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

This work proposes a framework for public decision processes with a monetary component, where voters can pledge donations to support preferred projects and influence the allocation of public funds. It captures a range of real-world scenarios, including participatory budgeting, blockchain-based funding platforms, charitable programs, and crowdfunding campaigns. Through both theoretical and experimental evaluation, our study addresses questions concerning the existence, structure, number, computation, and quality of Nash equilibria.

1 Introduction

Participatory budgeting (PB) enables communities to directly influence the allocation of public funds [22, 23, 3]. 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 (publicly available within the PABULIB library [14]). The total budget was 826,000 Polish złoty (PLN), and the four most supported projects were as shown in the following table.

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 ($\approx$ 0.25€) each to secure its selection, hence directing a substantial portion of the budget toward a (highly approved) project they favor.

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

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, 16]. A commonly used method follows the greedy approach described earlier for PB. In Project Catalyst's 13th funding round [21], 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 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 approach can lead to missed opportunities, as funds remain untapped for projects that fall just short of their goals. The Alpha Colony video game is a prime example often cited in relevant discussions [12]. It ended its campaign just €23 shy of its €42,000 goal—an amount that translates to less than five cents for each of its supporters. Yet, the platform's rules meant that no funds were released or opportunities given for top-ups and the project was labeled unsuccessful. Kickstarter reports over 7,000 projects that raised between 81% and 99% of their goal [17]. However, this limitation isn't only an issue for unfunded projects. The board game Reavers of Midgard ran an incredibly successful campaign, raising over 3000% of its original goal [15]. Thanks to this, several improvements and additional features beyond the initial plan became possible. Still, the total funding fell just short of one final stretch goal which would have required only about €1 from each supporter, a negligible amount compared to what was collected.

To our knowledge, the diverse scenarios above have not previously been systematically identified; bringing them together constitutes a first, modest-but-concrete contribution of this work. These illustrate that, in numerous settings, outcomes could have been substantially improved if participants were allowed to donate. We assume that donations are permitted and explore aspects of strategic 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 specified 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. Notably, our work is the first to provide an experimental study on real-world data examining the effects of allowing donations in public decisions.

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

1. We provide theoretical and empirical insights into questions such as when non-zero contributions from voters are necessary to reach an NE, when non-zero equilibria exist, and when equilibria may fail to exist altogether.
2. We identify conditions for NE determination. Moreover, 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. [25] and Lazos et al. [18]. 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. [25], which suggest per-project donation bounds, our model—similarly to the one examined by Lazos et al. [18]—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 value 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. Strategic aspects in PB related to pure Nash equilibria have also been considered in [13], where the strategic actors are the project proposers.

2 Framework and Preliminaries

We will refer to an instance of the framework as a *pledge game* $\mathcal{G}$. In such an instance, 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*, denoted by $\text{cost}(p) \in \mathbb{R}_+$. 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 $\text{supp}(p) := \{v_i \in V : p \in A_i\}$. We assume that $|\text{supp}(p)| \geqslant 1$ for every project $p \in \mathcal{P}$. In an instance, we are also given a publicly available *global budget* $\beta \in \mathbb{R}_{\geqslant 0}$ such that $\beta < \sum_{p \in \mathcal{P}} cost(p)$ 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. This setup captures the process used in the vast majority of real-world participatory budgeting applications. Nevertheless, the funding order need not be exogenous, as it may be derived from approval counts or induced by other voting rules. Thus, our framework captures a variety of selection methods. For example, beyond the classic greedy rule that orders projects by approval counts, our model also captures variants based on approval-to-cost ratios or other measures of voter satisfaction. The method of Equal Shares, however,

does not directly fit our framework, as its selection process depends on the available budget, which in our setting also depends on voters' donations. Studying strategic donations under rules with fairness guarantees is a particularly interesting direction for future work.

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 $j \in \{0, 1, \ldots, m\}$ such that $W(B) = \{p_1, p_2, \ldots, p_j\}$, that is, the outcome consists of the first j projects in the order defined by $\mathcal{P}$, without skipping any; for $j = 0$ the prefix is the empty set.

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 [24, 20, 19]. 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:

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

Under this model, a voter may donate even to projects they do not approve of, if doing so benefits projects they approve of. We say that 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 $u_i(\mathbf{d}) \geqslant u_i((\mathbf{d}_{-i}, d_i'))$, in words, when no voter can increase their utility by unilaterally changing their donation.

Example 1. Take 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. Let the supporters of p_2 collectively donate 50. Then, $B = 150$ and p_1 is considered first in the funding order (this can be because, e.g., $|\text{supp}(p_1)| > |\text{supp}(p_2)|$). As 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 among them). If a single voter from $\text{supp}(p_2)$ contributes the entire amount of 50, 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 0. Suppose now that $\text{cost}(p_1) < 150$. Then, 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.

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, which will be analyzed in the following sections.

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. We say the setting is *single-approval* if each voter approves exactly one project, that is, $|A_i| = 1$ for every voter $v_i \in V$.

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. However, 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 each voter v_i be associated with a weighting 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, as said, in contrast, we call the case where $\gamma_i(p) = 1$ for every $v_i \in V$ 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 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 $(V, \mathcal{P}, \beta, (A_i)_{i \in [n]})$. Moreover, $(\mathbf{b}, (\gamma_i)_{i \in [n]})$ can also be included in the tuple, when they apply.

2.2 Experimental Setup and Limitations

Throughout our work, we complement the theoretical findings with empirical results. For this, we analyze data from PB elections available in PABULIB—a suitable library for the questions we study which contains a large and diverse collection of instances with varying election sizes, voters' preference structures, and project characteristics. We evaluate 676 instances for which our heuristic method for identifying all NE completed within a reasonable time frame (see Section 3 and Appendix B.2 for details). To allow meaningful comparisons across instances involving monetary contributions, we focused on elections conducted in the same currency—we restricted attention to 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: projects are sorted by the number of approvals they receive. 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. We focus primarily on small values of γ, as these are particularly natural in PB where the voters are generally not expected to be willing to cover the full cost individually, which is precisely the motivation for collective budgeting.

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—refer to [26] for multiple such examples. 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) 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. Since 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.

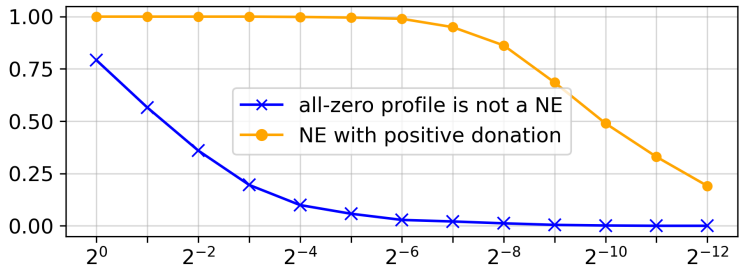


Figure 1: Fraction of instances that do not admit the all-zero profile as an NE (blue crosses) and that admit at least one profile other than the all-zero as an NE (yellow dots), for various values of γ.

