Shortlisting Rules and Incentives in an End-to-End Model for Participatory Budgeting

Simon Rey, Ulle Endriss and Ronald de Haan

October 20, 2020

Abstract

We introduce an end-to-end model of participatory budgeting grounded in social choice theory. This model accounts for both the first stage, in which participants propose projects to be shortlisted, and the second stage, in which they vote on which of the shortlisted projects should be funded. We introduce several shortlisting rules for the first stage and we analyse them in both normative and algorithmic terms. Our main focus is on the incentives of participants to engage in strategic behaviour, especially in the first stage, in which they need to reason about how their proposals will impact the range of strategies available to everyone in the second stage.

1 Introduction

Participatory budgeting (PB) is a loosely defined range of mechanisms designed to improve the involvement of ordinary citizens in public spending decisions [Cabannes, 2004, Shah, 2007]. In a first stage, participants are invited to propose projects, a selection of which are put on a shortlist. Then, in a second stage, everyone can vote on the shortlisted projects to decide which of them should receive funding. These are problems of social choice with a clear algorithmic component [Brandt et al., 2016, Aziz and Shah, 2020]. Prior formal work—in particular in the AI and the Economics & Computation communities—has concentrated almost exclusively on the second stage [Aziz et al., 2018, Benade et al., 2017, Fluschnik et al., 2019, Fain et al., 2016, Goel et al., 2019, Jain et al., 2020, Rey et al., 2020, Talmon and Faliszewski, 2019]. In this paper, we instead propose an *end-to-end model* of PB that accounts for both stages.

Our contribution, beyond the formulation of the model itself, is threefold. First, we propose and analyse several shortlisting rules for the first stage. Second, we analyse the incentives of engaging in strategic manipulation when making proposals during the first stage. Third, we clarify a number of issues regarding similar incentives during the second stage. Let us briefly discuss each of these in turn.

During the first stage, the *shortlisting stage*, participants can propose projects (e.g., planting a tree). In practice, proposals often are reviewed by experts to exclude projects that are technically infeasible. Beyond that, it is also often seen as desirable to significantly reduce the number of proposals entering the second stage.[1] An important objective at this point is *diversity*. Inspired by the Thiele rules for committee elections [Janson, 2016], we study diversity w.r.t. proposers, by minimising the number of participants without shortlisted proposals. Secondly, inspired by clustering algorithms [Jain and Dubes, 1988], we also explore diversity w.r.t. the shortlist, by avoiding selecting too many similar projects.

During the second stage, the *allocation stage*, participants vote on the shortlisted projects to decide which of them should be funded—under a given budget limit. Since the shortlisting stage determines the set of projects for the allocation stage, there is significant interaction between the two stages, which motivates the study of end-to-end models that can account for such effects. Our focus is on participants who strategise during the shortlisting stage to affect the set of shortlisted projects that people can vote for during the allocation stage. We introduce the notion of *first-stage strategyproofness* and analyse how it depends on both the information available to participants and the choice of aggregation rules used during the two stages.

[1] E.g., for the 2008 to 2012 PB exercises in Lisbon, around 30% of the projects were shortlisted [Allegretti and Antunes, 2014].

Finally, we complement known results regarding participants' incentives to strategise during the second stage. In particular, we observe that the approximate strategyproofness result of Goel et al. [2019] for the widely used *greedy-approval rule* does not extend to its 'optimised' variant, the *approval-maximising rule* [Talmon and Faliszewski, 2019]. We also observe that the possibility of achieving strategyproofness can depend heavily on the precise manner in which ties between equally good outcomes are broken.

Related work. Most prior research of a formal nature regarding normative concerns in PB has focused on the allocation stage only. For example, Goel et al. [2019] discuss strategic behaviour, Aziz et al. [2018] discuss proportionality, and Jain et al. [2020] enrich the basic framework to be able to account for complementarity and substitutability of projects. Much of this work takes inspiration from the literature on multiwinner voting [Faliszewski et al., 2017], exploiting the fact that electing a committee of k representatives is isomorphic to selecting projects for funding when each project costs 1 dollar and the budget limit is k dollars. In this paper, we discuss allocation rules proposed by Goel et al. [2019] and Talmon and Faliszewski [2019], and we complement the analysis of incentives to strategise given by the first group of authors.

