Participatory Funding Coordination: Model, Axioms and Rules

Haris Aziz[1] and Aditya Ganguly[1]

UNSW Sydney, Australia
{haris.aziz,a.ganguly}@unsw.edu.au

Abstract. We present a new model of collective decision making that captures important crowd-funding and donor coordination scenarios. In the setting, there is a set of projects (each with its own cost) and a set of agents (that have their budgets as well as preferences over the projects). An outcome is a set of projects that are funded along with the specific contributions made by the agents. For the model, we identify meaningful axioms that capture concerns including fairness, efficiency, and participation incentives. We then propose desirable rules for the model and study, which sets of axioms can be satisfied simultaneously. An experimental study indicates the relative performance of different rules as well as the price of enforcing fairness axioms.

Keywords: Social choice · participatory budgeting · fairness · crowd-funding.

1 Introduction

Consider a scenario in which a group of residents want to pitch in money to buy some common items for the house but not every item is of interest or use to everyone. Each of the items (e.g. TV, video game console, music system, etc.) has its price. The residents each have a maximum amount they can spend towards the common items. Residents would like to have as much money as possible used toward items that are useful to them. It is a scenario that is encountered regularly in shared houses or apartments.

As a second scenario, hundreds of donors want to fund charitable projects. Each of the projects (e.g. building a well, enabling a surgery, funding a scholarship, etc.) has a cost requirement. Donors have upper caps on their individual budgets and care about the amount of money that is used towards projects of which they approve. The question of how to coordinate the funding in a principled and effective way is a fundamental problem in crowdfunding and donor coordination. The model that we propose is especially suitable for coordinating donations from alumni at various universities.

Both of the settings above are coordination problems in which agents contribute money, and they have preferences over the social outcomes. A collective outcome specifies which projects are funded and how much agents are charged. For these problems, we consider the following question. *What is a desirable and* *principled way of aggregating the preferences and financial contributions of the agents?*

Contributions. We propose a formal model that we refer to as *Participatory Funding Coordination (PFC)* that captures many important donor coordination scenarios. In this model, agents have an upper budget limit. The outcome for the problem is a set of projects that are funded and the respective monetary contributions of the agents for the funded projects. The utility of the agents is the amount of money used for projects that are approved by them. It reflects the approved investment from the perspective of an individual agent. We lay the groundwork for work on the model by formulating new axioms for the model. The logical relations between the axioms are established and the following question is studied: which sets of axioms are simultaneously achievable? We propose and study rules for the problems that are inspired by welfarist concerns but satisfy participation constraints. In addition to an axiomatic study of the rules, we also undertake an experimental comparison of the rules. The experiment sheds light on the impact that various fairness or participation constraints can have on the social welfare. This impact has been referred to as the price of fairness in other contexts. In particular, we investigate the effects of enforcing fairness properties on instances that model real-world applications of PFC, including crowdfunding.

2 Related Work

Our model generally falls under the umbrella of a collective decision making setting in which agents' donations and preferences are aggregated to make funding decisions. It is a concrete model within the broad agenda of achieving effective altruism [18, 19, 20].

The model we propose is related to the discrete participatory budgeting model [5, 4, 3, 14, 16, 21]. In discrete participatory budgeting, agents do not make personal donations towards the projects. They only express preferences over which projects should be funded. We present several axioms that are only meaningful for our model and not for discrete participatory budgeting. Algorithms for discrete participatory budgeting cannot directly be applied to our setting because they do not take into account individual rationality type requirements.

Another related setting is multi-winner voting [13]. Multi-winner voting can be viewed as a restricted version of discrete participatory budgeting. The Participatory Funding Coordination (PFC) setting differs from multi-winner voting in some key respects: in our model, each project (winner) has an associated cost, and we select projects subject to a knapsack constraint as opposed to having a fixed number of winners.

Our PFC model relies on approval ballots in order to elicit agents' preferences. Dichotomous preferences have been considered in several important setting including committee voting [2, 17] and discrete participatory budgeting [4, 15].

Another related model that takes into account the contributions of agents was studied by Brandl et al. [8]. Just like in our model, an agent's utilities are based on how much money is spent on projects approved by the agent. However, their model does not have any costs and agents can spread their money over projects in any way. Our model has significant differences from the model of [7, 8]: (1) in our setting, the projects are indivisible and have a minimum cost to complete; and (2) agents may not be charged the full amount of their budgets. The combination of these features leads to challenges in even defining simple individual rationality requirements. Furthermore, it creates difficulties in finding polynomial-time algorithms for some natural aggregation rules (utilitarian, egalitarian, Nash product, etc.). Our model is more appropriate for coordinating donations where projects have short-term deadlines and a target level of funding which must be reached for the project to be successfully completed. We show that the same welfarist rules that satisfy some desirable properties in the model [7, 8], fail to do so in our model. Just as the work of Brandl et al. [7, 8], Buterin et al. [9] consider donor coordination for the divisible model in which the projects do not have costs and agents do not have budget limits. They also assume quasi-linear utilities, whereas we model charitable donors who are not interested in profit but want their money being used as effectively as possible towards causes that matter to them.