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.*

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 now turn our attention to the most interesting case of $\beta > 0$. First 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, one should 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 for NE 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. 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.*

We saw that under cost utilities there always exists a Nash equilibrium 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. As expected, 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 $\boldsymbol{d}, \boldsymbol{d}'$ 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.*

So, a characterizing family of NE of a pledge game captures all distinct total donation amounts arising in equilibria. 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. Specifically, we show that we can compute and concisely depict the polytope 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 $\mathcal{G}$, every NE $\boldsymbol{d}$ with a nonempty outcome satisfies the following conditions: 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.*

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. Surprisingly, we show that the same holds for the single-approval setting under cost utilities, in the boundless setting.

Theorem 4. *In an instance $\mathcal{G}$ of the single-approval setting under cost utilities and boundless contributions, a strategy profile $\boldsymbol{d}$ inducing a nonempty outcome is an NE if and only if, letting $B = \beta + \sum_{v_i \in V} d_i$ and $p^* = \ell(W_{\mathcal{P}}(B))$, it holds that $W_{\mathcal{P}}(B)$ is a prefix of $\mathcal{P}$, and the following conditions hold: $\sum_{v_i \notin \text{supp}(p^*)} d_i = 0$; $B = \sum_{p \in W_{\mathcal{P}}(B)} \text{cost}(p)$; $d_i \leqslant \text{cost}(p^*)$ for every $v_i \in \text{supp}(p^*)$. Moreover, a characterizing family of NE can be computed in polynomial time.*

The positive result of Theorem 4 relies on a set of assumptions. We show that all these assumptions are in fact necessary; dropping either can lead to exponentially many budget-nonequivalent equilibria.

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.*

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 d_i^1 is a single-contributor profile if all contributions come from voter v_i, i.e., $d_i = \delta > 0$ and $d_j = 0$ for every $j \neq i$. Consider a profile $\mathbf{d}_i^1$. Let $W := W_{\mathcal{P}}(\beta + \delta)$ 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 $\boldsymbol{d}$ 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

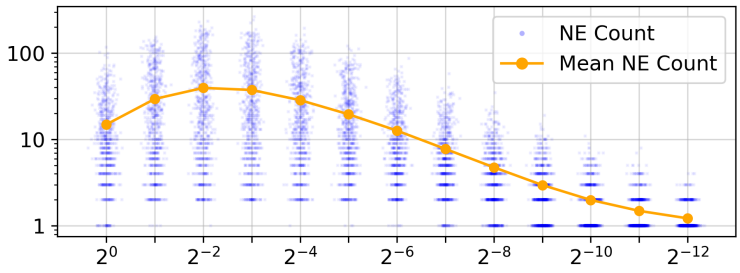


Figure 2: Number of budget-nonequivalent NE per instance for different values of γ. These numbers (NE count) per instance are shown as (small) blue dots, while the (large) yellow dots indicate averages.

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 by 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* (SW) 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 $SW(\mathbf{d}) = \sum_{v_i \in V} u_i(\mathbf{d})$. We 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, and

(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 cannot decrease social welfare. For general utilities or arbitrary funding order this is not true, as an NE $\boldsymbol{d} \neq \boldsymbol{d_0}$ might produce lower social welfare compared to $\boldsymbol{d_0}$.*

In real-life 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, as expected, 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 used in 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.

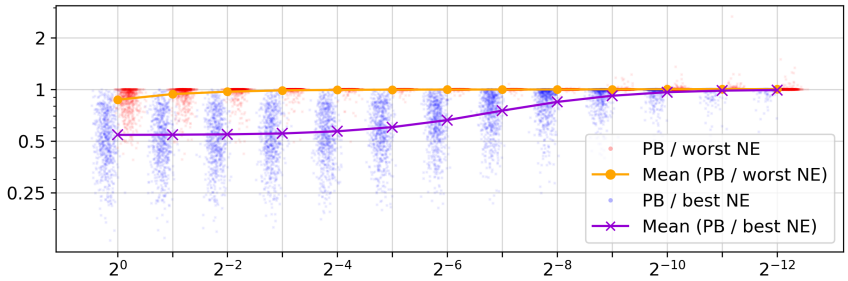


Figure 3: Welfare in the outcome without donations (PB) divided by the welfare in the best and worst NE, per instance for different values of γ. Blue (left) dots show the best NE, red (right) dots the worst. Purple cross and large yellow dot indicate their respective means.

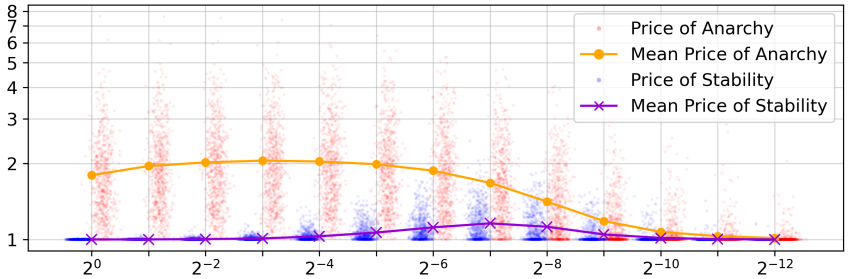


Figure 4: Welfare in the best and worst NE relative to the welfare maximizing solution, per instance for different values of γ. Blue (left) dots correspond to the best NE, red (right) dots to the worst NE. Purple cross and large yellow dot indicating their means.

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

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) per instance 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 appealing theoretically. Our results in this regard are discussed next.

On the Welfare Maximization Problem. 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. We show that the problem is NP-hard even under cost utilities for the setting of bounded contributions. Additionally, 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 Appendix B.1. Due to space constraints, we also defer the corresponding plot on NE quality to Appendix B.1 (Figure 6), noting that the observed patterns are not substantially different from those in Figure 4. In Appendix B.1 we also report how much the maximum welfare deviates from the discussed upper bound (Figure 7). 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.

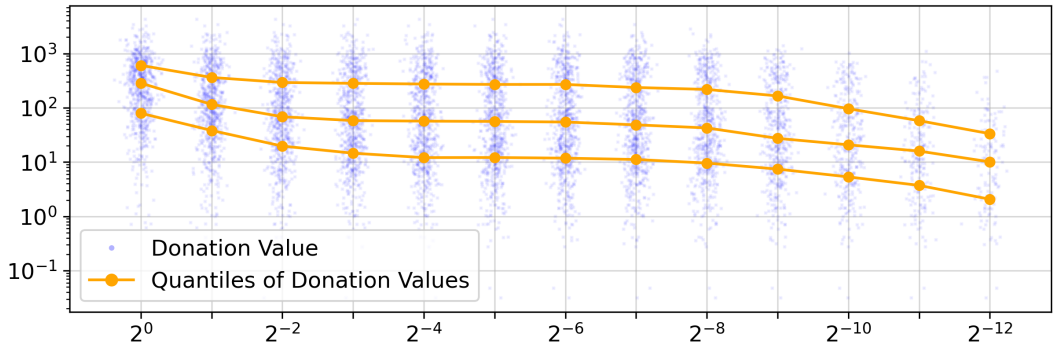


Figure 5: Minimum donation values in PLN currency across equilibria, shown per instance for different values of γ; large yellow dots indicate the corresponding 1st, 2nd, 3rd quartiles.

