On Nash Equilibria in Participatory Budgeting with Donations and Beyond

Grzegorz Lisowski
University of Groningen
Groningen, The Netherlands
g.a.lisowski@rug.nl

Georgios Papasotiropoulos
University of Warsaw
Warsaw, Poland
gpapasotiropoulos@gmail.com

Grzegorz Pierczyński
University of Warsaw
Warsaw, Poland
g.pierczynski@gmail.com

Krzysztof Rogowski
University of Warsaw
Warsaw, Poland
kr418382@mimuw.edu.pl

ABSTRACT

Background: Real-world public decision processes increasingly combine collective budget allocation with voluntary monetary contributions. Existing models either restrict donations to projects approved by the donor or may require individuals to contribute even to projects they dislike, for the public good. The models studied in the literature account for benefits from funded projects without explicitly incorporating the personal cost of contributing.

Objectives and Research Questions: We aim to identify real-world funding scenarios across various domains where allowing participants to make private donations could have substantially improved outcomes. Motivated by these, we study strategic aspects of a model in which voters donate to increase their utility, depending on the funded projects and their contributions.

Methods: We introduce a formal framework for public funding with a global budget and individual contributions, where allocation follows a predefined funding order over projects, derived either from voters' input or externally imposed. Taking a game-theoretic approach in which participants strategically choose the magnitude of their donations, we analyze the existence, quality, structure, and further properties of Nash equilibria both theoretically and through computational experiments on participatory budgeting datasets.

Results: We provide conditions for equilibrium existence and for when contributions are necessary in equilibria under the examined model. We characterize subclasses where all equilibria can be efficiently identified and structurally described. We analyze the size of the equilibrium set, propose a heuristic to detect all donation strategies that satisfy minimal strategic incentive requirements, evaluate their quality across multiple metrics, and study how personal contribution bounds affect the outcomes.

Conclusions: Our framework offers a systematic understanding of strategic donations in public funding mechanisms.

ACM Reference Format:
Grzegorz Lisowski, Georgios Papasotiropoulos, Grzegorz Pierczyński, and Krzysztof Rogowski. 2026. On Nash Equilibria in Participatory Budgeting with Donations and Beyond. In *Appears at the 8th Games, Agents, and Incentives Workshop (GAIW-26). Held as part of the Workshops at the 25th International Conference on Autonomous Agents and Multiagent Systems., Paphos, Cyprus, May 2026*, IFAAMAS, 9 pages.

Appears at the 8th Games, Agents, and Incentives Workshop (GAIW-26). Held as part of the Workshops at the 25th International Conference on Autonomous Agents and Multiagent Systems., Armstrong, Curry, Hosseini, Mattei, Tsang, Wąs (Chairs), May 2026, Paphos, Cyprus.

1 INTRODUCTION

Participatory budgeting (PB) enables communities to directly influence the allocation of public funds [3, 18, 19]. Residents submit project proposals and express preferences over all proposals, e.g., by approving those they support. In the vast majority of applications, a simple greedy method is used for selecting projects: they are considered in order of approval count and funded as long as their cost is affordable. Otherwise—no matter how small the shortfall or how strong the support is—a project whose cost exceeds the remaining budget is discarded. But what if that gap could be covered by a (small) contribution from (some of) its supporters? Perhaps it would be worth giving them the opportunity to step in with a contribution, so that the project can still be realized, rather than redirecting funds to less popular projects or leaving funds unused.

As a particularly illustrative example, consider the 2022 local PB election in Łódź, Teofilów-Wielkopolska (available within the PABULIB library [12]). The total budget was 826,000 Polish złoty (PLN), and the four most supported projects were as follows:

Project	Cost (PLN)	Approvals	Budget Left if Funded
A	150,000	13,390	676,000
B	550,000	13,314	126,000
C	140,000	13,166	-14,000
D	15,000	3,695	111,000

When Project C was considered, the remaining budget was not enough to fund it. Project D, which had more than three times fewer supporters than Project C, was funded instead. The supporters of Project C, if given the option, would only need to contribute about 1 PLN (≈ 0.25€) each to secure its selection, directing a substantial portion of the budget toward a (highly approved) project they favor.

Similar situations can also be identified in a wide range of domains of public funding beyond PB. Several *blockchain-based funding platforms* have recently gained traction. There, a fixed budget is used to fund community-proposed projects based on the preferences of token holders [10, 13]. A commonly used method follows the greedy approach described earlier for PB. In Project Catalyst's 13th funding round [17], a €[1]48,000 project was not funded because, after projects with higher support were selected, the remaining budget fell 0.6% short of covering its cost. With over 400

[1]The (crypto)currency used differs across the discussed applications. For clarity, amounts are converted and roughly rounded.

backers, its supporters might have preferred to contribute less than €1 each to make up the shortfall—if given the option—rather than see the funds redirected elsewhere. Similarly, in DeepFunding's 3rd round, a project for an AI app on kidney disease, supported by nearly a quarter of the electorate, required €5,000 but only €4,700 was available after more popular projects took their share of the budget [11]. Participants did not have the option to contribute directly, and the available funds went unused.

A further motivation of our study comes from *charitable funding programs*. Several countries (e.g., Italy, Poland, Slovakia) allow taxpayers to redirect a portion of their personal income tax (up to 0.5%, 1.5%, and 3%, respectively) to a charitable organization of their choice through their annual tax declaration form, instead of directing this amount to the state. Citizens select an organization from a list of eligible entities. Imagine one of these organizations publishing a prioritized list of actions it will carry out depending on how much funding it will ultimately receive. This order of priorities combined with the expectations of the citizens about the total funds the organization could collect could significantly influence whether and how much people decide to allocate to it.

Our work also applies to *crowdfunding platforms*, which have become popular ways for creators to pitch ideas and secure funding by setting financial goals and relying on supporter contributions. In these platforms, the campaign creator sets a funding order: reaching an initial goal guarantees funding of a base project, while additional funds unlock subsequent projects (e.g., "stretch goals" in Kickstarter) funded as higher thresholds are met in a predefined sequence. The commonly used all-or-nothing model means that if a target is not met by the deadline, no matter how small the shortfall, the relevant funds are not released. This can lead to missed opportunities, as funds remain untapped for projects that fall just short of their goals. We refer to the full version of our work for an outline of striking real-life cases where extra contributions—as little as a handful of cents—could have bridged the shortfall.

To our knowledge, the diverse scenarios above have not previously been systematically identified; bringing them together constitutes our first, modest but concrete contribution. These, illustrate that in numerous settings, outcomes could have been substantially improved if participants were allowed to make donations. We assume that donations are permitted and explore aspects of strategically donating.