The features of our PFC model enable the model to translate smoothly to a number of natural settings. Crowdfunding, in particular, is a scenario in which we would like to capitalise upon commonalities in donors' charitable preferences [11]. Furthermore, crowdfunding projects (e.g. building a well, funding a scholarship, etc.) often have provision points (see e.g. Agrawal et al. [1], Chandra et al. [10], Damle et al. [12]), and it can be critical for these targets to be met. For example, a project to raise funds for a crowdfunding recipient to pay for a medical procedure would have to raise a minimum amount of money to be successful, otherwise all donations are effectively wasted.

Crowdfunding projects have been discussed in a broader context with various economic factors and incentive issues presented [1]. Bagnoli and Lipman [6] discuss additional fairness and economic considerations for the related topic of the division of public goods. The discrete model that we explore, where projects have finite caps, has the potential to coordinate donors and increase the effectiveness of a crowdfunding system.

3 Participatory Funding Coordination

A *Participatory Funding Coordination (PFC)* setting is a tuple (N, C, A, b, w) where N is the set of agents/voters, C is the set of projects (also generally referred to as candidates). The function $w : C \to \mathbb{R}^+$ specifies the cost $w(c)$ of each project $c \in C$. The function $b : N \to \mathbb{R}_{\geq 0}$ specifies the budget b_i of each agent $i \in C$. The budget b_i can be viewed as the maximum amount of money that agent i is willing to spend. For any set of agents $M \subseteq N$, we will denote $\sum_{i \in M} b_i$ by $b(M)$. The approval profile $A = (A_1, \ldots, A_n)$ specifies for each agent, her set of acceptable projects A_i. An *outcome* is a pair (S, x) where $S \subseteq C$ is the set of funded projects and x is a vector of payments that specify for each $i \in N$, the payment x_i that is charged from agent i. We will restrict our attention to feasible outcomes in which $x_i \leq b_i$ for all $i \in N$ and only those projects get financial contributions that receive their required amount. Also, note that the projects that are funded are only those that receive the entirety of their price in payments from the agents. For any given PFC instance, a mechanism F returns an outcome. We will denote the set of projects selected by F as F_C and the payments by F_x. For any outcome (S, x), since $x_i \leq b_i$, the money $b_i - x_i$ can either be kept by the agent i or it can be viewed as going into some common pool. The main focus of our problem is to fund a maximal set of projects while satisfying participation constraints.

We suppose that an agent's preferences are *approval-based*. For any set of funded projects S, any agent i's utility is

$$u_i(S) = \sum_{c \in S \cap A_i} w(c).$$

That is, an agent cares about how many dollars are *usefully* used on his/her approved projects. Our preferences domain is similar to the one used by Brandl et al. [8] who considered a continuous model in which projects do not have target costs. In their model, agents also care about how much money is used for their liked projects.

4 Axiom design

In this section, we design axioms for outcomes of the PFC setting. We consider an outcome (S, x). For any axiom **Ax** for outcomes, we say that a mechanism satisfies **Ax** if it always returns an outcome that satisfies **Ax**.

We first present three axioms for our setting that are based on the principle of participation:

- **Minimal Return (MR)**: each agent's utility is at least as much as the money put in by the agent: $u_i(S) \geq x_i$. In other words, the societal decision is as good for each agent i as i's best use of the money x_i that she is asked to contribute. We will use this as a minimal condition for all feasible outcomes.
- **Implementability (IMP)** : There exists a payment function $y : N \times C \to \mathbb{R}_{\geq 0}$ such that $\sum_{c \in C} y(i, c) = x_i$ for all $i \in N$, $\sum_{i \in N} y(i, c) \in \{0, w(c)\}$ for all $c \in C$ and there exists no $i \in N$ and $c \notin A_i$ such that $y(i, c) > 0$. Here $y(i, c)$ represents the money paid by i to project c. IMP captures the requirement that an agent's contribution should only be used on projects that are approved by the agent.
- **Individual Rationality (IR):** the utility of an agent is at least as much as an agent can get by funding alone: $u_i(S) \geq \max_{S' \subseteq A_i, w(S') \leq b_i}(w(S'))$. Note that IR is easily achieved if the project costs are high enough: if for $i \in N$ and $c \in C$, $w(c) > b_i$, then every outcome is IR.