Concluding our work, we note that, following our results in Appendix B.1 (and also Theorem 2), 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 now 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. 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 is also a worthwhile direction. 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 datasets from other application domains (e.g., blockchain-based grant programs) would be valuable for examining whether the observed trends extend beyond PB. 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 The authors thank Piotr Skowron for his input on the design and motivation of the model studied in this work. 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 have been supported by the European Union (ERC, PRO-DEMOCRATIC, 101076570). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them. Additionally, Grzegorz Pierczyński was supported by Poland's National Science Center grant no. 2022/45/N/ST6/00271. Krzysztof Rogowski was supported by Poland's National Science Center grant no. 2021/42/E/HS4/00196.



References

[1] H. Aziz and A. Ganguly. Participatory funding coordination: Model, axioms and rules. In *Proceedings of the International Conference on Algorithmic Decision Theory*, 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*, 2023.

[3] L. Bartocci, G. Grossi, S. G. Mauro, and C. Ebdon. The journey of participatory budgeting: A systematic literature review and future research directions. *International Review of Administrative Sciences*, 2023.

[4] D. Baumeister, L. Boes, and J. Hillebrand. Complexity of manipulative interference in participatory budgeting. In *International Conference on Algorithmic Decision Theory*, pages 424–439. Springer, 2021.

[5] G. Birmpas, E. Markakis, and G. Schäfer. Cost sharing over combinatorial domains: Complement-free cost functions and beyond. In *Proceedings of the Annual European Symposium on Algorithms*, 2019.

[6] N. Boehmer, P. Faliszewski, Ł. Janeczko, D. Peters, G. Pierczyński, Š. Schierreich, P. Skowron, and S. Szufa. Evaluation of project performance in participatory budgeting. *arXiv preprint arXiv:2312.14723*, 2023.

[7] 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*, page 102585, 2022.

[8] V. Buterin, Z. Hitzig, and E. G. Weyl. A flexible design for funding public goods. *Management Science*, pages 5171–5187, 2019.

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

[10] P. De Filippi, S. Cossar, M. Mannan, K. Nabben, T. Merk, and J. Kamalova. Report on blockchain governance dynamics. Technical report, Project Liberty Institute, BlockchainGov, 2024.

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

[12] DreamQuest Games. Alpha colony: An exploration, building and trading game. https://www.kickstarter.com/projects/dreamquest/alpha-colony-an-exploration-building-and-trading-g, 2012. Accessed: 2026-01-15.

[13] P. Faliszewski, Ł. Janeczko, A. Kaczmarczyk, G. Lisowski, P. Skowron, S. Szufa, and M. Szwagierczak. Strategic cost selection in participatory budgeting. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.

[14] P. Faliszewski, J. Flis, D. Peters, G. Pierczyński, P. Skowron, D. Stolicki, S. Szufa, and N. Talmon. Participatory budgeting: Data, tools, and analysis. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 2023.

[15] Grey Fox Games. Reavers of midgard: The sequel to champions of midgard. https://www.kickstarter.com/projects/152730994/reavers-of-midgard-the-sequel-to-champions-of-midg, 2023. Accessed: 2026-01-15.

[16] S. C. B. Jaganathan, M. Veerasamy, A. Mubarakali, C. Dhasarathan, V. Ramasamy, R. Kalpana, and N. Marina. The future of web crowdfunding: An ethereum blockchain approach. In *Intelligent Technologies for Sensors*. Apple Academic Press, 2023.

[17] Kickstarter. Stats. https://www.kickstarter.com/help/stats, 2025. Accessed: 2026-01-15.

[18] P. Lazos, E. Markakis, and G. Papasotiropoulos. Participatory budgeting with donations: The case of selective voters. In *Proceedings of the European Conference on Artificial Intelligence*, 2025.

[19] G. Papasotiropoulos, S. Z. Pishbin, O. Skibski, P. Skowron, and T. Wąs. Method of equal shares with bounded overspending. In *Proceedings of the ACM Conference on Economics and Computation*, 2025.

[20] D. Peters, G. Pierczyński, and P. Skowron. Proportional participatory budgeting with additive utilities. In *Proceedings of the Advances in Neural Information Processing Systems*, 2021.

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

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

[23] A. Shah. *Participatory Budgeting*. Public Sector Governance and Accountability. World Bank, 2007.

[24] N. Talmon and P. Faliszewski. A framework for approval-based budgeting methods. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2019.

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

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

A Proofs of the Theoretical Results

Proof of Observation 1

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 $\mathbf{d}$ it holds that $B = d_i'$ 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. Then, v_1 would receive the utility of ${}^{cost(p_1)}/_2 > 0 = u_1(\mathbf{d_0})$. □

Proof of Theorem 2

Consider an instance $\mathcal{G}$ of the boundless setting and with cost utilities. Suppose first that voters donate nothing, i.e., they donate according to $\mathbf{d_0}$. Then, the global budget will fund some projects, precisely the set $W(\beta)$. Say that $\hat{c}$ is the first project (in the funding order of $\mathcal{P}$) that does not belong to $W(\beta)$. At the moment of consideration of $\hat{c}$ there was not enough budget to cover $cost(\hat{c})$. Say that r is the available budget at that point and $r' := \text{cost}(\hat{c}) - r$, i.e., the budget that remains to be covered in order to fund $\hat{c}$ and having no available funds after this purchase. The project $\hat{c}$ has at least one supporter by the corresponding modeling assumption. If the voters in $\text{supp}(\hat{c})$ pay r' in total then the outcome $W(\beta + r')$ will consist of all the projects in the ordering of $\mathcal{P}$ up until $\hat{c}$. This is because after the consideration and purchase of $\hat{c}$ the remaining budget will be zero and $\hat{c}$ was the first that has not been bought from the global budget under $\mathbf{d_0}$.

We will show that the donation profile $\mathbf{d}$ where d_i equals ${}^{r'}/_{|\text{supp}(\hat{c})|}, \forall v_i \in \text{supp}(\hat{c})$ or 0 otherwise, is a Nash equilibrium. In the setting of cost utilities, and under $\mathbf{d}$, every voter in $\text{supp}(\hat{c})$ pays ${}^{r'}/_{|\text{supp}(\hat{c})|} \leqslant r' = \text{cost}(\hat{c}) - r \leqslant cost(\hat{c})$ so their utility is non-negative. Consider first the case where one of them, say v_i reduces their donation, say to $d_i' < {}^{r'}/_{|\text{supp}(\hat{c})|}$. The change in the outcome will be as follows: the new outcome $W(\beta + r' - d_i + d_i'))$ will have all projects before $\hat{c}$ in the ordering, and perhaps some additional projects which come after $\hat{c}$ of cost at most $r + d_i' + \sum_{v_j \in (\text{supp}(\hat{c}) \setminus \{v_i\})} d_j$.

