that experiences a satisfaction of $1/2$ from the outcome.

Consider now the case where v_1 acts truthfully with respect to the utility (i.e. $b_1 = u_1$) but casts $q_1 = 0.4$, i.e., declares a willingness to donate half of d_1. This results to a voting power of $w_1 = 2.4$. Hence, the support vector of v_1 now becomes $(1.2, 0, 1.2)$. Then, the first round of DA-Greedy will select project 1 since it receives more total support than project 2, and project 3 is not affordable by the public budget. Comparing the support-to-cost ratios for the remaining projects, namely p_2 and p_3, once again it holds that project 3 will be selected so the winning bundle will now be $\{p_1, p_3\}$. As a result, the change in what voter 1 decided to donate resulted in a satisfaction of 1 from the outcome.

Now, we turn our attention to proving that DA-Greedy is also manipulable by misreporting ballots. We focus once again in the instance created for proving that DA-Pareto is manipulable by misreporting ballots. We repeat the specifics of the instance below for ease of reference. Let $s_1 = 4.1, s_2 = 3.5, w_1 = w_2 = 3.1$ 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

We begin by supposing that v_1 declares her true preferences, and then, DA-Greedy selects $\{p_2\}$ to be funded by the public budget. As a result, under this method, v_1 can receive a utility of at most $1/3.1$ since either p_1 or p_3 can be then bought, but not both as their total cost exceeds the remaining budget.

Say now that v_1 changes her declared ballot to $b_1 = (2/3.1, 0, 0, 1.1/3.1)$. In this case, project 1 will be selected in the first round of the method as it has the same cost as project 2 (and these are the only affordable options) but receives more support by the electorate. Then, computing the support-to-cost ratio of project 2 we have that it is less than that of project 3, making project 3 a winning project as well. Hence, the utility of v_1 now equals $2/3.1$. Therefore, once again, we observe that the change at what v_1 declared resulted to an increase of her satisfaction.

Proof of Theorem 10.

The following problem, which we will call Π', has been proven to be NP-hard by Meir et al. [15]: We are given a set P' of candidates, a set V' of voters who have already cast cumulative votes by distributing b' points each, among candidates of P', a special voter v' (the manipulator), a specified number of winners k', a utility vector u' that indicates the true utility of v' regarding candidates of P', and a parameter t'; we are asked whether v' can cast a cumulative vote summing up to b' such that in the resulting election her utility (as specified by the vector u') obtained from the k' candidates maximizing the support of voters in $V' \cup \{v'\}$ is at least t'.

Given an instance I' of Π' we create an instance I of U-MANIP as follows: Say that $P = P'$, all candidate projects in P are of unit cost, $V = V' \cup \{v\}$, where v is the manipulator in I and will correspond to the manipulator v' of I'. For a voter $v_i \in V \setminus \{v\}$, her ballot u_i is given by $\frac{x'_i(p_j)}{b'}$, where x'_i is the cumulative vote of the corresponding voter from P'. We note that for every voter $v_i \in V \setminus \{v\}$ it holds that $\sum_{p_j \in P} u_i(p_j) = \sum_{p_j \in P'} \frac{x'_i(p_j)}{b'} = 1$. Let $s_i = b', \forall i \in V$. Moreover say that $L = k'$ and $d_i = 0$, for every voter $i \in V$. The utility of the manipulator in I towards any project $p_j \in P$ (i.e., her ballot vector) is given by $\frac{u'(p_j)}{\sum_{p_z \in P} u'(p_z)}$. Again, notice that $\sum_{p_j \in P} \frac{u'(p_j)}{\sum_{p_z \in P} u'(p_z)} = 1$. All voters in V have a voting weight of b', so the support of any voter in $V \setminus \{v\}$ is given by their ballot scaled by b'. Finally say that the optimization parameter of U-MANIP is set to $t = \frac{t'}{\sum_{p_j \in P} u'(p_j)}$.