Our Contribution. We propose a formal framework for analyzing settings where individuals not only influence the allocation of public funds but may also voluntarily contribute additional money to support the implementation of their preferred options. It captures a range of real-world scenarios, from public budgeting to charitable giving and crowdfunding. For this purpose, we assume the existence of a *funding order*. It may either be derived from voters' input (such as approval counts in the greedy method commonly used in city-level community funding or blockchain-based applications), or externally defined by an authority (such as a charitable organization or a crowdfunding campaign creator). The election organizer allocates funds to projects based on this order, considering both the global public budget and individuals' donations. We adopt the game-theoretic perspective: assuming that voters are aware of the funding order, as well as others' preferences and donation tendencies, each individual selects a donation strategy aimed at serving their own interests. Within the proposed framework:

(1) We provide theoretical and empirical insights into questions like when non-zero contributions are necessary to reach an Nash equilibrium (NE), when non-zero NE exist, and when NE may fail to exist altogether.
(2) We identify conditions for NE determination. We characterize subclasses of the framework in which all equilibria can be efficiently found, and their structure can be well understood.
(3) We analyze the order of magnitude of the number of equilibria in theory and practice. Then, we propose a heuristic procedure for identifying all equilibria, and we evaluate their quality empirically across distinct metrics.
(4) We also study the impact of bounds on voters' contributions from the theoretical and empirical perspective. Such a restriction is motivated by practical considerations, as voters are often subject to intrinsic limits on how much they are willing or able to donate.

Related Literature. Frameworks for public decisions that allow voters to contribute monetarily have already been explored by Chen et al. [9], Wang et al. [21] and Lazos et al. [14]. In the former two, a voter's contribution may end up being used for enabling funding for a project they disapprove of. At the other extreme, in the latter, voters contribute only if this results exclusively in the election of projects they like. The model we examine lies in between: Voters aim to contribute monetarily to increase their utility, even if this partly funds projects they disapprove of. Unlike the frameworks studied by Chen et al. [9] and Wang et al. [21], which suggest per-project donation bounds, our model—similarly to the one examined by Lazos et al. [14]—assumes a single global bound per voter. Also, our analyses apply to any voting rule inducing a funding order, e.g., the greedy method does, rather than designing new ones. Importantly, unlike all the mentioned works where participants' utility depends solely on the funded projects, our utility model incorporates the cost of donating, capturing the personal burden of contributing. Specifically, we adopt the natural notion of quasi-linear utility functions, as used in related contexts by Aziz et al. [2], Buterin et al. [8] and Birmpas et al. [5], which increase with the number of approved and funded projects and decrease with the amount of money a voter donates. The works of Brandl et al. [7], Aziz et al. [2] and Aziz and Ganguly [1] are also conceptually close to ours, however, there is no exogenous budget in their setting. Finally, Boehmer et al. [6] and Baumeister et al. [4] study projects' evaluation and interventions in PB that quantify, among others, how much a project's cost must be reduced to make it electable under various rules. While their models abstract away from strategic contributions, such cost reductions could also, in practice, be implemented through donations.

2 FRAMEWORK AND PRELIMINARIES

We refer to an instance of the framework described as a *pledge game* $\mathcal{G}$. In such, there is a set of *voters* $V = \{v_1, \ldots, v_n\}$ and a set of *projects* $\mathcal{P} = \{p_1, \ldots, p_m\}$. Each project $p \in \mathcal{P}$ has a *cost of implementation*, $\mathrm{cost}(p) \in \mathbb{R}_+$. Each voter $v_i \in V$ *approves* of the set $A_i \subseteq \mathcal{P}$. For each project $p \in \mathcal{P}$, we define its set of *supporters* as $\mathrm{supp}(p) := \{v_i \in V : p \in A_i\}$. We assume that $|\mathrm{supp}(p)| \geqslant 1, \forall p \in \mathcal{P}$.We are also given a publicly available *global budget* $\beta \in \mathbb{R}_{\geqslant 0}$, upper bounded by $\sum_{p \in \mathcal{P}} cost(p)$ and will be used for funding projects. In addition to the global budget, voters may choose to contribute *donations* in an effort to increase the chances of funding projects they support. Each voter $v_i \in V$ is associated with a monetary value, i.e., a donation $d_i \in \mathbb{R}_{\geqslant 0}$. We refer to the tuple $\mathbf{d} = (d_1, \ldots, d_n)$ as a *donation profile*. We write $(\mathbf{d}_{-i}, d'_i)$ to denote the profile in which voter v_i donates $d'_i \neq d_i$ while all other voters donate according to $\mathbf{d}$. To compare with the classical setting in the literature of budgeting processes where donations are absent, we often refer to the *all-zero profile* $\mathbf{d_0} := (0)^n$, in which no voter contributes any funds.

The projects in $\mathcal{P}$ are considered for funding according to a strict total *funding order* and we assume that the set $\mathcal{P}$ is ordered according to it (i.e., p_i is higher in the order than p_j if $i < j$). Given the funding order, the projects are processed sequentially and, when a project p is considered, it is selected for implementation if its cost does not exceed the remaining funds. Otherwise, it is skipped. This process continues until all projects have been considered. Note that such a funding order is not necessarily exogenous as it can also be directly derived from approval counts or induced by other voting rules, making its presence an advantage rather than a limitation.

Given a pledge game $\mathcal{G}$, and a donation profile $\mathbf{d}$, the *outcome* of $\mathcal{G}$ can be inferred. Specifically, using B to denote the *total available funds*, i.e., $B := \beta + \sum_{v_i \in V} d_i$, the outcome depends on the candidates, their costs and funding order as well as on B, and we use $W_{\mathcal{P}}(B)$ to denote it, or simply $W(B)$ when $\mathcal{P}$ is clear from the context. We note that, following conventions from the Computational Social Choice literature, the description of $\mathcal{P}$ specifies not only the identity of the projects and their position in the funding order, but also a mapping from each project to its cost. For an outcome $W(B)$, we let $\ell(W(B))$ be the project $p_i \in W(B)$ that appears last in the funding order inferred by $\mathcal{P}$, i.e., the one in $W(B)$ having the maximum index. We say that $W(B)$ is a *prefix* of $\mathcal{P}$ if there exists an index $j \in 1, \ldots, m$ such that $W(B) = \{c_1, c_2, \ldots, c_j\}$, that is, the outcome consists of the first j projects in the order defined by $\mathcal{P}$, without skipping any.

Given an outcome $W(B)$ we evaluate voters' satisfaction using a *quasi-linear utility function*. We first assume that a voter's utility increases with the total cost of approved projects included in the outcome—an approach that aligns with the cost utility models in PB and beyond [15, 16, 20]. Additionally, we account for each voter's personal contribution, which reduces their utility according to the amount donated. Given a pledge game $\mathcal{G}$ and a profile $\mathbf{d}$, we define the *cost utility* of v_i under $W(B)$ induced by $\mathbf{d}$ and we refer to Section 2.1 for a generalization:

$$u_i(\mathbf{d}) := \sum_{p \in W(B) \cap A_i} cost(p) - d_i. \tag{1}$$

A voter may donate even to projects they do not approve of, if doing so benefits projects they approve of. A strategy profile $\mathbf{d}$ is a (pure) *Nash equilibrium* (NE) if for every voter v_i and every strategy $d'_i \in \mathbb{R}_{\geqslant 0}$ we have that $u_i(\mathbf{d}) \geqslant u_i((\mathbf{d}_{-i}, d'_i))$, i.e., when no voter can increase their utility by unilaterally changing their donation.