The change in utility for voter i will be as follows, compared to their utility under $W(\beta + r')$: it will increase by $d_i - d_i'$ due to the decrease of their payment; it will decrease by $cost(\hat{c})$ due to the fact that $\hat{c}$ is not bought anymore as the obtained funds are not enough for it and then increase by at most $r + d_i' + \sum_{v_j \in (\text{supp}(\hat{c}) \setminus \{v_i\})} d_j$, due to the cost utilities setting and the fact that after $\hat{c}$ this is the remaining amount of money that can be spent on purchasing further projects. Therefore, the change in their utility will be upper bounded by

$$d_i - d_i' - cost(\hat{c}) + r + d_i' + \sum_{v_j \in (\text{supp}(\hat{c}) \setminus \{v_i\})} d_j = -cost(\hat{c}) + r + \sum_{v_j \in \text{supp}(\hat{c})} d_j =$$

$$-cost(\hat{c}) + r + |\text{supp}(\hat{c})| \frac{r'}{|\text{supp}(\hat{c})|} = -cost(\hat{c}) + r + r' = -cost(\hat{c}) + r + cost(\hat{c}) - r = 0.$$

On the other hand, if any voter increases their payment to a value larger than the one specified by $\mathbf{d}$, they need to buy a project on their own to increase their satisfaction further. This is due to the fact that $W(\beta + r')$ is a prefix and after the purchase of the last project of it, no remaining money is left. Buying a project on their own will not give them additional utility.

Next, it is trivial to show that in the bounded setting or for general utilities the zero profile might be the only one. Specifically, for the bounded setting, consider the case of a single project p_1 such that $\text{cost}(p_1) > \beta + \sum_{i \in [n]} b_i$. For general utilities, consider the case of a single project p_1 and one voter v_1 approving p_1 and having $\gamma_1(p_1) = {}^1/_2$ and $\beta < {}^{\text{cost}(p_1)}/_2$. In the first instance, the only feasible solution is the empty one and the only NE resulting to this solution is the all-zero profile. In the second instance, asking v_1 to donate any positive amount will either not be sufficient to fund p_1 so they will have incentive to donate less. Otherwise, they need to pay strictly more than ${}^{\text{cost}(p_1)}/_2$ to fund p_1 resulting again in a negative utility, with an incentive to decrease to 0 donation once again.

We now turn to proving that there are instances with no equilibria under the bounded setting or general utilities. First, for the bounded setting. Consider a case of two projects p_1 and p_2 such that $\text{cost}(p_1) = 3, \text{cost}(p_2) = 2$ and two voters v_1 (supporting only p_1) and v_2 (supporting only p_2) with $b_1 = 1$ and $b_2 = 1.5$. Say that $\beta = 0.5$. In this instance the maximum budget, say B, that can be used is 3, consisting of $b_1 + b_2 + \beta$. Therefore, one cannot fund both projects. If $B < 2$ then v_2 would like to donate more in order to fund p_2. If $3 > B > 2$ then v_1 is donating but p_1 is not funded, so they would prefer to donate nothing. If $B = 3$, this can only happen if v_2 is donating but then, as p_2 is not funded, they prefer to deviate to donating 0. Finally, we consider the case where $B = 2$. If only v_2 is donating, then v_1 prefers to donate 1 more to reach a total budget of 3 and fund p_1. On the other hand, if v_1 is donating as well, the fact that only p_2 can get funded under $B = 2$ means that they would prefer to decrease their donation to 0. Therefore, there is no NE in the described instance.

Finally, we consider the case of general utilities and we describe an instance with no NE as well. Consider again the case of two projects p_1 and p_2 such that $\text{cost}(p_1) = 3, \text{cost}(p_2) = 2$ and two voters v_1 (supporting only p_1) and v_2 (supporting only p_2) with $\gamma_1(p_1) = \gamma_2(p_2) = 1/2$. Say that $\beta = 1.4$. Let B be the total budget that can be gathered by both the public budget and voters' contributions. Hence it holds that $B \geqslant 1.4$. If $1.4 \leqslant B < 2$ then no project is bought. But in this case v_2 would like to donate more in order to fund p_2 by paying at most 0.6; this would result in a strictly positive utility. Say that $3 > B \geqslant 2$. In this case, if v_1 is donating zero, then they would prefer to increase their contribution in order to fund p_1 by paying at most 1; if v_1 is paying, as p_1 is not funded, it is better to reduce to 0 donation. In the case that $5 > B \geqslant 3$, it holds that only project p_1 is getting funding, so if v_2 donates they prefer to reduce to a zero donation; if v_2 does not donate, v_1 has to donate more than 1.6 but this results in negative utility so they prefer to donate nothing as well. We conclude with the case of $B \geqslant 5$. Then, both projects are being funded, but, in this case, the two voters need to contribute in total at least 3.6, so at least one of them contributes at least 1.8; this results in a negative utility, hence they would prefer to deviate to donating nothing. □

Proof of Theorem 3

For any project $p_i \in \mathcal{P}$, let $P(p_i) = \{p_j \in \mathcal{P} : j \leqslant i\}$ denote the set of projects up to and including p_i. We also define $P(p_0) = \emptyset$. Note that in any Nash equilibrium $\mathbf{d}$, if $\beta + \sum_{v_i \in V} d_i = B$ then $W_{\mathcal{P}}(B) = P(p)$ for some $p \in \mathcal{P}$. This follows from the funding rule of the crowdfunding setting: each project can only be funded if all projects of higher priority are already funded.

Let $\mathbf{d}$ be a Nash equilibrium where the last funded project under it is p^*. Then every voter with positive contribution must support p^*. To prove this, assume otherwise, i.e., that a voter $v_i \notin \text{supp}(p^*)$ has $d_i > 0$. By reducing d_i by an infinitesimal positive amount of money, only project p^* becomes unfunded. Since v_i does not approve p^* and they reduced their cost, the acquired utility from the new outcome is larger than before, contradicting the fact that $\mathbf{d}$ is a Nash equilibrium. Consequently, every Nash equilibrium $\mathbf{d}$ can be characterized by a unique quantity which denotes the total voters' contribution and from which an outcome $W_{\mathcal{P}}(B)$ can be directly inferred which equals $P(\ell(W_{\mathcal{P}}(B)))$. Moreover, this contribution needs to be made only from supporters of $\ell(W_{\mathcal{P}}(B))$, or a subset of them, and should be equal to $\sum_{p_i \in W_{\mathcal{P}}(B)} \text{cost}(p_i) - \beta$. Since there are $O(m)$ possible prefixes of $\mathcal{P}$ that can work as outcomes $W_{\mathcal{P}}(B)$, we can specify a characterizing family of NE with at most $O(m)$ distinct strategy profiles. Knowing the structure of each NE, the total cost of each and an upper bound on their number, we proceed with formulating a linear program that induces the polytope of all possible NE. First, consider an NE profile $\mathbf{d}$ and say that p_t is the last funded project under $\mathbf{d}$ and that $D(\mathbf{d}) \subseteq \text{supp}(p_t)$ is the set of supporters of p_t that contribute a positive amount in $\mathbf{d}$. As noted earlier, each voter not in $D(\mathbf{d})$ donates nothing. If any voter v_i increases their donation, this can only lead to the purchase of projects after p_t in the ordering of $\mathcal{P}$. By the fact that the total available budget that comes from the public budget increased by the contribution of voters in $D(\mathbf{d})$ equals exactly $\sum_{p_i \in W_{\mathcal{P}}(B)} \text{cost}(p_i)$ we have that any funded project after p_t can only be funded by individual contribution of v_i and only them. Therefore, they do not contribute positively to their utility, since the acquired utility of v_i from having an approved project p_j in the solution funded only by themselves is always non-positive. Therefore no profitable deviation of v_i can come from an increase of payment.