We note that MR is specified with respect to the amount x_i charged to the agent. It can be viewed as a participation property: an agent would only want to participate in the market if she gets at least as much utility as the money she spent. We will show IMP is stronger than MR. IMP can also be viewed as a fairness property: agents are made to coordinate but they only spend their money on the projects they like.

Remark 1. If (S, x) is an IMP outcome with associated payment function y, then for any subset of projects $S' \subseteq S$, there is an IMP outcome that funds only the set of projects S'. In particular, the payment function y' for one such implementable outcome is obtained by setting (for each agent i) $y'(i, c) = y(i, c)$ for all $c \in S'$ and $y'(i, c) = 0$ for all $c \in S \setminus S'$.

Next, we present axioms that are based on the idea of efficiency.

- **Exhaustive (EXH)**: An outcome (S, x) satisfies EXH if there exists no set of agents $N' \subseteq N$ and unfunded project $c \in C \setminus S$ such that $c \in \cap_{i \in N'} A_i$ with $w(c) \leq \sum_{i \in N'}(b_i - x_i)$. In words, agents in N' cannot pool in their unspent money and fund another project liked by all of them.
- **Pareto optimality (PO)-X**: An outcome (S, x) is Pareto optimal within the set of outcomes satisfying property X if there exists no outcome (S', x') satisfying X such that $u_i(S') \geq u_i(S)$ for all $i \in N$ and $u_i(S') > u_i(S)$ for some $i \in N$. Note that Pareto optimality is a property of the set of funded projects S irrespective of the payments.
  - PO is Pareto optimal among the set of all outcomes.
  - PO-IMP: Pareto optimal among the set of IMP outcomes.
  - PO-MR: Pareto optimal among the set of MR outcomes.
- **Payment constrained Pareto optimality (PO-Pay)**: An outcome is PO-Pay if it is not Pareto dominated by any outcome of at most the same price. Formally, there exists no (S', x') such that $\sum_{i \in N} x'_i \leq \sum_{i \in N} x_i$, $u_i(S') \geq u_i(S)$ for all $i \in N$ and $u_i(S') > u_i(S)$ for some $i \in N$.
- **Weak Payment constrained Pareto optimality (weak PO-Pay)**: An outcome is weakly PO-Pay if it is not Pareto dominated by any outcome that charges at most the same cost from each agent. Formally, there exists no (S', x') such that $x'_i \leq x_i$ and $u_i(S') \geq u_i(S)$ for all $i \in N$ and $u_i(S') > u_i(S)$ for some $i \in N$.

A concept that can be viewed in terms of participation, efficiency, and fairness is the adaptation of the principle of core stability for our setting.

- **Core stability (CORE)**: There exists no set of agents who can pool in their budget and each gets a strictly better outcome. In other words, an outcome (S, x) is CORE if for every subset of agents $N' \subseteq N$, for every subset of projects $C' \subseteq C$ such that $w(C') \leq \sum_{i \in N'} b_i$, the following holds for some agent $i \in N'$: $u_i(S) \geq w(C' \cap A_i)$.

We also describe a basic fairness axiom for outcomes and rules based on the idea of proportionality.

- **Proportionality (PROP)**: Suppose a set of agents $N' \subseteq N$ each have approval set that is *exactly* some (common) set of projects $C' \subseteq C$ such that $\sum_{i \in N'} b_i \geq w(C')$. In that case, all the projects in C' are selected.

Finally, we consider an axiom that is defined for mechanisms rather than outcomes. We say that a mechanism satisfies **strategyproofness** if there exists no instance under which some agent has an incentive to misreport her preference relation.

We conclude this section with some remarks on computation. The following proposition follows via a reduction from the Subset Sum problem.

Proposition 1. *Even for one agent, computing an IR, PO, PO-MR, or PO-IMP outcome is NP-hard.*

Note that IMP is a property of an outcome rather than a set of projects. We say that a set of projects S is IMP if there exists a feasible vector of charges to agents x such that the outcome (S, x) is IMP. The property IMP can be tested in polynomial time via reduction to network flows.

Proposition 2. *For a given set of projects S, checking whether there exists a vector of charges x such that (S, x) is implementable can be done in polynomial time.*

5 Axioms: Compatibility and Logical Relations

In this section, we study the compatibility and relations between the axioms formulated.

Remark 2. Note that IR and MR are incomparable. Any outcome in which every agent is not charged any money trivially satisfies MR. However, it will not satisfy IR if any agent could afford one of their approved projects by themselves. On the other hand, an IR outcome may not be MR. Consider a profile with one agent and one project. Say the agent has budget greater than the cost of the project, but does not approve of the project. Then, the outcome where the agent is forced to fund the project is IR but not MR.