Example 1. Consider a pledge game where $\beta = 100$ and $\mathcal{P} = \{p_1, p_2\}$, with $\text{cost}(p_1) = 200$ and $\text{cost}(p_2) = 150$, where $\text{supp}(p_1) \cap \text{supp}(p_2) = \emptyset$. Under $\mathbf{d_0}$, i.e., with no donations, no project can be funded. Suppose that the supporters of p_2 collectively donate 50. Then, $B = 150$ and p_1 is considered first in the funding order (because, e.g., $|\text{supp}(p_1)| > |\text{supp}(p_2)|$). Since the available amount is insufficient to fund p_1, the process moves to p_2, which is affordable. Hence, by contributing a total of 50, the supporters of p_2 secure funding for the project they approve (regardless of how the contributions are distributed). In particular, the larger the set $\text{supp}(p_2)$ is, the lower the donation from each of its members can be. If a single voter from $\text{supp}(p_2)$ contributes the entire amount of 50, then their utility is 100, while the utility of the others in $\text{supp}(p_2)$ is 150, and the utility of voters in $\text{supp}(p_1)$ is zero. Suppose instead that $\text{cost}(p_1) < 150$. The donation of 50 by the supporters of p_2 leads to a different outcome. Specifically, $W(B) = \{p_1\}$. Then, at least one voter in $\text{supp}(p_2)$ has a negative utility and prefers to deviate to a zero donation, hence the considered profile is not an NE.

Limitations of the Work. (i) Our study is centered around (pure) NE as it is the standard starting point for strategic reasoning, providing a robust and interpretable stability baseline; any strategically acceptable outcome should at least satisfy individual rationality. Our main goal is to identify the full spectrum of outcomes satisfying this criterion. That said, individual incentives are also behaviorally realistic in settings where communication or cooperation is cumbersome or unreliable. Indicatively, first, in blockchain governance, participants act through pseudonymous wallets with little possibility of verified coordination. Second, in PB, voting is anonymous, making pre-election agreements inherently fragile, as no binding commitments are possible. In line with standard NE analysis, we assume complete-information. (ii) Additionally, we view contributions as going to the public budget rather than to individual projects—an alternative that would also be reasonable to study (particularly for PB, though less so for the other settings captured by our model). Our choice is supported by real-world cases where the overall total budget has been increased after the vote, either by individuals or by the organizing authority [22]. Nevertheless, as several of our proofs show, due to the funding order, donors effectively know where their funds will be allocated, even if only implicitly and indirectly. (iii) Moreover, allowing donations may introduce some imbalance and wealth-based inequality among voters, which we acknowledge as a limitation. However, the funding order mitigates this effect, at least partly, since, for example, when determined by approval counts, additional influence tends to favor more widely supported projects, while such concerns are even less pronounced in crowdsourcing or charitable funding contexts. (iv) Finally, while misreporting project preferences is also a natural behavior to consider, we abstract from it mainly to obtain a cleaner and more focused understanding of the strategic role of contributions. Furthermore, because the funding order may be independent of voters' ballots—and ballots may even be absent in settings such as crowdsourcing or charitable funding—misreporting could in some cases not affect the outcome.

2.1 Special Cases, Generalizations and Variants

Our model independently generalizes two natural and important application domains: In *participatory budgeting*, the funding order is obtained from voters' preferences and some fixed voting rule. Most often, the order reflects project popularity—e.g., by ranking projects according to the number of supporters, yet the use of more complex voting rules could be covered by our model as well. In the context of *charitable donations*, the order may be exogenously specified by the organization, reflecting its internal priorities. Having introduced our base model, we now present several natural special cases, generalizations, and variants. These will be analyzed in the following sections to provide a comprehensive understanding.

Bounded-Contributions Setting. Consider the case where each voter v_i has a *donation bound* b_i, representing the maximum amount they are willing to contribute. We denote the vector of these bounds as $\mathbf{b} = (b_1, \ldots, b_n)$. Then, only strategy profiles $(d_i)_{i\in[n]}$ satisfying $d_i \leqslant b_i$, for each voter v_i are feasible. We refer to the base model that does not involve donation bounds as *boundless*.

Single-Approval Case. The setting is *single-approval* if each voter approves exactly one project, that is, $|A_i| = 1, \forall v_i \in [n]$.

General- and Cost Utilities. In the base model each voter v_i experiences a utility that directly relates approved and selected projects' cost to the voter's contribution. In reality, the value a voter gains from a project does not always correspond exactly to its cost. To account for such discrepancies, we let v_i be associated with a function $\gamma_i : A_i \mapsto (0, 1]$. The weight inferred by the function reflects how much of the project's cost translates into actual value for the voter and this flexibility is useful since the benefit a voter derives from implemented projects and the cost of their contributions are not always comparable in the same scale. The utility then is $u_i(\mathbf{d}) = \sum_{p\in W\cap A_i} \gamma_i(p)cost(p) - d_i$. We refer to this setting as *general utilities* and, in contrast, we call the case where $\gamma_i(p) = 1, \forall v_i \in V$ and $p \in A_i$ (i.e., Expression (1)) the *cost utilities* setting.

Crowdfunding Setting. Consider a single fundraising campaign which consists of several individual projects. Those projects are typically funded in a sequential and dependent manner, where funding a project depends on earlier milestones being met. This structure creates a dependency such that no project can be funded unless all preceding projects in the order have already been funded. This naturally induces a specific funding relation that applies to our framework. To capture this, we define the *crowdfunding setting* in which the following additional feasibility constraint is being enforced: for any two projects $p_i, p_j \in \mathcal{P}$ with $i < j$, having p_j in the outcome requires that p_i is also included in the outcome.

A pledge game $\mathcal{G}$ can be written as a tuple $(V, \mathcal{P}, \beta, (A_i)_{i\in[n]})$; $(\mathbf{b}, (\gamma_i)_{i\in[n]})$ can be included in the tuple, when they apply.

2.2 Experimental Setup

Throughout our work, we complement the theoretical findings with empirical results. For this, we analyze data from PB elections which are available in PABULIB. Specifically, we evaluate 676 instances (for which our heuristic method for identifying all NE completed within a reasonable time frame—refer to the full version for details. Moreover, to allow meaningful comparisons across instances involving monetary contributions, we focused on elections conducted in the same currency—we restricted attention to elections using PLN, which account for around 80% of the approval-based instances in PABULIB. While our theoretical results apply to arbitrary funding orders, in the experiments we adopt a specific order which is in line with the greedy method commonly used in PB processes globally: projects are sorted by the number of approvals they receive. Finally, for simplicity and clarity, we assume $\gamma_i(p) = \gamma_j(p')$ for every pair of voters (v_i, v_j) and any pair of projects (p, p'). In such cases, we refer to the weighting function as γ. We selected such values as inverse powers of 2. We note that this assumption applies only to our experiments and we use it as PB datasets lack valuation information, and it helps avoid confounding effects ensuring interpretability.