We are not aware of any formal work regarding the shortlisting stage in PB, but there are still links to multiwinner voting, where the term 'shortlisting' is used in two different senses: either to emphasise that exactly k candidates are to be elected [Faliszewski et al., 2017], or to refer to the problem of electing a set of variable size [Kilgour, 2016, Faliszewski et al., 2020, Lackner and Maly, 2020]. Only the latter problem is related to shortlisting for PB, and also this problem does not involve costs or a budget limit, as is the case for PB. We note that Lackner and Maly [2020] also discuss connections to clustering algorithms.

Paper outline. We develop our end-to-end model for PB in Section 2. Then, Section 3 is dedicated to shortlisting rules, Section 4 to first-stage strategyproofness, and Section 5 to strategic issues arising during the second stage.

2 The Model

In this section, we introduce the two stages of our model for PB and fix our assumptions regarding agent preferences.

2.1 Basic Notation and Terminology

Let $\mathbb{P} = \{p_1, \ldots, p_m\}$ be the (finite) set of all conceivable projects. The cost of each project is given by $c : \mathbb{P} \to \mathbb{N}$. The total cost of any set $P \subseteq \mathbb{P}$ is written $c(P) = \sum_{p \in P} c(p)$. The budget limit is denoted by $B \in \mathbb{N}$. For every project $p \in \mathbb{P}$, we assume w.l.o.g. that $c(p) \leq B$. The set of agents participating in the PB exercise is denoted by $\mathcal{N} = \{1, \ldots, n\}$.

We shall make use of the following generic procedure to choose a "best" subset (fitting the budget) of a given set of projects in view of a given ranking of those projects.

Definition 1 (Greedy selection)**.** *For a set $P \subseteq \mathbb{P}$ of projects and a strict linear order $\gg$ on P, the greedy selection procedure returns the set of projects $GREED(P, \gg)$ defined iteratively as follows. Projects are examined following $\gg$. A project $p \in P$ is selected if and only if doing so keeps the total cost of the selected projects within the budget limit B. The next project, if any, is then considered.*

We are going to require a means for breaking ties, both between alternative projects and between alternative sets of projects. For any $P \subseteq \mathbb{P}$, let $idx(P) = \{i \in \mathbb{N} \mid p_i \in P\}$ be the set of indices of the projects in P. The *canonical tie-breaking rule* T returns $T(P) = p_i$ with $i = \min(idx(P))$ for any nonempty set $P \subseteq \mathbb{P}$. We also use T to transform weak orders on projects into strict orders. Take any weak order $\geq$ on $\mathbb{P}$. Then for every indifference class $P \subseteq \mathbb{P}$ of $\geq$, we break ties as follows: $p = T(P)$ is the first project, then comes $T(P \setminus \{p\})$, and so forth. Overloading notation, we denote by $T(\geq)$ the strict order thus obtained. Finally, we extend T to nonempty sets $\mathfrak{P} \subseteq 2^{\mathbb{P}}$ of sets of projects in a lexicographic manner: $T(\mathfrak{P})$ is the unique set $P \in \mathfrak{P}$ such that $T((P \setminus P') \cup (P' \setminus P)) \in P$ for all $P' \in \mathfrak{P} \setminus \{P\}$. Thus, we require that, amongst all the projects on which P and P' differ, the one with the lowest index must belong to P.

2.2 The Shortlisting Stage

In the first stage, agents are asked to propose projects. An *shortlisting instance* is a tuple $\langle \mathbb{P}, c, B \rangle$. Because of bounded rationality, an agent may not be able to conceive of all the projects she would approve of if only she were aware of them. We denote by $C_i \subseteq \mathbb{P}$ the set of projects that agent i can conceive of (her *awareness set*), and we call the vector $\boldsymbol{C} = (C_1, \dots, C_n)$ the *awareness profile*. Agent i knows the cost of the projects in C_i as well as the budget limit B.

We denote by $P_i \subseteq C_i$ the set of projects agent $i \in \mathcal{N}$ chooses to actually propose, and we call the resulting vector $\boldsymbol{P} = (P_1, \dots, P_n)$ a *shortlisting profile*. We use $(\boldsymbol{P}_{-i}, P'_i)$ to denote the profile we obtain when, starting from profile $\boldsymbol{P}$, agent i changes her proposal to $P'_i \subseteq C_i$.