Next, we point out that that PO-Pay is equivalent to weak PO-Pay.

Proposition 3. *PO-Pay is equivalent to weak PO-Pay.*

Proof. Suppose an outcome (S, x) is not weakly PO-Pay. Then, it is trivially not PO-Pay. Now suppose (S, x) is not PO-Pay. Then, there exists another outcome (S', x') such that $\sum_{i \in N} x'_i \leq \sum_{i \in N} x_i$ and $u_i(S') \geq u_i(S)$ for all $i \in N$ and $u_i(S') > u_i(S)$ for some $i \in N$. Note that S' can be funded with total amount $\sum_{i \in N} x'_i$ irrespective of who paid what. So S' is still affordable if $x'_i \leq x_i$ for each agent i. □

The next proposition establishes further logical relations between the axioms.

Proposition 4. *The following logical relations hold between the properties.*

1. *IMP implies MR.*
2. *PO implies PO-Pay.*
3. *PO-X implies PO-Y if Y implies X.*
4. *PO-IMP implies EXH.*
5. *PO-IR implies EXH.*
6. *CORE implies IR.*
7. *The combination of PO-IMP and IMP imply PROP.*

Next, we show that MR is compatible with PO-Pay.

Proposition 5. *Suppose an outcome is MR and there is no other MR outcome that Pareto dominates it. Then, it is PO-Pay.*

Proof. Suppose the outcome (S, x) is MR and PO-MR. We claim that (S, x) is PO-Pay. Suppose it is not PO-Pay. Then there exists another outcome (S', x') such that $\sum_{i\in N} x'_i \leq \sum_{i\in N} x_i$, $u_i(S') \geq u_i(S)$ for all $i \in N$ and $u_i(S') > u_i(S)$ for some $i \in N$. Note that S' is affordable with total amount $\sum_{i\in N} x'_i$ irrespective of who paid what. So S' is still affordable if $x'_i \leq x_i$. Therefore, we can assume that $x'_i \leq x_i$ for all $i \in N$. Note that since S' Pareto dominates S and since (S, x) is MR, $u_i(S') \geq u_i(S) \geq x_i \geq x'_i$ for all $i \in N$. Hence (S', x') also satisfies MR. Since (S', x') is MR and since S' Pareto dominates S, it contradicts the fact that (S, x) is PO-MR. □

Proposition 6. *There always exists an outcome that satisfies IMP, IR, PO-IMP and hence also MR and EXH.*

Proof. For each $i \in N$ compute $S_i = \arg\max_{S'\subseteq A_i, w(S')\leq b_i} w(S')$, i.e. a maximum total weight set of approved projects that has weight at most b_i. Then, observe that any outcome that funds all of $S = \bigcup_{i\in N} S_i$ is necessarily IR. Thus, in order to construct an IMP and IR outcome, we can construct a payment function $y : N \times C \to \mathbb{R}_{\geq 0}$ that funds S. For each project $c \in C$, let n_c denote the number of agents i with $c \in S_i$. Note that if $c \in S$, then $n_c \geq 1$. Then, for each $i \in N$ let $y(i, c) = \frac{w(c)}{n_c}$ if $c \in S_i \subseteq S$ and $y(i, c) = 0$ if $c \notin S_i$. It is then simple to check that each agent's total cost is affordable to them, each project in S is fully paid for, and each agent i only pays for projects in $S_i \subseteq A_i$ (i.e. projects that they approve of). Therefore, we can let $x = (\sum_{c\in C} y(1, c), \ldots, \sum_{c\in C} y(n, c))$ and see that the outcome (S, x) is IR and IMP.

Now, observe that if any IMP outcome (S', x') Pareto dominates (S, x), then it must be also be IR because the utility of each agent is at least as high as before. There can only be a finite number of Pareto improvements to (S, x) since the number of possible subsets of projects to be funded is finite, and Pareto dominance depends only on the projects funded, not the costs to the agents. Hence, there must exist such a Pareto improvement which is IR, IMP and PO-IMP. Finally, Proposition 4 gives that this outcome must be MR and EXH. □

Note that PO-Pay and IMP are both satisfied by an empty outcome with zero charges. PO-IMP and IMP are easily satisfied by computing a PO outcome from the set of IMP outcomes. PO-Pay and PO-IMP are easily satisfied by computing a PO outcome which may not necessarily satisfy IMP.