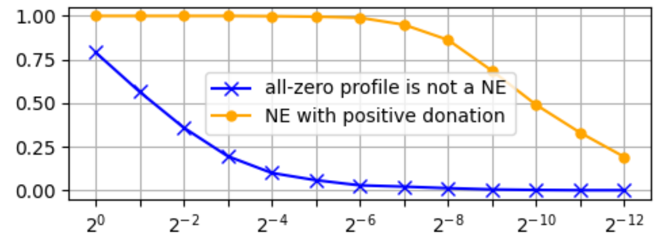

Figure 1: Fraction of instances not admitting the all-zero profile as an NE (blue crosses) and admitting at least one profile other than the all-zero as an NE (yellow dots), for various values of γ.

3 ANALYSIS OF PURE NASH EQUILIBRIA

We begin by identifying when the all-zero profile is an equilibrium in the proposed framework, that is, when no voter is incentivized to contribute. We identify a simple condition under which the all-zero profile is guaranteed to be an NE and we also show that this is tight, in the sense that if it is not satisfied, the profile $\mathbf{d_0}$ is not necessarily an NE. This result implies that when a (positive) global budget is available, donations might be necessary in order to reach an equilibrium.

Observation 1. *If $\beta = 0$, the all-zero profile is always a Nash equilibrium; for $\beta > 0$, it might not be.*

PROOF. First, take an instance where $\beta = 0$ and its corresponding all-zero profile $\mathbf{d_0}$. Observe that no projects are funded under $\mathbf{d_0}$ and that none of the voters made any contributions, and hence for each i we have $u_i(\mathbf{d_0}) = 0$. Suppose now that $\mathbf{d_0}$ is not an NE. Then, there is a voter v_i that prefers to donate d_i' instead of 0. Say that such a deviation results in a donation profile $\mathbf{d}$ where all voters but v_i donate according to $\mathbf{d_0}$ and v_i donates d_i'. Under $\mathbf{d}$, $B = d_i'$ and hence $\text{cost}(W_{\mathcal{P}}(B)) \leqslant d_i'$. Then, $u_i(\mathbf{d}) = \sum_{p\in W_{\mathcal{P}}(B)} \gamma_i(p)\text{cost}(p) - d_i' \leqslant \sum_{p\in W_{\mathcal{P}}(B)} \text{cost}(p) - d_i' \leqslant 0 = u_i(\mathbf{d_0})$. It follows that $\mathbf{d_0}$ is an NE.

Towards proving that there are instances of $\beta > 0$ in which the all-zero profile is not an equilibrium, consider the simple scenario of a single project p_1, a budget $\beta = \text{cost}(p_1)/2$, and a voter v_1 who approves p_1, under cost utilities. Then, $\mathbf{d_0}$ is not an NE, as v_1 would prefer to donate $cost(p_1)/2$ instead of 0 and thus allowing for the purchase of p_1. In this case, v_1 would receive the utility of $cost(p_1)/2 > 0 = u_1(\mathbf{d_0})$. □

Therefore, when $\beta = 0$, as for instance in charitable settings where no exogenous budget is available, no voter has an incentive to donate on their own. In contrast, this is not always the case in public budgeting. We now present empirical observations on the existence of incentives for unilateral deviation from the all-zero profile. In more than 3/4 of the examined real-world instances, there exists at least one voter willing to donate unilaterally to change the outcome in their favor. This holds under cost utility preferences and, as shown in Figure 1, the percentage gradually decreases as the value of γ is reduced, since unilateral deviations become then less beneficial. For moderate and lower values of γ, the all-zero profile is already an equilibrium in many cases.

We turn our attention to the, most interesting, case of $\beta > 0$. We note that for certain families of pledge games, determining whether the all-zero profile is a Nash equilibrium can be done efficiently. This holds in particular for single-approval instances and under cost utilities, as well as in the crowdfunding setting, even for general utilities. In the latter case, it suffices to check whether any supporter of the first project that remains unfunded under $\mathbf{d_0}$ would benefit by donating the shortfall between the project's cost and the available budget at the time the project is considered—that is, whether the utility gain outweighs the required donation. In the former, we check the same condition not only for the first unfunded project but also for any further project of no higher cost. These observations follow directly from the characterization results in the corresponding settings, which we present later as Theorems 3 and 4.

We next examine the existence of equilibria. We first identify a class of pledge games in which an equilibrium other than the all-zero is guaranteed to exist. Additionally, we showcase tightness of this class by proving that if either of the conditions characterizing the class is not met, there can be scenarios where the all-zero equilibrium is the only NE, as well as pledge games where no NE exists at all.

Theorem 2. *In the boundless setting and with cost utilities there is always a non-zero NE. In the setting of bounded contributions or under general utilities an NE may not exist or the zero-profile might be the only one.*

Proof Sketch. Let $\hat{c}$ be the first project in the funding order of $\mathcal{P}$ that does not belong to $W(\beta)$ and r be the available budget at the point of arrival of $\hat{c}$ and $r' := \text{cost}(\hat{c}) - r$. We show that the donation profile $\mathbf{d}$ where d_i equals $r'/|\text{supp}(\hat{c})|$, for each $v_i \in \text{supp}(\hat{c})$ or 0 otherwise, is an NE. If a voter v_i reduces their donation to $d'_i < d_i$ we show that the change in their utility will be non-positive. Also, if a voter increases their payment beyond the amount specified by $\mathbf{d}$, they would need to fund a project on their own to further increase their satisfaction, which would not increase their utility. Regarding the second part of the statement, for the bounded setting, it suffices to consider a case of $\beta = 0.5$, projects p_1 and p_2 such that $\text{cost}(p_1) = 3, \text{cost}(p_2) = 2$ and voters v_1 (supporting only p_1) and v_2 (supporting only p_2) with $b_1 = 1$ and $b_2 = 1.5$. □

We saw that under cost utilities there always exists an NE where voters donate a positive amount, which is not necessarily the case for general utilities. Empirically, as also illustrated in Figure 1, we observe that such equilibria continue to exist across all the examined instances, even as we move away from the cost utility setting, for large and up to moderate values of γ. The further we move away, beyond a certain point, the lower the percentage becomes.

It is important to note that pledge games that have at least one equilibrium other than $\mathbf{d}_0$ may have an infinitely large number of equilibria that differ on how the overall donation is being split among contributors: take, e.g., a boundless instance of cost utilities with 2 voters and a project p_1 approved by both, and with $\beta = 0$. There are infinitely many ways to split $cost(p_1)$ among the voters and it is easy to see that each of them is an equilibrium. In this simple example, although there are infinitely many distinct equilibria, they all lead to the same outcome. Since a potentially large number of equilibrium profiles can be indistinguishable from the perspective of their outcomes, we will hereinafter mostly focus on detecting classes of profiles such that all members of which result in the same outcome. Profiles belonging in such a class are equivalent up to a redistribution of donations among voters which lead to the same aggregate donation amount, and, consequently, the same set of funded projects.

Definition 1. For an instance $\mathcal{G}$, two strategy profiles $\mathbf{d}, \mathbf{d'}$ are *budget-equivalent* if $\sum_{v_i \in V} d_i = \sum_{v_i \in V} d'_i$. Else, they are budget non-equivalent. A family of budget non-equivalent NE such that all other NE of $\mathcal{G}$ are budget-equivalent with one from the family is called a *characterizing family of NE*.