A *shortlisting rule* R maps any given shortlisting instance $I = \langle \mathbb{P}, c, B \rangle$ and shortlisting profile $\boldsymbol{P}$ to a shortlist, i.e., a set $R(I, \boldsymbol{P}) \subseteq \bigcup \boldsymbol{P} = P_1 \cup \dots \cup P_n$ of shortlisted projects.

2.3 The Allocation Stage

In the second stage, agents vote on the shortlisted projects to decide which projects should get funded. An *allocation instance* is a tuple $\langle \mathcal{P}, c, B \rangle$, where $\mathcal{P} \subseteq \mathbb{P}$ is the set of shortlisted projects. Contrary to the shortlisting stage, the agents now know about all the projects they can vote for. They vote by submitting approval ballots, denoted by $A_i \subseteq \mathcal{P}$ for each agent $i \in \mathcal{N}$, giving rise to a *profile* $\boldsymbol{A} = (A_1, \dots, A_n)$. The *approval score* of a project p in profile $\boldsymbol{A}$ is defined as $n_p^{\boldsymbol{A}} = |\{i \in \mathcal{N} \mid p \in A_i\}|$. Moreover, we define the weak order $\geq_{app}^{\boldsymbol{A}}$ on $\mathbb{P}$ by stipulating that $p \geq_{app}^{\boldsymbol{A}} p'$ holds if and only if $n_p^{\boldsymbol{A}} \geq n_{p'}^{\boldsymbol{A}}$. Finally $(\boldsymbol{A}_{-i}, A'_i)$ is the profile obtained from $\boldsymbol{A}$ when agent $i \in \mathcal{N}$ changes her ballot to $A'_i \subseteq \mathcal{P}$.

The goal of this stage is to choose a *budget allocation* $A \subseteq \mathcal{P}$. Such an allocation is *feasible* if $c(A) \leq B$. It is *exhaustive* if it is feasible and there exists no project $p \in \mathcal{P} \setminus A$ with $c(A \cup \{p\}) \leq B$. For a given instance I, we denote the the set of feasible budget allocations and exhaustive budget allocations by $\mathcal{A}(I)$ and $\mathcal{A}_{EX}(I)$, respectively.

An *allocation rule* F maps any given allocation instance I and profile $\boldsymbol{A}$ to a feasible budget allocation $F(I, \boldsymbol{A}) \in \mathcal{A}(I)$. Two well-known allocation rules try to maximise the approval score of their output, either greedily [Goel et al., 2019] or exactly [Talmon and Faliszewski, 2019]. Note that both can be computed in polynomial time.

Definition 2 (Greedy-approval rule)**.** *The greedy-approval rule F returns, for any given allocation instance $I = \langle \mathcal{P}, c, B \rangle$ and profile $\boldsymbol{A}$, the following budget allocation:*

$$F(I, \boldsymbol{A}) = \mathit{GREED}\left(\mathcal{P}, T\left(\geq_{app}^{\boldsymbol{A}}\right)\right)$$

Definition 3 (Approval-maximising rule)**.** *The approval-maximising rule F returns, for any given allocation instance $I = \langle \mathcal{P}, c, B \rangle$ and profile $\boldsymbol{A}$, the following budget allocation:*

$$F(I, \boldsymbol{A}) = T\left(\operatorname*{argmax}_{A \in \mathcal{A}(I)} \sum_{p \in A} n_p^{\boldsymbol{A}}\right)$$

We say that an allocation rule F is *exhaustive* if, for all instances I and all profiles $\boldsymbol{A}$, we have $F(I, \boldsymbol{A}) \in \mathcal{A}_{EX}(I)$. We furthermore say that F is *unanimous* if, for every instance $I = \langle \mathcal{P}, c, B \rangle$ and every profile of the form $\boldsymbol{A} = (A, \dots, A)$ with $A \in \mathcal{A}(I)$, we have $F(I, \boldsymbol{A}) \supseteq A$. We will once need the following strengthening of unanimity.