Proposition 7. *There always exists an outcome that satisfies MR, IR, PO-MR and hence also EXH.*

Proof. Existence of an outcome that satisfies MR, IR, PO-MR: From the proof of Proposition 6, we know that an IMP and IR outcome always exists. Also, from Proposition 4, we know that every IMP outcome is MR, so there always exists an MR and IR outcome. Now suppose the outcome satisfying MR and IR does not satisfy PO-MR. Then there exists another outcome satisfying MR that Pareto dominates the original outcome, which is still IR. There cannot exist an infinite number of Pareto improvements because there are only finitely many possible subsets of projects that can be funded. Hence we can reach a PO-MR outcome that is also IR and MR. Proposition 4 gives that this outcome is EXH. □

We note that if no agent can individually fund a project, then every outcome is IR. In crowdfunding settings in which projects have high costs, the IR requirement is often easily satisfied.

6 Aggregation Rules

In this section, we take a direct welfarist view to formalize rules that maximize some notion of welfare. We consider three notions of welfare: utilitarian, egalitarian, and Nash welfare; and we define the following rules.

- UTIL: define the utilitarian welfare derived from an outcome (S, x) as $\sum_{i\in N} u_i(S)$. Then, UTIL returns an outcome that maximises the utilitarian welfare.
- EGAL: given some outcome (S, x), write the sequence of agents' utilities from that outcome as a tuple $u(S) = (u_i(S))_{i\in N}$, where u is sorted in non-decreasing order. Then, EGAL returns an outcome (S, x) such that $u(S)$ is lexicographically maximal among the outcomes.
- NASH: maximises the Nash welfare derived from an output (S, x), i.e. $\prod_{i\in N} u_i(S)$.

Proposition 8. *UTIL, EGAL, and NASH satisfy PO and hence PO-MR, PO-IMP, PO-Pay, and EXH.*

One notes that the rules UTIL, EGAL, and NASH do not satisfy minimal guarantees such as MR. The reason is that an agent may donate her budget to a widely approved project even though she may not approve any of such projects.

Given that the existing aggregation rules do not provide us with guarantees that the outcomes they produce will satisfy our axioms, we can instead define rules that optimize social welfare within certain subsets of feasible outcomes. For a property X, we can define UTIL-X, EGAL-X, and NASH-X as rules that maximise the utilitarian, egalitarian and Nash welfare respectively among only those outcomes that satisfy property X.

Next, we analyse the properties satisfied by rules EGAL/UTIL/NASH constrained to the set of MR or IMP outcomes. In the continuous model introduced by Brandl et al. [8], there is no need to consider the rule NASH-IMP, as the NASH rule in the case where projects can be funded to an arbitrary degree (given there is sufficient budget) already satisfies IMP.

Before we study the axiomatic properties, we note that most meaningful axioms and rules are NP-hard to achieve or compute. The following result follows from Proposition 1.

Proposition 9. *Even for one agent, computing an UTIL, UTIL-MR, UTIL-MR, EGAL, EGAL-MR, EGAL-IMP, NASH, NASH-MR, NASH-IMP outcome is NP-hard.*

Similarly, the next result follows from Proposition 5.

Proposition 10. *UTIL-MR, EGAL-MR, and NASH-MR satisfy PO-Pay.*

From Proposition 5, it follows that UTIL-MR, EGAL-MR, and NASH-MR satisfy PO-Pay. In contrast, we show that UTIL-IMP, EGAL-IMP, and NASH-IMP do not satisfy PO-Pay. In order to show this, we prove that it is possible in some instances for the set of jointly IMP and PO-IMP outcomes to be disjoint from the set of PO-Pay outcomes.

Proposition 11. *UTIL-IMP, EGAL-IMP and NASH-IMP do not satisfy PO-Pay. In fact it is possible that no IMP and PO-IMP outcome satisfies PO-Pay.*

Similarly, the following also holds.

Proposition 12. *EGAL, EGAL-MR and EGAL-IMP are not strategyproof.*

Table 1 shows the axioms that are satisfied by restricting the aggregation rules to optimising within the space of MR or IMP outcomes.

7 Experiment

In addition to the axiomatic study of the welfare-based rules, we undertake a simulation-based experiment to gauge the performance of different rules with respect to utilitarian and egalitarian welfare. Our study shows the impact of fairness axioms such as MR and IMP on welfare.

We generate random samples of profiles in order to simulate two potential real-world applications of PFC.