A characterizing family of NE of a pledge game captures all NE inducing different outcomes. We turn our attention to the crowdfunding setting. For this, a particularly strong result can be established both for the unbounded and bounded contributions setting. We show that we can compute and concisely depict the simplex of all possible NE of such an instance. Moreover, NE in this setting share a certain characteristic structure. A profile $\mathbf{d}$ is an NE only if supporters of the last project of $W_{\mathcal{P}}(B)$ are contributing under $\mathbf{d}$.

Theorem 3. *In an instance of the crowdfunding setting $\mathcal{G}$, a strategy $\boldsymbol{d}$ is an NE if and only if there exists a value $B \geqslant 0$ such that $\beta + \sum_{v_i \in \text{supp}(p^*)} d_i = B$ and $\sum_{v_i \notin \text{supp}(p^*)} d_i = 0$ where $p^* = \ell(W_{\mathcal{P}}(B))$. Moreover, a characterizing family of NE can be computed in polynomial time.*

Proof Sketch. A first observation that plays a central role in the proof is that in any NE $\mathbf{d}$ for the crowdfunding setting, the outcome is a prefix; this follows directly from the design of the setting. Since there are $O(m)$ possible prefixes of $\mathcal{P}$ that can serve as outcomes $W_{\mathcal{P}}(B)$, we can define a characterizing family of NE with at most $O(m)$ distinct strategy profiles. Moreover, we show that every voter with a positive contribution under $\mathbf{d}$ must support p^*, and the total contribution must equal $\sum_{p_i \in W_{\mathcal{P}}(B)} \text{cost}(p_i) - \beta$. It then suffices to formulate a linear program capturing the simplex of all possible NE. To do so, we specify a linear incentive compatibility constraint, which—together with constraints inferred directly from the observations on the structure and total cost of each NE—provides necessary and sufficient conditions for $\mathbf{d}$ to constitute an NE with outcome equal to the prefix up to p^*. Solving the program for a fixed project p^* yields either: the simplex of all budget-equivalent NE that fund the corresponding prefix, or infeasibility, indicating that no NE funding this specific prefix exists. The union of feasible simplices across all such candidate projects forms the set of all NE of $\mathcal{G}$. □

The main technical tool leading to the proof of Theorem 3 was the observation that the setting of crowdfunding itself forces each outcome to be a prefix. Although not as obvious at first sight, the same holds for the single-approval setting under cost utilities, in the boundless setting. This leads to an analogous positive result.

Theorem 4. *In an instance of the single-approval setting $\mathcal{G}$ and under cost utilities in the boundless setting, a strategy profile $\boldsymbol{d}$ is an NE in $\mathcal{G}$ if and only if there exists a value $B \geqslant 0$ such that $\beta + \sum_{v_i \in \text{supp}(p^*)} d_i = B$ and $\sum_{v_i \notin \text{supp}(p^*)} d_i = 0$ where $p^* = \ell(W_{\mathcal{P}}(B))$. Moreover, a characterizing family of NE can be computed in polynomial time.*

Proof. We begin by proving an important structural property of any outcome under the single-approval setting and cost utilities. Specifically, in such an instance, only a prefix of $\mathcal{P}$ can be funded in a Nash equilibrium. To establish this, suppose for contradiction that there exists a Nash equilibrium $\mathbf{d}$ where a project p_i is not funded but a subsequent project p_k, for $k > i$, is funded. It is easy to see that under the single-approval setting, any voter $v_j \in \text{supp}(p_i)$

satisfies $d_j = 0$. Given that p_i is not elected, it also holds that $u_j(\mathbf{d}) = 0$. Though, in this case, v_j would prefer to contribute any amount necessary to fund p_i. This amount would be strictly less than $\text{cost}(p_i)$ since funds previously directed toward p_k or any project in between the two will now be used for p_i. Under cost utilities, this is profitable for the v_j as it will result in a strictly positive utility, which contradicts the fact that $\mathbf{d}$ is an equilibrium. Note that in any Nash equilibrium that realizes an outcome that is the prefix $P(p_t)$ for some $t \in [m]$, voters that do not belong in $\text{supp}(p_t)$ will contribute zero. This follows from the same arguments as in the proof of Theorem 3.

Consider a profile $\mathbf{d}$ where only supporters of the last funded project, say p_t are contributing positive amounts. Moreover say that they contribute exactly $\sum_{p_i \in P(p_t)} \text{cost}(p_i) - \beta$. We will show that $\mathbf{d}$ is an NE. We examine the two potential deviations, i.e., decrease of their donation or increase of the donation of any voter. First, if any contributing voter reduces their donation, this results in p_t no longer being funded. Under $\mathbf{d}$ this voter, say v_i experiences a utility that is zero if they are the only contributing voter or strictly positive otherwise. We now consider what happens when v_i reduces their donation. Since we are in the single-approval setting and p_t is the only one providing utility to v_i, their utility does not increase after the deviation as they lose the benefit of their approved project and hence the utility will be upper bounded by zero then. If any voter increases their donation (either a non-contributor becoming a contributor or an existing contributor increasing their payment), they would need to fund additional projects beyond the current prefix entirely on their own. Since these additional projects provide zero additional utility to the deviating voter in the single-approval setting, such deviations are not profitable either.

Consequently, since no profitable deviations exist, the equilibrium conditions reduce to simple budget balance that adds up to the cost of the prefix $P(p_t)$ among the supporters of the last funded project p_t. Hence, for each possible prefix $P(p_t)$, the simplex of Nash equilibria can be found by specifying the participation and budget constraints from the proof of Theorem 3. □

The positive result of Theorem 4 is relying on a set of assumptions. The result that follows shows that all these are in fact necessary to ensure that a polynomially large characterizing family of NE exists.

Theorem 5. *The number of budget-nonequivalent NE might be exponential in m for (1) single-approval ballots under general utilities, (2) general ballots under cost utilities. Both results hold even for the boundless setting.*

Proof Sketch. For (1) we use $\gamma = 1/2$ and provide an instance on m projects of cost $3^m, 3^{m-1}, \ldots, 3$. Take $3^{m+1} - 3$ voters who support (only) the last in order project, namely c_m. All projects in $\{c_1, \ldots, c_{m-1}\}$ have a single supporter, so the total number of voters is $3^{m+1} + m - 4$. Say also that $\beta = 0$. For this instance it holds that asking all voters that support the last project to buy any subset of projects by splitting the cost equally among them is an NE. This is because, first, whichever set of projects they buy, they get a utility that is at least $3/2 - \frac{3^m + 3^{m-1} + \cdots + 3}{3^{m+1} - 3} = 1$. Second, if a voter v_i that is paying 0 increases the donation to d_i in order to buy their approved project, say the $(m-j+1)$-th project in $\mathcal{P}$, then they will get a utility of $3^j/2 - d_i \leqslant 3^j/2 - 3^j + (3^{j-1} + 3^{j-2} + \cdots + 3) = 3/2 - 3^j < 0$. □

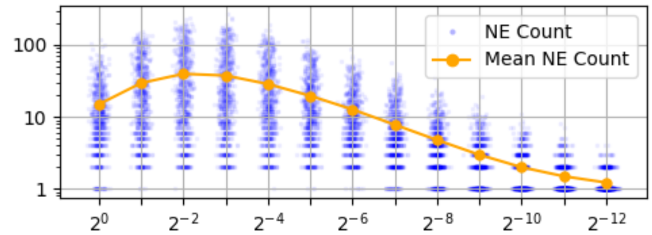

Figure 2: Number of budget-nonequivalent NE per instance for different values of γ. These numbers per instance are shown as (small) blue dots, while the (large) yellow dots indicate averages.