Since all projects in I are of unit cost, and also since $L = k'$ and $d_i = 0$ for every $i \in V$, it holds that exactly the k' projects of maximum total support will be selected in the winning bundle both under DA-Pareto and under DA-Greedy. Towards proving the forward direction of the reduction, say that there is a way for v' in I' to cast a cumulative vote summing up to b' among the projects of P' in a way that the k' projects of maximum total score will give a utility of at least t' to her. Say that this results in a vector χ such that $\sum_{p_j \in P'} \chi(p_j) = b'$. Let the manipulator in I cast a ballot for each project p_j according to $\frac{\chi(p_j)}{b'}$. Taking into account that her voting power equals b', the support vector of the manipulator is identical to χ. Moreover, the support vector of each other voter v_i equals x'_i. Therefore, the outcomes of DA-Pareto and DA-Greedy in I will be exactly the set B of the k' projects that produce the maximum utility to the voters from I'. Then, the utility that the manipulator receives is equal to

$$\sum_{p_j \in B} \frac{u'(p_j)}{\sum_{p_z \in P} u'(p_z)} \geq \frac{t'}{\sum_{p_z \in P} u'(p_z)} = t.$$

Therefore, the constructed instance has also an affirmative answer to the U-MANIP problem. The reverse direction is identical.

Proof of Theorem 11.

Suppose that in a given scenario $n - 1$ voters have already submitted their cumulative ballots and contribution parameters and let v be the remaining voter. We can assume that there is a fixed ordering of the projects with respect to the preferences of v (based on her utility vector), regardless of how much the donation of v is, and that her cumulative ballot will be in line with this ordering. The cost of each project as well as the support of each voter other than v towards each project also remain fixed regardless of the donation of v. This leads to a consistent ranking of projects with respect to support-to-cost ratio, that is independent of v's donation. More precisely, if a project p_i is ranked before a project p_j when v submits a claimed contribution of q' and the support of v towards p_i is greater than the one towards p_j (the other case is identical), then p_i will still be ranked before p_j after v changes her donation to q''. This is because the support of voters other than v towards these projects remains the same, as well as their costs, and the new support of v towards p_i (after changing her donation to q'') is again greater than the support of v towards p_j. As a result, DA-Greedy examines all projects in a predetermined order that remains unaltered regardless of the donation of v, and funds some projects based on the remaining global budget first, and then, on supporters' remaining money. Suppose first that v donates $q = 0$. Then, we can run DA-Greedy and compute her utility with respect to the outcome. Let's call R_1 this execution of DA-Greedy. Note that for any possible donation of v, the outcome of the first round of DA-Greedy, which corresponds to the selection of the projects that will be funded by the global budget, remains the same. We call B^* the set of these projects. We now focus on projects not in B^* that are supported by v, and not bought by others in R_1, denoted as $\hat{B}$. Increasing the donation of v can only lead to funding projects in $\hat{B}$.

For each project in $\hat{B}$, we can compute, in polynomial time, the shortfall between its cost and intended allocation in R_1. This shortfall equals the difference between the project's cost and the amount of money that its supporters have left (possibly increased by some remaining global budget) at the time the project is being considered in R_1. Choose the project requiring the least additional donation from v in order to get funded and ask v to contribute that amount, say q'. It holds that for any contribution value between q and q' no change will happen in the utility that v will experience from the outcome of DA-Greedy. Run the procedure again under the assumption that v contributes q', call this run R_2, and compute the total utility that v gains from the outcome of DA-Greedy in R_2.

In a similar manner to before, we can compute the minimal donation q'' that v can do in order to see one more of the projects she supports being selected and repeat the procedure for q''. Actually, we can repeat the process until all projects in $\hat{B}$ are bought or v's newly computed donation exceeds her stake by calling R_i the $i-$th execution of the mechanism. Voter v should select to submit a contribution value equal to the donation $q^{(j)}$ that was used in the run of the DA-Greedy mechanism R_j, for the value of $j \leq m$ that achieves to maximize her utility among the executions of the mechanism that have been checked.