	UTIL-MR	EGAL-MR	NASH-MR	UTIL-IMP	EGAL-IMP	NASH-IMP
MR	✓	✓	✓	✓	✓	✓
IMP	–	–	–	✓	✓	✓
PROP	–	–	–	✓	✓	✓
IR	–	–	–	–	–	–
PO	–	–	–	–	–	–
PO-MR	✓	✓	✓	–	–	–
PO-IMP	✓	✓	✓	✓	✓	✓
PO-Pay	✓	✓	✓	–	–	–
EXH	✓	✓	✓	✓	✓	✓
CORE	–	–	–	–	–	–
SP	–	–	–	–	–	–

Table 1: Properties satisfied by UTIL-MR, EGAL-MR, NASH-MR, UTIL-IMP, EGAL-IMP and NASH-IMP.

1. Share-house setting: In this example, we can imagine a group of house-mates pooling their resources to fund communal items for their house. We operate under the following assumptions:
   - Number of agents from 3-6: this represents a reasonable number of house-mates in a share-house.
   - Number of projects from 5-12: projects may include buying items such as tables, chairs, sofas, televisions, lights, kitchen appliances, washing machines, dryers, etc.
   - Agent budgets are from 300-600 and project costs are from 50-1000. We base these costs on typical rent and furniture costs in Australia as well as costs of the above items in first and second-hand retailers. We expect that each agent brings some money to the communal budget, and would spend around one or two weeks' worth of rent on one-time communal expenses.
2. Crowdfunding setting: In this example, we imagine a relatively small number of expensive projects to be funded, and a large number of philanthropic donors, and make the following assumptions.
   - Number of agents from 20-50: A review of crowdfunding websites such as Kickstarter and GoFundMe shows that the most promoted projects are typically funded by thousands of donors, and smaller projects can attract tens of donors. For the purposes of our simulation, we use between 20-50 donors, which is still relatively large compared to the number of available projects.
   - Number of projects from 3-8: In crowdfunding, there are far more projects available than a donor actually sees. However, we can estimate that in a browsing session, a donor might view the top 3-8 promoted projects.
   - Agent budgets from 0-400 and project costs from 1000-10000: Projects in real-life crowdfunding can have vastly varying costs. For our simulation, we want for the agents with all their money combined to be able to afford

some, but not all of the available projects in order to create instances that are not trivially resolved by funding all or none of the projects.

The results of the experiments are shown in Figures 1, 2, 3, 4, 5, 6, 7 and 8.

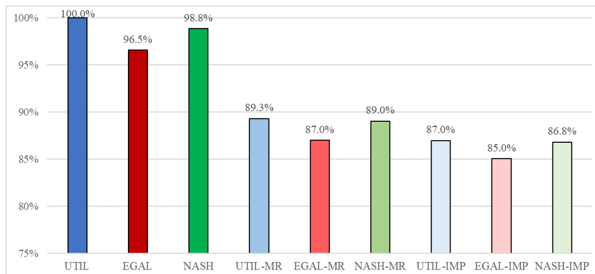

Fig. 1: Average performance of rules with respect to utilitarian welfare in share-house simulations as a percentage of the maximum achievable utilitarian welfare.

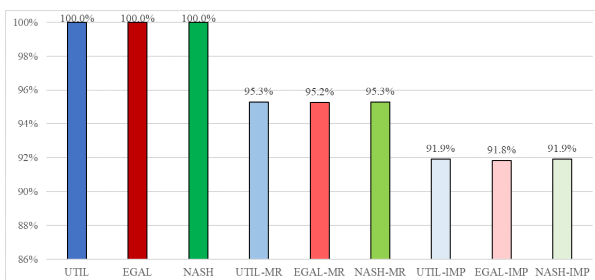

Fig. 2: Average performance of rules with respect to utilitarian welfare in crowdfunding simulations as a percentage of the maximum achievable utilitarian welfare.

Imposing MR on a rule seems to have a significant impact on both utilitarian and egalitarian welfare on average. Of course, since IMP implies MR, we expect that imposing IMP as a constraint will have an even greater cost on welfare, but from our experiment, this cost is a relatively small increase on top of the cost of imposing MR. It is worth noting that in worst-case scenarios, it is always possible that there are no non-trivial outcomes that satisfy the constraints, and so there is a risk that a rule subject to a constraint could produce an outcome that gives all agents zero utility.

When considering average performance, rules are more resilient to the imposition of fairness constraints for instances that simulate crowdfunding scenarios compared to share-house scenarios. When the number of agents is high and

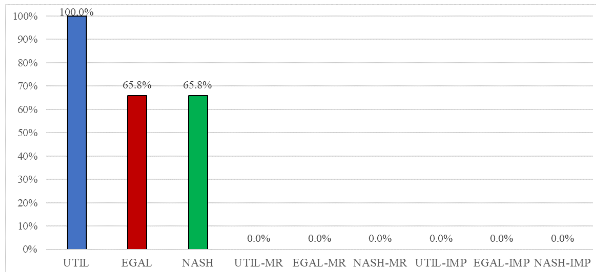

Fig. 3: Worst-case performance of rules with respect to utilitarian welfare in share-house simulations as a percentage of the maximum achievable utilitarian welfare.