Complementing Theorem 5, we experimentally examine the number of equilibria in the considered instances. In Figure 2 we observe an interesting trend in how the value of γ affects the average and maximum number of equilibria: as γ deviates from cost utility, the number of equilibria initially grows, but then gradually declines toward 1. For moderate values of γ, the number of equilibria is larger as, when γ is high, voters are more inclined to contribute individually to secure high utility, while when γ is low, voters tend to donate nothing, as the project costs become harder to split among supporters in a way that gives everyone a significant level of satisfaction.

Beyond the settings of crowdfunding (Theorem 3) and single-approval under cost utilities (Theorem 4), we provide a condition that characterizes NE in arbitrary pledge games. This relies on the notion introduced below.

We say that a strategy profile, denoted by $\mathbf{d}_i^1$, is a *single-contributor profile* if all contributions come from a single voter v_i, In $\mathbf{d}_i^1$, if $\sum_{v_j \in V} d_j = \delta > 0$, the contribution vector is defined as: $d_j = \delta$, if $j = i$ and 0 otherwise. Consider a profile $\mathbf{d}_i^1$ with $\beta = 0$. Let $W := W_{\mathcal{P}}(\delta)$ be the outcome under $\mathbf{d}_i^1$, and let $u_i^1(\delta)$ denote the utility of v_i, i.e., $u_i^1(\delta) := u_i(\mathbf{d}_i^1) = (\sum_{p \in W \cap A_i} \gamma_i(p) \cdot cost(p)) - \delta$.

Proposition 6. *Given an instance $\mathcal{G}$, a donation profile $\mathbf{d}$ is an NE if and only if for each voter $v_i \in V$ it holds that:*

$$u_i^1(\delta) = \sup_{x \in [0, b_i]} u_i^1(\delta - d_i + x),$$

where $b_i \in \mathbb{R}_{\geqslant 0} \cup \{+\infty\}$ and $\delta = \sum_{v_i \in V} d_i$.

One immediate strength of Proposition 6 is that the condition ensuring a voter v_i has no incentive to deviate from d_i is independent of how the remaining donations (i.e., $\delta - d_i$) are distributed among other voters. Unlike the characterizations in Theorems 3 and 4, the condition provided in Proposition 6 does not yield a polynomial-time method for recognizing equilibria. Still, it is not only of theoretical interest; it also serves as the building block of the heuristic approach we develop and use in the experimental component of our work.

In particular, we can exploit the above property to check the existence of equilibria with a fixed total budget and find a representative among them. Suppose we are given an oracle ζ that, for a value δ and a voter v_i, returns the maximal interval $R \subseteq \mathbb{R}$ such that $\delta \in R$ and $u_i^1(\delta) = \sup_{x \in R} u_i^1(x)$. Combining this with Proposition 6, we can determine for which values of d_i the interval $[\delta - d_i, \delta - d_i + b_i]$ is contained in R. This yields a pair of inequalities for d_i characterizing when v_i has no incentive to deviate. Applying this reasoning to all voters yields a full characterization of the set of equilibria with total donation δ, using a linear number of calls to ζ. The oracle ζ can be implemented efficiently for selection rules whose outcomes change only finitely many times as the available funds increase from 0 to $+\infty$. Then, each u_i^1 is piecewise linear, and the interval R can be identified in time linear in the number of outcome changes. For the greedy rule that selects candidates by approval count while funds allow (as in our experiments), this number is $O(2^m)$ (see Theorem 5). Yet, it performs well on most of the examined real-world instances, enabling us to conduct extensive empirical evaluations.

We conclude the discussion on characterizations of equilibria with noting that the approach of proving Theorems 3 and 4 using LPs has a further advantage. It is possible to further add an objective or additional constraints to the program, e.g., to find the NE among all that maximizes the minimum utility or that splits the donations as equally as possible.

We now evaluate the quality of NE, using the social welfare that voters obtain from a given profile. As standard, we define it as the sum of individual voters' utility. For a profile $\mathbf{d}$ and an outcome $W_{\mathcal{P}}(\beta + \sum_{v_i \in V} d_i)$, we define the *social welfare* as $SW(\mathbf{d}) = \sum_{v_i \in V} u_i(\mathbf{d})$. We mainly consider three evaluation metrics:

(i) Social welfare in equilibria compared to that of the outcome when donations are not allowed.
(ii) Social welfare in equilibria compared to the optimal in (terms of welfare) profile.
(iii) Social welfare in equilibria compared to a theoretical upper bound on the maximum welfare.

As the first baseline for comparison, we consider the classic PB setting where donations are not allowed. We begin by analyzing how the welfare of NE compares to that setting theoretically, before quantifying the difference empirically to get a high-level sense of the gains from allowing donations.

Theorem 7. *When the projects in $\mathcal{P}$ are in decreasing order of votes and for cost utilities, allowing for donations can only result in an outcome of increased social welfare. For general utilities or arbitrary funding order this is not true, as an NE $\mathbf{d} \neq \mathbf{d_0}$ might produce lower social welfare compared to $\mathbf{d_0}$.*

Proof. Consider a strategy profile $\mathbf{d} \neq \mathbf{d_0}$ and say that some project $p_i \in \mathcal{P}$ is funded, which would not have been bought under $\mathbf{d_0}$. Then this does not affect any elected project p_j such that $j < i$. On the other hand, a project $p_k \in W(\beta)$ for $k > i$ might now be excluded from selection. However, the joint total cost of all such projects, say a set $\hat{C} \subseteq \mathcal{P}$, cannot be greater than $\text{cost}(p_i)$. or cost utilities, funding p_i results in the following change in social welfare:

$$|\text{supp}(p_i)|\text{cost}(p_i) - \sum_{p_k \in \hat{C}} |\text{supp}(p_k)|\text{cost}(p_k) \geqslant$$
$$\max_{p_k \in \hat{C}} \{|\text{supp}(p_k)|\}\text{cost}(p_i) - \sum_{p_k \in \hat{C}} |\text{supp}(p_k)|\text{cost}(p_k) \geqslant$$
$$\max_{p_k \in \hat{C}} \{|\text{supp}(p_k)|\} \sum_{p_k \in \hat{C}} \text{cost}(p_k) - \sum_{p_k \in \hat{C}} |\text{supp}(p_k)|\text{cost}(p_k) \geqslant 0,$$

hence, when allowing for donations, the sum of individual voters' utility cannot decrease.

For general utilities, consider an instance with 2 projects, where p_1 costs 4 and p_2 costs 3. The first project is approved by v_1 with $\gamma_1(p_1) = 1/2$ and the second is approved by v_2 with $\gamma_2(p_2) = 1$. Say that $\beta = 3$, and that tie-breaking is lexicographic, i.e. p_1 comes first in $\mathcal{P}$. With no donations, p_2 is being bought, hence $SW(\mathbf{d_0}) = 3$. However, allowing for donations, v_1 has the option to donate 1. Say that $\mathbf{d} = (1, 0)$. Then $W_{\mathcal{P}}(4) = \{p_1\}$.