Consider deviations that decrease the payment of v_i. The only such deviations that need to be considered are exactly those that lead to the exclusion of a project from selection. Hence, such a deviation will lead to a change in the set of bought projects, but the outcome will always need to remain a prefix. It suffices to consider, for each shorter prefix (including the empty one), the minimum nonnegative contribution by which a voter can obtain that outcome, whenever such a deviation is feasible.

In summary, for each $t \in [m]$, consider funding exactly $P(p_t)$ by donating exactly $\sum_{p_i \in P(p_t)} \text{cost}(p_i) - \beta$. A strategy profile $\mathbf{d}$ constitutes a Nash equilibrium with this outcome if and only if the following constraints are simultaneously satisfied:

- Participation constraint: $d_i = 0$ for all $i \notin \text{supp}(p_t)$
- Budget constraint: $B = \beta + \sum_{v_i \in \text{supp}(p_t)} d_i = \sum_{p \in P(p_t)} \text{cost}(p)$ and $P(p_t) = W_{\mathcal{P}}(B)$
- Incentive compatibility: For every $v_i \in \text{supp}(p_t)$ and every $j \in \{0, \ldots, t-1\}$, let S_j denote the total cost of projects in $P(p_t) \setminus P(p_j)$ and L_{ij} the utility that voter v_i derives from these projects, with $P(p_0) = \emptyset$. Whenever $L_{ij} < S_j$, we require $d_i \leqslant \max L_{ij}, S_{j+1}$, where $S_t = 0$.

To see that the incentive-compatibility constraints are sufficient, consider a voter v_i reducing their contribution so that the funded outcome becomes $P(p_j)$, for some $j < t$. Such a deviation is possible only if $d_i > S_{j+1}$, and the maximum amount that the voter can save is $\min d_i, S_j$. Since this deviation decreases their utility from funded projects by L_{ij}, it is profitable exactly when $\min d_i, S_j > L_{ij}$. The stated constraints exclude precisely these deviations. They also cover deviations resulting in the empty outcome by taking $j = 0$. We additionally check whether the all-zero profile is an NE, as such an equilibrium may leave part of the public budget unspent and hence may not be captured by the above LPs. This can be done in polynomial time by checking whether any voter can profitably fund an additional prefix through a unilateral donation.

For each $t \in [m]$, solving the linear program formulated above yields either: the polytope of all budget-equivalent Nash equilibria funding $P(p_t)$, or infeasibility, meaning that no NE funding $P(p_t)$ exists. The union of feasible polytopes for each $t \in [m]$, together with the all-zero profile whenever it is an NE, forms the class of all NE of the instance. Since each linear program has polynomial size and there are $O(m)$ such programs that need to be solved, the entire class of NE is polynomially computable.

For the bounded-contributions setting, we additionally impose to the linear program the constraint $0 \leqslant d_i \leqslant b_i$, for each voter v_i. □

Proof of Theorem 4

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. By the individual-rationality condition, this voter obtains nonnegative utility under d. 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, 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 polytope of Nash equilibria can be found by specifying the constraints from the proof of Theorem 3. □

Proof of Theorem 5

For the first statement, assume $\gamma = 1/2$. Consider m projects of cost $3^m, 3^{m-1}, \ldots, 3$. There are $3^{m+1} - 3$ voters who support the last in order project, namely p_m, (and only this). All projects in $\{p_1, \ldots, p_{m-1}\}$ have a single supporter, so the total number of voters is $3^{m+1} + m - 4$. Say also that $\beta = 0$. Asking all voters that support the last project to buy any subset of projects (together with the last one) by splitting the cost equally among them is an NE. Whichever set of projects they buy, they get a utility that is at least

$$3/2 - \frac{3^m + 3^{m-1} + \cdots + 3}{3^{m+1} - 3} = 3/2 - 1/2 = 1.$$

Therefore, none of them has an incentive to reduce their spending as then they will get a non-positive utility because their approved project will not be selected. Increasing their donation can only result in funding further projects they do not approve, so, again, this does not improve their utility. Second, suppose that 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 $3^j/2 - d_i$. It holds that $d_i \geqslant 3^j - (3^{j-1} + 3^{j-2} + \cdots + 3)$. Therefore, their utility is upper bounded by

$$3^j/2 - 3^j + (3^{j-1} + 3^{j-2} + \cdots + 3) = -\frac{3^j}{2} + \sum_{i=1,\ldots,j-1} 3^i = -\frac{3^j}{2} + \frac{3 - 3^j}{2} = 3/2 - 3^j < 0.$$

Hence, any subset of projects can be bought at an equilibrium, which proves the statement.

For the second statement, consider an instance of cost utilities with 1 voter and m projects where $cost(p_i) = 2^{m-i+1}$, for $i \in [m]$. Every possible subset of the set of candidate projects is a feasible solution by asking the voter to donate the following: for the set of projects with indices in $X \subseteq [m]$ they need to donate $\sum_{i \in X} 2^{m-i+1}$ and this will result in a utility of 0, while donating any more or less than that will not result in increasing their utility, hence it is an equilibrium. □