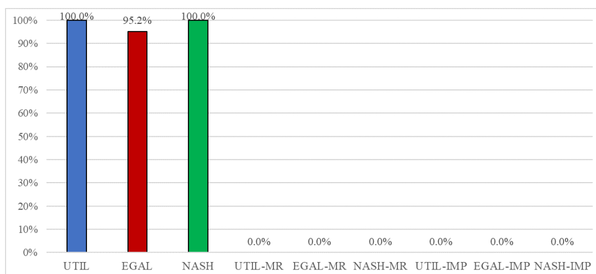

Fig. 4: Worst-case performance of rules with respect to utilitarian welfare in crowdfunding simulations as a percentage of the maximum achievable utilitarian welfare.

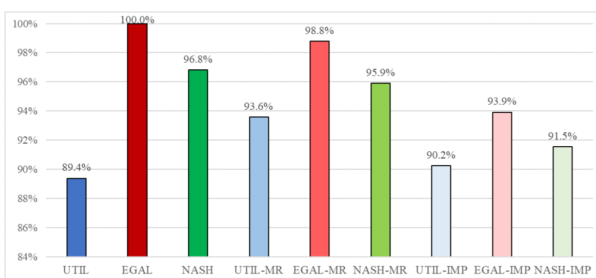

Fig. 5: Average performance of rules with respect to egalitarian welfare in share-house simulations as a percentage of the maximum achievable egalitarian welfare.

the number of projects is small, and project costs are high compared to agent budgets, it seems to be easier to achieve fairness properties.

We typically expect the NASH rule to be a compromise between UTIL and EGAL. This manifests in the results, where the performance losses for NASH

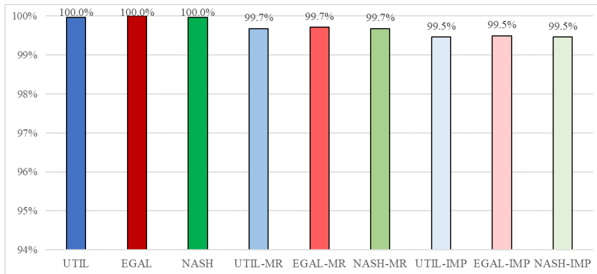

Fig. 6: Average performance of rules with respect to egalitarian welfare in crowdfunding simulations as a percentage of the maximum achievable egalitarian welfare.

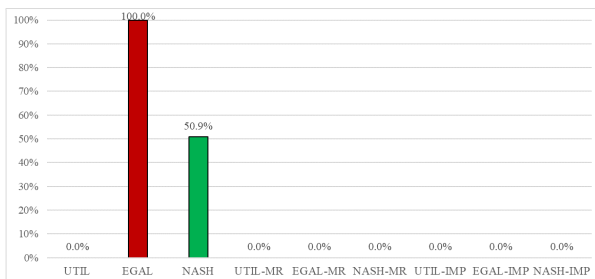

Fig. 7: Worst-case performance of rules with respect to egalitarian welfare in share-house simulations as a percentage of the maximum achievable egalitarian welfare.

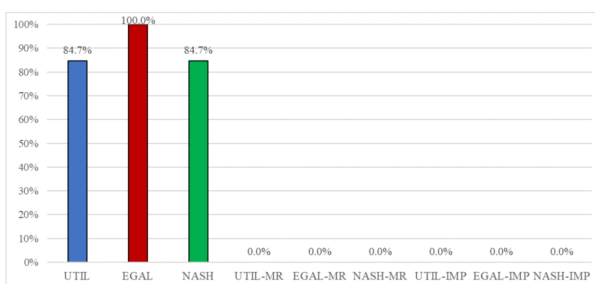

Fig. 8: Worst-case performance of rules with respect to egalitarian welfare in crowdfunding simulations as a percentage of the maximum achievable egalitarian welfare.

with respect to utilitarian welfare are considerably less than those for EGAL. NASH loses considerably less with respect to egalitarian welfare than UTIL.

8 Conclusions

We proposed a concrete model for coordinating funding for projects. A formal approach is important to understand the fairness, participation, and efficiency requirements a system designer may pursue. We present a detailed taxonomy of such requirements and clarify their properties and relations. We also analyse natural welfarist rules both axiomatically and experimentally.