It is not hard to see that $\mathbf{d}$ is indeed an NE and that it satisfies $SW(\mathbf{d}) = 2 < SW(\mathbf{d_0})$. For proving the same statement for cost utilities but without the assumption on having the candidates ordered by the number of approvals it suffices to consider the same instance but with $|\text{supp}(p_1)| = 1$ and $|\text{supp}(p_2)| = 2$. □

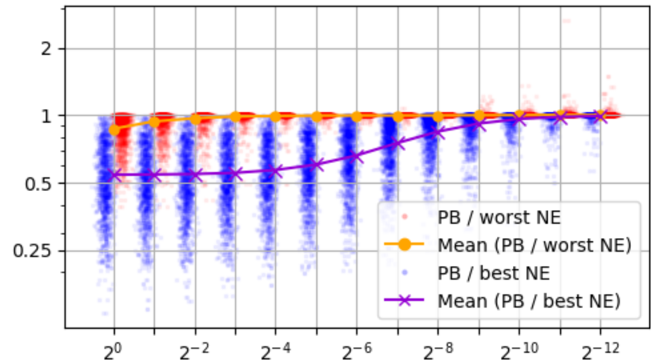

Figure 3: Welfare in the outcome without donations (PB) divided by the welfare in the best and worst NE, per instance for different values of γ. Blue (left) dots show the best NE, red (right) dots the worst. Purple cross and large yellow dot indicate means.

In datasets, we observe that the positive result of Theorem 7 showing that allowing donations improves social welfare continues to hold even beyond the cost utility setting. As shown in Figure 3, this is especially evident when considering the best equilibrium per instance, where we see an average improvement by a factor of around 2 for values of γ up to 2^{-4}, with the effect gradually diminishing as γ decreases. However, even for the worst equilibria, no significant welfare loss is observed for these values. In line with the negative example (Theorem 7) some instances may show a drop in welfare. This appears mainly for very small values of γ, but on average, for NE, the effect of allowing for donations then remains negligible.

Turning to a second natural baseline for evaluating the quality of NE, we study the classic notion of the Price of Anarchy (PoA), defined as the ratio between the maximal achievable social welfare and that of the worst Nash equilibrium. First, we show that, in principle, it can be arbitrarily large.

Theorem 8. *PoA is in $\Omega(nmt)$, even for cost utilities and when $\mathcal{P}$ is sorted in decreasing order of votes, for any $t \in \mathcal{R}_+$.*

Proof. Consider an instance with $m+1$ projects and $n+1$ voters. Voter, v_1 approves p_1. All other voters approve all projects in $\mathcal{P}$. Say that $cost(p_1) = 1/n{+}1$ and $cost(p_i) = t$ for all $i > 1$, for some fixed value of t. Let $\beta = 1/n{+}1$. Project p_1 has the maximum number of supporters, the rest projects are ordered lexicographically in $\mathcal{P}$. Donating 0 is an NE which results in a social welfare of 1. On the other hand, if we consider a profile $\mathbf{d}$ where v_1 donates tm then $SW(\mathbf{d}) = 1/n{+}1 - tm + nmt$. So the ratio between the welfare maximizing solution and the worst NE is at least $(n{-}1)mt{+}1/n{+}1$. □

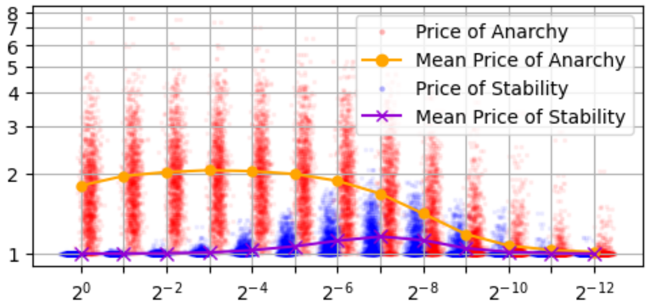

Figure 4: Welfare in the best and worst NE relative to the welfare maximizing solution, per instance for different values of γ. Blue (left) dots correspond to the best NE, red (right) dots to the worst NE. Purple cross and large yellow dot indicating means.

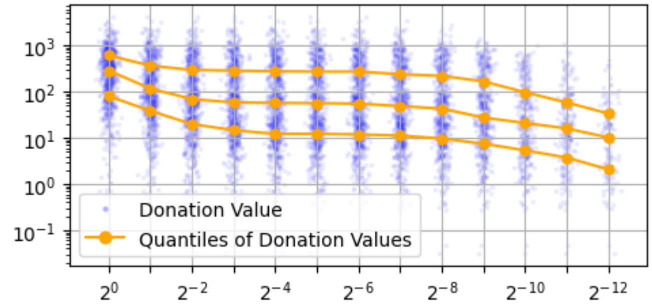

Figure 5: Minimum donation values in PLN currency across equilibria, shown per instance for different values of γ; large yellow dots indicate the corresponding 1st, 2nd, 3rd quartiles.

In Figure 4 we empirically show that, on average, the worst NE achieves at least around half of the optimal welfare for large and moderate values of γ, and comes very close to optimality for small values. Even beyond averages, the welfare loss in the worst equilibria is bounded by a small constant. Looking at the best equilibria (Price of Stability) yields even more encouraging results: on average, the achieved welfare is very close to the optimal one (with a slight deviation for moderate values of γ), while in almost all instances, the best equilibrium reaches at least half of the optimal welfare.

As a third and final metric for evaluating the quality of equilibria, and motivated by the fact that computing the exact optimal social welfare is nontrivial for general instances of the examined pledge game, we compare the social welfare in equilibria to an upper bound of the maximal achievable overall utility. We define this bound by neglecting the funding order of an instance, so we can fund any subset of projects exploiting both global and individuals' budgets, aiming solely to maximize the total satisfaction of the electorate. The solution to this optimization problem upper bounds the maximum social welfare under the ordering of the projects inferred by $\mathcal{P}$. Apart from serving as a meaningful benchmark for comparing utility in equilibria, the welfare maximization problem without a funding order is also theoretically appealing in its own right. Our results in this regard are discussed in the next paragraph.

For maximizing the welfare, in the absence of a funding order, we obtain a set of (both positive and negative) results regarding its computational complexity. First, we show that the problem is NP-hard even under cost utilities for the setting of bounded contributions. Nevertheless, we provide an FPT algorithm and an FPTAS for computing the optimal solution for general utilities. We identify a simple restriction of the pledge game that makes the problem polynomial-time solvable for the boundless setting. Since ignoring the funding order goes beyond the core model of our study, we defer the full discussion and the corresponding additional results to the full version due to space constraints. Likewise, we also defer the corresponding plot on NE quality, noting that the observed patterns are not substantially different from those in Figure 4. We also report how much the maximum welfare deviates from the discussed upper bound. This deviation is negligible for large values of γ, grows to an average of less than 5% for moderate values, and increases smoothly as γ decreases further.