Proof of Proposition 6

Consider a donation profile $\mathbf{d}$ and assume that it is an NE. For contradiction, assume that the condition from the statement does not hold. In particular, first say that $u_i^1(\delta) > \texttt{sup}_{x \in [0,b_i]} u_i^1(\delta - d_i + x)$. This is trivially false since the right-hand side is a supremum over a set of values which contains the left-hand side. Therefore for some voter v_i the following holds: $u_i^1(\delta) < \texttt{sup}_{x \in [0,b_i]} u_i^1(\delta - d_i + x)$. Then there exists some $x^* \in [0, b_i]$ with $x^* \neq d_i$ such that $u_i^1(\delta) < u_i^1(\delta - d_i + x^*)$. We will show that the deviation from d_i to x^* is beneficial for v_i. Recall that $\delta = \sum_{v_i \in V} d_i$. We will use $\mathbf{0}$ to denote $\mathbf{d_0}$. Then:

$$u_i((\mathbf{d}_{-i}, x^*)) \overset{(\star)}{=} u_i((\mathbf{0}_{-i}, \sum_{j \neq i} d_j + x^*)) - \sum_{j \neq i} d_j = u_i^1(\sum_{j \neq i} d_j + x^*) - \sum_{j \neq i} d_j >$$
$$u_i^1(\sum_j d_j) - \sum_{j \neq i} d_j = u_i((\mathbf{0}_{-i}, \sum_j d_j)) - \sum_{j \neq i} d_j \overset{(\star)}{=} u_i((\mathbf{d}_{-i}, d_i)),$$

where the relations marked with $(\star)$ follow from the fact that the pair of profiles $(\mathbf{d}_{-i}, x^*)$ and $(\mathbf{0}_{-i}, \sum_{j \neq i} d_j + x^*)$ as well as the pair of profiles $(\mathbf{d}_{-i}, d_i)$ and $(\mathbf{0}_{-i}, \sum_j d_j)$, amount to the same outcome, but in the latter ones the voter v_i donates $\sum_{j \neq i} d_j$ more, hence their utility lowers by that amount. The obtained inequality that $u_i((\mathbf{d}_{-i}, x^*)) > u_i((\mathbf{d}_{-i}, d_i))$ contradicts the assumption that $\mathbf{d}$ is an NE.

We now move to the reverse direction and we assume that the condition $u_i^1(\sum_j d_j) = \texttt{sup}_{x \in [0,b_i]} u_i^1(\sum_{j \neq i} d_j + x)$ holds for a profile $\mathbf{d}$. Assume that $\mathbf{d}$ is not an NE. There exists some voter v_i such that $u_i((\mathbf{d}_{-i}, x^*)) > u_i(\mathbf{d}) = u_i((\mathbf{d}_{-i}, d_i))$. Then:

$$u_i^1(\sum_{j \neq i} d_j + x^*) = u_i((\mathbf{0}_{-i}, \sum_{j \neq i} d_j + x^*)) = u_i((\mathbf{d}_{-i}, x^*)) + \sum_{j \neq i} d_j >$$
$$u_i((\mathbf{d}_{-i}, d_i)) + \sum_{j \neq i} d_j = u_i((\mathbf{0}_{-i}, \sum_j d_j)) = u_i^1(\sum_j d_j).$$

Since x^* has to be a valid donation for voter v_i it has to be in $[0, b_i]$, which contradicts the assumption $u_i^1(\sum_j d_j) = \texttt{sup}_{x \in [0,b_i]} u_i^1(\sum_{j \neq i} d_j + x)$, and concludes the proof. □

Proof of Theorem 7

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 $\text{cost}(p_i)$. Hence, for cost utilities, funding p_i results in the following change in the total utility obtained from funded projects:

$$\begin{aligned} |\text{supp}(p_i)|\text{cost}(p_i) - \sum_{p_k \in \hat{C}} |\text{supp}(p_k)|\text{cost}(p_k) \geqslant \\ \max\{|\text{supp}(p_k)| : p_k \in \hat{C}\}\text{cost}(p_i) - \sum_{p_k \in \hat{C}} |\text{supp}(p_k)|\text{cost}(p_k) \geqslant \\ \max\{|\text{supp}(p_k)| : p_k \in \hat{C}\} \sum_{p_k \in \hat{C}} \text{cost}(p_k) - \sum_{p_k \in \hat{C}} |\text{supp}(p_k)|\text{cost}(p_k) \geqslant 0, \end{aligned}$$

The same argument holds when replacing $|\text{supp}(p)|$ by $|\text{supp}(p)| - 1$ for every project p. Moreover, in any NE with positive donations, the entire available budget must be spent, as otherwise some voter could profitably reduce their contribution. Consequently, accounting for the cost of donations, the social welfare cannot decrease compared to the all-zero profile.

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 an NE and that it satisfies $SW(\mathbf{d}) = 1 < SW(\mathbf{d_0})$. For proving the same statement for cost utilities but without the assumption that projects are ordered by approval counts, it suffices to consider the same instance with $\gamma_1(p_1) = 1$ and $|\text{supp}(p_1)| = 1$, $|\text{supp}(p_2)| = 2$. □

Proof of Theorem 8

Consider an instance with $m + 1$ projects and $n + 1$ voters. The first voter, v_1 approves the first project 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 $SW(\mathbf{d}) = 1 - tm + nmt$. So the ratio between the welfare maximizing solution and the worst NE is at least $(n - 1)mt + 1$. □

B Additional Results and Details

B.1 The Social Welfare Maximization Problem in the Absence of a Funding Order

In this section, we discuss the social welfare maximization problem in the absence of a funding order. We first show that when the requirement to follow the funding order is removed—and thus substantially deviating from the core model presented in Section 2—the problem becomes easy to solve in the boundless case: the optimal outcome simply includes each project p which satisfies $\sum_{v_i \in V} \gamma_i(p) \geqslant 1$. In the special case of cost utilities, this corresponds to selecting all projects supported by at least one voter. Any donation profile $\mathbf{d}$ that meets the required budget can be used to realize this outcome.

Theorem 9. *In the boundless setting, and in the absence of a funding order, a profile $\mathbf{d}$ maximizing social welfare can be computed in polynomial time as long as $\beta \leqslant cost(W)$, for $W := \{p \in \mathcal{P} : \sum_{v_i \in V} \gamma_i(p) \geqslant 1\}$. In particular the outcome maximizing social welfare of such an instance is the set W and any distribution of donations across voters that realizes W and satisfies $\beta + \sum_{v_i \in V} d_i = \sum_{p \in W} cost(p)$ maximizes social welfare.*

Proof. Consider an arbitrary set of projects W, bought after voters donating $\mathbf{d}$. We compute the social welfare under W and $\mathbf{d}$ and we will then determine the set W that maximizes welfare and how the overall donations should be distributed among the voters of the electorate. It holds that

$$SW(\mathbf{d}) = \sum_{v_i \in V} (\sum_{p \in W} \gamma_i(p) cost(p) - d_i) = \sum_{v_i \in V} \sum_{p \in W} \gamma_i(p) cost(p) - \sum_{v_i \in V} d_i.$$

In order to compute $\sum_{v_i \in V} d_i$, we observe that the total payment needs to cover the cost of all projects in W decreased by the global budget. Hence, $\sum_{v_i \in V} d_i \geqslant \sum_{p \in W} \text{cost}(p) - \beta$. Consequently,

$$SW(\mathbf{d}) \leqslant \beta + \sum_{p \in W} \text{cost}(p) \left(\sum_{v_i \in V} \gamma_i(p) - 1 \right). \tag{2}$$

For the cost utilities setting, we have that $SW(\mathbf{d}) = \beta + \sum_{p \in W} cost(p)(|\text{supp}(p)| - 1)$, therefore the set of projects that maximizes social welfare includes all projects that have at least one supporter. Similarly, for general utilities, it includes each project p satisfying $\sum_{v_i \in V} \gamma_i(p) \geqslant 1$.