In practical applications of PFC, it is important to balance welfare demands with fairness conditions. Our experiment investigated the cost of fairness when imposing MR or IMP on UTIL, EGAL and NASH rules over instances that model crowdfunding and share-house scenarios. We find that imposing MR alone significantly reduces welfare on average, but imposing IMP as well produces a relatively small additional cost on welfare. The costs of imposing any fairness condition are much more pronounced on instances that model a share-house setting than a crowdfunding setting, suggesting that for a large number of agents and large project costs, fairness conditions are more easily met.

Our model is not just a rich setting to study collective decision making. We feel that the approaches considered in the paper go beyond academic study and can be incorporated in portals that aggregate funding for charitable projects. We envisage future work on online versions of the problem. We studied a utility model in which agents want as much money spent on their approved projects. It will be interesting to examine utility models in which agents care about which unapproved projects are funded or factor in the payments they have been made.

Acknowledgements

The authors thank Barton Lee and the anonymous reviewers of ADT 2021 for their helpful comments. They also thank the UNSW Taste of Research program under which this research was conducted.

References

1 Agrawal, A., Catalini, C., Goldfarb, A.: Some Simple Economics of Crowdfunding. In: Innovation Policy and the Economy, Volume 14, pp. 63–97, NBER Chapters, National Bureau of Economic Research, Inc (June 2013)
2 Aziz, H., Brill, M., Conitzer, V., Elkind, E., Freeman, R., Walsh, T.: Justified representation in approval-based committee voting. Social Choice and Welfare pp. 461–485 (2017)
3 Aziz, H., Lee, B.E.: Proportionally representative participatory budgeting with ordinal preferences. In: Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI) (2021)
4 Aziz, H., Lee, B.E., Talmon, N.: Proportionally representative participatory budgeting: Axioms and algorithms. In: Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS 2018, Stockholm, Sweden, July 10-15, 2018, pp. 23–31 (2018)

5 Aziz, H., Shah, N.: Participatory budgeting: Models and approaches. In: Rudas, T., Gábor, P. (eds.) Pathways between Social Science and Computational Social Science: Theories, Methods and Interpretations, Springer (2020)
6 Bagnoli, M., Lipman, B.L.: Provision of Public Goods: Fully Implementing the Core through Private Contributions. The Review of Economic Studies **56**(4), 583–601 (10 1989)
7 Brandl, F., Brandt, F., Peters, D., Stricker, C., Suksompong, W.: Donor coordination: Collective distribution of individual contributions. In: GAIW (Games, Agents and Incentives Workshops) (2019)
8 Brandl, F., Brandt, F., Peters, D., Stricker, C., Suksompong, W.: Funding public projects: A case for the nash product rule. Working paper (2020)
9 Buterin, V., Hitzig, Z., Weyl, E.G.: A Flexible Design for Funding Public Goods. Management Science **65**(11), 5171–5187 (2019)
10 Chandra, P., Gujar, S., Narahari, Y.: Crowdfunding public projects with provision point: A prediction market approach. In: Proceedings of the 22nd European Conference on Artificial Intelligence (ECAI), pp. 778–786 (2016)
11 Corazzini, L., Cotton, C., Valbonesi, P.: Donor coordination in project funding: Evidence from a threshold public goods experiment. Journal of Public Economics **128**(1), 16–29 (2015)
12 Damle, S., Moti, M.H., Chandra, P., Gujar, S.: Civic crowdfunding for agents with negative valuations and agents with asymmetric beliefs. In: Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI), pp. 208–214 (2019)
13 Elkind, E., Faliszewski, P., Skowron, P., Slinko, A.: Properties of multiwinner voting rules. Social Choice and Welfare (2017)
14 Fain, B., Goel, A., Munagala, K.: The core of the participatory budgeting problem. In: Proceedings of the 12th International Conference on Web and Internet Economics (WINE '16), pp. 384–399 (2016)
15 Fluschnik, T., Skowron, P., Triphaus, M., Wilker, K.: Fair knapsack. In: Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI) (2019)
16 Goel, A., K., A.K., Sakshuwong, S., Aitamurto, T.: Knapsack voting for participatory budgeting. ACM Transactions on Economics and Computation (TEAC) **7**(2), 8:1–8:27 (2019), ISSN 2167-8375
17 Lackner, M., Skowron, P.: A quantitative analysis of multi-winner rules. In: Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI), pp. 407–413 (2019)
18 MacAskill, W.: Doing Good Better: How Effective Altruism Can Help You Make a Difference. Avery (2015)
19 MacAskill, W.: Effective altruism: Introduction. Essays in Philosophy **18**(1) (2017)
20 Peters, D.: Economic Design for Effective Altruism, pp. 381–388 (2019), URL https://app.dimensions.ai/details/publication/pub.1122621382
21 Talmon, N., Faliszewski, P.: A framework for approval-based budgeting methods. In: Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI) (2019)