We have observed that the theoretical guarantees for the examined framework can change depending on whether there is a bound on each voter's maximum contribution. Aiming to better understand the monetary considerations, we finally empirically explore how large the donations might need to be in an equilibrium. For each instance we calculate the *donation value* which is the minimal quantity such that there exists a non-zero NE in which every voter donates at most that much. We plot the donation value for each instance admitting a non-zero equilibrium, in Figure 5. We highlight that in half of the instances, donations between 10 and 300 PLN (roughly €2.5 to €70, depending on the value of γ) can lead to an equilibrium, while over a quarter of the instances admits an equilibrium with donation value of at most 100PLN for every γ. The number of instances where very small donations (on the order of 10PLN or even 1PLN≈0.25€) are enough to reach an NE is also substantial.

4 CONCLUSION

Our work reveals several cases in real-life public funding processes where allowing voters to donate can strongly influence the outcome. We explore this effect both theoretically and experimentally, primarily through the lens of Nash equilibria.

Considering group deviations is a natural next step, as groups with shared preferences could coordinate donations to steer outcomes in their favor. Exploring other equilibrium notions or analyzing theoretical guarantees in specific variants of the model are also promising directions. We intentionally separate the contribution allocation process from the election phase to remain manageable towards clear and interpretable insights. A promising direction is to relax this separation and study the interplay between the two stages. One could also consider addressing other limitations of our study, among those discussed in Section 2.

Our experiments focused on real-world PB instances, but applying similar analyses to other datasets (e.g., from grant programs in the blockchain universe) would be valuable. On the theoretical side, the computational complexity of deciding the existence of equilibria is an interesting open question. Our study serves as a clear suggestion to policymakers overseeing public funding decisions to consider enabling donations.

Acknowledgments. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 101002854). Grzegorz Lisowski acknowledges support by the European Union under the Horizon Europe project Perycles (Participatory Democracy that Scales). Georgios Papasotiropoulos and Grzegorz Pierczyński were 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. This work was supported by the Polish National Science Centre through grant 2021/42/E/HS4/00196.

Funded by the European Union

European Research Council
Established by the European Commission

Funded by the European Union

REFERENCES

[1] Haris Aziz and Aditya Ganguly. 2021. Participatory Funding Coordination: Model, Axioms and Rules. In *Proceedings of the International Conference on Algorithmic Decision Theory*.

[2] Haris Aziz, Sujit Gujar, Manisha Padala, Mashbat Suzuki, and Jeremy Vollen. 2023. Coordinating Monetary Contributions in Participatory Budgeting. In *Proceedings of the International Symposium on Algorithmic Game Theory*.

[3] Luca Bartocci, Giuseppe Grossi, Sara Giovanna Mauro, and Carol Ebdon. 2023. The Journey of Participatory Budgeting: A Systematic Literature Review and Future Research Directions. *International Review of Administrative Sciences* (2023).

[4] Dorothea Baumeister, Linus Boes, and Johanna Hillebrand. 2021. Complexity of Manipulative Interference in Participatory Budgeting. In *International Conference on Algorithmic Decision Theory*. Springer, 424–439.

[5] Georgios Birmpas, Evangelos Markakis, and Guido Schäfer. 2019. Cost sharing over Combinatorial Domains: Complement-Free Cost Functions and Beyond. In *Proceedings of the Annual European Symposium on Algorithms*.

[6] Niclas Boehmer, Piotr Faliszewski, Łukasz Janeczko, Dominik Peters, Grzegorz Pierczyński, Šimon Schierreich, Piotr Skowron, and Stanisław Szufa. 2023. Evaluation of Project Performance in Participatory Budgeting. *arXiv preprint arXiv:2312.14723* (2023).

[7] Florian Brandl, Felix Brandt, Matthias Greger, Dominik Peters, Christian Stricker, and Warut Suksompong. 2022. Funding Public Projects: A Case for the Nash Product Rule. *Journal of Mathematical Economics* (2022), 102585.

[8] Vitalik Buterin, Zoë Hitzig, and E. Glen Weyl. 2019. A Flexible Design for Funding Public Goods. *Management Science* (2019), 5171–5187.

[9] Jiehua Chen, Martin Lackner, and Jan Maly. 2022. Participatory Budgeting with Donations and Diversity Constraints. In *Proceedings of the AAAI Conference on Artificial Intelligence*.

[10] Primavera De Filippi, Sofia Cossar, Morshed Mannan, Kelsie Nabben, Tara Merk, and Jamilya Kamalova. 2024. *Report on Blockchain Governance Dynamics*. Technical Report. Project Liberty Institute, BlockchainGov.

[11] Deep Funding. 2023. Funding Round 3. https://deepfunding.ai/dfr3-results. Accessed: 2026-01-15.

[12] Piotr Faliszewski, Jarosław Flis, Dominik Peters, Grzegorz Pierczyński, Piotr Skowron, Dariusz Stolicki, Stanisław Szufa, and Nimrod Talmon. 2023. Participatory Budgeting: Data, Tools, and Analysis. In *Proceedings of the International Joint Conference on Artificial Intelligence*.

[13] Subash Chandra Bose Jaganathan, Murugesh Veerasamy, Azath Mubarakali, Chandramohan Dhasarathan, Velmani Ramasamy, R. Kalpana, and Ninoslav Marina. 2023. The Future of Web Crowdfunding: An Ethereum Blockchain Approach. In *Intelligent Technologies for Sensors*. Apple Academic Press.

[14] Philip Lazos, Evangelos Markakis, and Georgios Papasotiropoulos. 2025. Participatory Budgeting with Donations: The Case of Selective Voters. In *Proceedings of the European Conference on Artificial Intelligence*.

[15] Georgios Papasotiropoulos, Seyedeh Zeinab Pishbin, Oskar Skibski, Piotr Skowron, and Tomasz Wąs. 2025. Method of Equal Shares with Bounded Overspending. In *Proceedings of the ACM Conference on Economics and Computation*.

[16] Dominik Peters, Grzegorz Pierczyński, and Piotr Skowron. 2021. Proportional Participatory Budgeting with Additive Utilities. In *Proceedings of the Advances in Neural Information Processing Systems*.

[17] Project Catalyst. 2024. Funding Round 13. https://projectcatalyst.io/funds/13/voting-results. Accessed: 2026-01-15.

[18] Simon Rey and Jan Maly. 2023. The (Computational) Social Choice Take on Indivisible Participatory Budgeting. *arXiv preprint arXiv:2303.00621* (2023).

[19] Anwar Shah. 2007. *Participatory Budgeting*. World Bank.

[20] Nimrod Talmon and Piotr Faliszewski. 2019. A Framework for Approval-Based Budgeting Methods. In *Proceedings of the AAAI Conference on Artificial Intelligence*.

[21] Shiwen Wang, Chenhao Wang, Tian Wang, and Weijia Jia. 2023. Approval-Based Participatory Budgeting with Donations. In *International Computing and Combinatorics Conference*. Springer, 404–416.

[22] Wikipedia Contributors. [n.d.]. List of Participatory Budgeting Votes. https://en.wikipedia.org/wiki/List_of_participatory_budgeting_votes.