By assumption, the total cost of this set is at least β. Therefore, in the boundless setting, we can choose donations satisfying $\beta + \sum_{v_i \in V} d_i = \sum_{p \in W} \text{cost}(p)$, making the bound tight. Hence, this outcome maximizes social welfare, independently of how the required donations are distributed among voters. □

The proof of Theorem 9 highlights the significance of the boundless setting in achieving polynomial-time solvability for the welfare maximization problem. This is because, in this case, once the outcome W maximizing social welfare is identified, it is always feasible to assign donations that realize W. This is no longer true in the bounded setting: the welfare-maximizing outcome may not be achievable within voters' donation limits. When this happens, one must search for the welfare-maximizing outcome among those that are feasible, and determine the corresponding donation profile to support it. This leads to the computational intractability established in the next result which reveals an interesting contrast between the two settings. While this hardness is already known from the work of Aziz et al. [2] (and applies also to the case with no global budget and without a funding order), it pertains to the setting of general utility functions. Since cost utilities are central to our study, we establish an NP-hardness for this specific case as well.

Theorem 10. *In the bounded-contributions setting and in the absence of a funding order, a profile* $\mathbf{d}$ *maximizing social welfare is NP-hard to be computed, even for cost utilities and a single voter.*

Proof. We consider the decision version of the welfare maximization problem in the bounded setting. Given a set of projects $\mathcal{P}$, each with an associated cost, a set of voters V with approval ballots and cost utilities, a public budget β and a parameter θ, the problem asks whether there exists an outcome $W \subseteq \mathcal{P}$ whose total cost does not exceed $\beta + \sum_{v_i \in V} b_i$ and for which there exists a feasible donation profile yielding social welfare at least θ. We prove the result by reducing from the PARTITION problem: Given a set of positive integers $E = \{a_1, a_2, \ldots, a_m\}$, determine whether there exists a set $S \subseteq E$ such that $\sum_{a_i \in S} a_i = \alpha$, for $\alpha = \sum_{a_i \in E} a_i/2$.

Given an instance I of PARTITION we introduce an instance of the welfare maximization problem with m projects where the i-th project has a cost of a_i. We consider a single voter v_1 who approves all projects and we set $\beta = \theta = \alpha$ and $b_1 = 0$.

For the forward direction, assume that I is a yes-instance of PARTITION, meaning there exists a subset $S \subseteq E$ such that $\sum_{a_i \in S} a_i = \alpha$. Consider the corresponding set of projects in the welfare maximization instance. This set, say W has total cost exactly α, so it is affordable by the global budget while respecting the contribution bound of v_1. Moreover, $u_1(\mathbf{d}) = cost(W) = \alpha \geqslant \theta$, meeting the required threshold. Therefore, the constructed welfare maximization instance is also a yes-instance.

For the reverse direction, assume that I is a no-instance, meaning that for every subset $S \subseteq E$, we have $\sum_{a_i \in S} a_i \neq \alpha$. Consider any corresponding set of projects in the welfare maximization instance. If it holds that $\sum_{a_i \in S} a_i > \alpha$, then this set exceeds the global budget β and is therefore not affordable without violating the contribution bound of v_1. On the other hand, if $\sum_{a_i \in S} a_i < \alpha$, then, while the set is affordable, it holds $u_1(\mathbf{d}) < \alpha = \theta$, failing to meet the required utility threshold. Therefore, no set of projects provides a feasible solution to the constructed instance, hence it is a no-instance. □

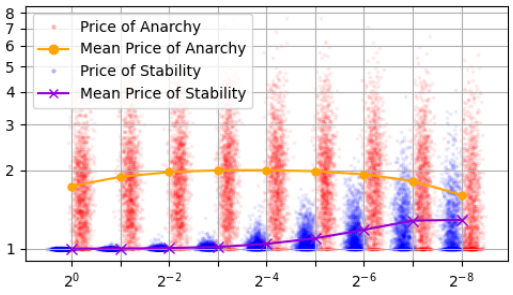


Figure 6: Comparison of welfare in the best and worst NE relative to the upper bound of the welfare maximizing solution defined according to Theorem 9, per instance for different values of γ. Blue (left) dots correspond to the best NE, red (right) dots to the worst NE, with the purple crosses and large yellow dots indicating their respective means.

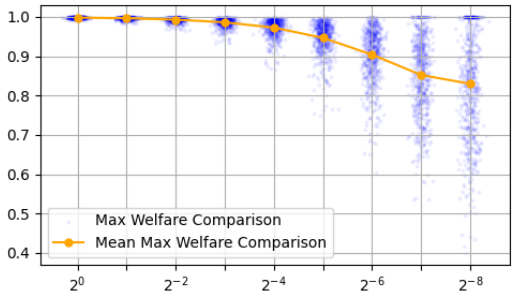


Figure 7: Blue dots show the ratio of the maximum social welfare to its upper bound defined according to Theorem 9, per instance for different values of γ, large yellow dots indicate their respective means.

We conclude the part on the setting without projects' ordering by visualizing how far the maximum achievable social welfare in this setting (computed as per Theorem 9) is from the best and worst NE of instances (Figure 6), as well as to the welfare of the solution maximizing the social welfare in the setting that a funding order is present (Figure 7). We note that for the examined values of γ up to 2^{-8}, only very few instances of the examined dataset violated the assumption in Theorem 9, so the welfare-maximizing solution (without a funding order) could be efficiently computed either as in Theorem 9 or via brute force. For smaller γ, the computation became significantly more demanding, so we excluded those values.

B.2 Further Experimental Details

Our experiments were conducted on a Debian Linux system (kernel 6.1.140) equipped with a 13th Gen Intel Core i5-1335U processor (12 threads) and 16GB of RAM, using Python (v3.13.1). The full experimental process completes in under two hours.

For the dataset, we began with all PABULIB instances not marked as fully funded or experimental, based on the state of the library as of 28-06-2025. To ensure meaningful comparisons of donation magnitudes, we excluded 140 instances from PB elections not in Polish złoty, as the vast majority of the library's instances are from Poland. One additional instance was omitted due to a peculiar project (proposed in the 2020 Gdynia Grabówek PB election and removed by the organizers after the election) with a cost of approximately 1 billion, which significantly distorted results. Finally, 48 more instances were excluded due to prohibitively long-running times under our procedures; since our experiments are primarily an illustrative proof of concept aimed at qualitative insights, developing optimized implementations to handle more computationally demanding instances is a valuable direction for future work. More precisely for each instance we calculated a value referred to as "estimated memory usage", which is equal to the number of feasible subsets—a notion that will be described later—multiplied by the number of voters and by 32 bytes. The instances excluded in this last step are the ones for which the estimated memory usage exceeded 512 MiB. We include the list of all processed instances in the supplementary material, along with the code used for the experiments.

Algorithmic approach. In the main text, we claim that the greedy selection rule can change the outcome at most 2^m times as the budget increases from 0 to $+\infty$. This follows from the observation that the greedy rule never revisits a previously selected project: formally, there is no instance and budgets $b_1 < b_2 < b_3$ such that $W(b_1) = W(b_3) \neq W(b_2)$. Thus, the number of possible outcomes is at most 2^m. While this upper bound is tight (see Theorem 5), we did not encounter such worst-case instances in our experiments and were able to process nearly all PABULIB instances.

We refer to the budget β specified in the description of a pledge game as the *initial budget*, and the sum of the initial budget and all voter donations as the *total budget*. Subsets of projects that can be selected by the greedy rule for some total budget are called *feasible subsets*, and we order these by their total cost, denoted as *feasible costs*. Nash equilibria can only exist for total budget values equal

to the initial budget or a feasible cost; for other values, leftover money and nonzero donations would always create an incentive to deviate. Thus, there are only finitely many (up to 2^m) total budget values for which NE can exist, and we process them one by one.

The procedure for analyzing a single total budget B is as follows:

1. Identify the set S of feasible subsets and their total costs for the instance.
2. For each voter, check if increasing their donation could improve their utility (if so, no NE exists for profile with total budget B).
3. For each voter, determine an upper bound on their possible donation in any NE at total budget B.
4. Sum these bounds and check if the total (plus the initial budget) suffices to reach B (if not, no NE exists).
5. If all checks pass, an NE exists for total budget B.

Details of these steps are as follows:

Step 1: Identifying feasible subsets is straightforward for the greedy rule and can be done in $O(m2^m)$ time. We repeatedly run an extended greedy algorithm that, for each budget, returns both the selected projects and the minimal budget increase needed to change the outcome. Starting from $b_0 = 0$, we find $W(b_0)$ and the minimal increment α_0. We then proceed to $b_1 = b_0 + \alpha_0$, and so on, until all projects are selected.

Steps 2 and 3: These rely on analyzing the utility function u_i^1 for each voter v_i. The function u_i^1 is piecewise linear, with breakpoints at the feasible costs. By examining each linear segment, we can determine the maximal interval R such that $u_i^1(B) = \sup_{x \in R} u_i^1(B - d_i + x)$. If R is finite on the right, voter v_i can always increase their donation in any profile with total budget B to improve their utility since we don't consider any bounds on their donations in experiments, so no NE exists for B (step 2). The left endpoint of R gives an upper bound on d_i in any NE at B; if d_i is too large ($B - d_i \notin R$), v_i would have an incentive to decrease their donation (step 3). This procedure serves as an explicit implementation of the ζ oracle described in the main text, specialized for the greedy selection rule.

Step 4: Summing the upper bounds from all voters gives the maximum total donations possible in any NE at B. If the initial budget plus these donations is less than B, no NE exists (step 4). If all checks pass, we obtain a system of inequalities characterizing all NE profiles at B.

Experimental results presented in the paper are based mainly on the characterization of NE described above with only a little bit of transformations on top of it. In particular:

- Figure 1 and Figure 2 are obtained by iterating this procedure for all instances and all total budgets equal to the initial budget (original instance budget from PABULIB) or a feasible cost.
- For Figure 3 and Figure 4, we also compute the social welfare of the NE profiles found above. Since the greedy rule's outcome is fixed for a given total budget and all voters experience the same utility loss from donations, social welfare is independent of how donations are distributed and can be computed directly. The maximal social welfare in these figures is simply the maximum over all feasible subsets.
- For Figure 5, we compute donation values for all values of total budget for which NE exists and then find minimum of those. The procedure for calculating donation values is described below.
- For Figure 6 and Figure 7, we compute the global maximum social welfare (over all project subsets, not just feasible ones) for each instance. This is straightforward when the condition from Theorem 9 holds (as it does for all instances with $\gamma \geqslant 2^{-4}$). For instances that only slightly violate this condition (i.e., with a small remaining budget after funding all projects with utility exceeding cost), we use brute-force search for the best subset (for $\gamma \geqslant 2^{-8}$).

Calculating donation values. To calculate donation values in NE for a given instance and fixed total budget B, we intuitively aim to distribute donations as evenly as possible among voters. The procedure is as follows:

1. Order voters by their maximum possible donation in any NE at total budget B.
2. Take the smallest nonzero bound d_i from this list.
3. Ask all voters with nonzero bounds to donate d_i (or up to reaching B; then terminate).
4. Decrease the bound for all such voters by d_i and repeat from step 2 until the total budget is reached (since an NE exists, this process is guaranteed to terminate).

Maximal donations paid by a single voter in the resulting NE profile corresponds to the donation value for that specific instance under the total budget B.

Performance optimizations. Several optimizations were implemented to improve experimental performance: Frequently used values (e.g., voter utility functions) were precomputed for each instance and feasible subset. The implementation of ζ oracle, as described above, runs in time linear in the number of feasible subsets. Since voter utility functions are typically decreasing, and the maximal range R is often unbounded to the right, we implemented heuristics to quickly detect such cases. Where possible, we used numpy and matrix operations to accelerate calculations. We also leveraged parallelism throughout.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The abstract and the introduction contain a clear description of our results, techniques, and motivation for our research.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: Our manuscript includes a clear description of our modeling choices, as well as motivation for making them (see Sections 1 and 2). Furthermore, in Conclusion (Section 4) we discuss potential future research that could mitigate some of those limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [Yes]

   Justification: We provide sketches of the proofs for all of our theoretical claims in the main part of the paper. Due to space restrictions, full proofs are deferred to the technical appendix of our submission.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: We provide a robust description of the experimental setup that we used, which is entirely sufficient for replication of our empirical results (see Section 2.2 and Appendix B.2).

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [Yes]

   Justification: We include our code in the supplementary material of our submission. Regarding the data that we used in our study, we relied on the publicly available dataset PABULIB (see https://pabulib.org/).

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes]

   Justification: We provide a robust description of the experimental setup that we used, which is entirely sufficient for understanding our empirical results (see Section 2.2 and Appendix B.2).

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [Yes]

   Justification: All of our plots present clearly visible trends and most of them are in the form of scatter plots, which allow for a robust justification of our conclusions.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: We provide a robust description of the experimental setup that we used, which is entirely sufficient for replication of our empirical results (see Section 2.2 and Appendix B.2).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

   Answer: [Yes]

   Justification: Our research respects the NeurIPS Code of Ethics.

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes]

    Justification: Our manuscript clearly describes the practical applications and the societal impact that are directly related to our results.

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A]

    Justification: No such risks have been identified.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification: The code used in our research was written by the authors, while the data that we used is publicly available and has been credited appropriately.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [Yes]

    Justification: The code used in our research is well documented; see the supplementary material of the work.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification: Our research did not involve human subjects.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [N/A]

Justification: Our research did not involve human subjects.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

Answer: [N/A]

Justification: The core method development in this research does not involve LLMs as an important, original, or non-standard component.