Definition 4 (Strong unanimity)**.** *An allocation rule F is strongly unanimous if, for every allocation instance $I = \langle \mathcal{P}, c, B \rangle$, every agent $i \in \mathcal{N}$, every feasible set $A \in \mathcal{A}(I)$, and every profile $\boldsymbol{A}$ with $|\boldsymbol{A}| \geq 3$ and $A_{i'} = A$ for all agents $i' \in \mathcal{N} \setminus \{i\}$, we have $F(I, \boldsymbol{A}) \supseteq A$.*

Observe that both of the rules defined above are exhaustive and strongly unanimous (and thus also unanimous).

2.4 Agent Preferences

Suppose agent $i \in \mathcal{N}$ has preferences over all individual projects in $\mathbb{P}$ expressed as a strict linear order $\rhd_i$. For $\mathcal{P} \subseteq \mathbb{P}$, we denote by $\rhd_{i|\mathcal{P}}$ the restriction of $\rhd_i$ to $\mathcal{P}$. Moreover, amongst the projects in $\mathcal{P}$, agent i has an *ideal set* $top_i(\mathcal{P})$ of projects, assumed to be determined by the greedy selection procedure: $top_i(\mathcal{P}) = GREED(\mathcal{P}, \rhd_{i|\mathcal{P}})$. This approach will permit us to model what constitutes a *truthful* vote by an agent for varying shortlists $\mathcal{P}$. We call the vector $\boldsymbol{top}(\mathcal{P}) = (top_1(\mathcal{P}), \ldots, top_n(\mathcal{P}))$ the *ideal profile* given $\mathcal{P}$.

Following Goel et al. [2019] and Talmon and Faliszewski [2019], we make use of two *preference models* an agent can use to derive preference relations from her ideal set. For any ideal set $P \subseteq \mathbb{P}$, we denote by $\succeq_P$ the induced weak preference relation and by $\succ_P$ its strict part. For any two budget allocations A and A', under the *overlap preference model* we have $A \succeq_P A'$ if and only if $|A \cap P| \geq |A' \cap P|$, while under the *cost preferences model* we have $A \succeq_P A'$ if and only if $c(A \cap P) \geq c(A' \cap P)$.

Finally, for any preference relation $\succeq$ and any family of budget allocations $\mathfrak{P} \subseteq 2^{\mathbb{P}}$, we use $best(\succeq, \mathfrak{P})$ to denote the set of budget allocations that are undominated in $\mathfrak{P}$ w.r.t. $\succeq$.

3 Shortlisting Rules

While several allocation rules have been defined in the literature, this has not been the case for shortlisting rules. In this section, we therefore propose several shortlisting rules.

The first of these is what arguably is the simplest of all rules, the *nomination rule*, where every agent acts as a nominator, i.e., someone whose proposals are always all accepted.

Definition 5 (Nomination rule)**.** *The nomination rule R returns, for every shortlisting instance $I = \langle \mathbb{P}, c, B \rangle$ and shortlisting profile $\boldsymbol{P}$, the shortlist $R(I, \boldsymbol{P}) = \bigcup \boldsymbol{P}$.*

Although very natural, the nomination shortlisting rule is not effective in reducing the number of projects.

3.1 Equal-Representation Shortlisting Rules

Since the budget limit is not a hard constraint at the shortlisting stage, we might want to try to ensure that every participant has their say. Building on this idea, we introduce the *k-equal-representation shortlisting rules*, inspired by the Thiele rules for multiwinner voting [Janson, 2016].

Definition 6 (k-equal-representation shortlisting rules)**.** *Let $k \in \mathbb{N}$. The k-equal-representation shortlisting rule R returns, for any given shortlisting instance $I = \langle \mathbb{P}, c, B \rangle$ and profile $\boldsymbol{P} = (P_1, \ldots, P_n)$, the following shortlist:*

$$R(I, \boldsymbol{P}) = T\left(\underset{\substack{P \subseteq \bigcup \boldsymbol{P} \\ c(P) \leq kB}}{\text{argmax}} \sum_{i \in \mathcal{N}} \sum_{\ell=0}^{|P_i \cap P|} \frac{1}{n^\ell}\right)$$

The weight $1/n$ in the rule ensures that the rule will always select, if possible, a project proposed by the agents with the smallest number of thus-far-selected projects.

While intuitively attractive, computing shortlists under this rule is a computationally demanding task.

Proposition 1. *For any $k \in \mathbb{N}$, computing the outcome of the k-equal-representation shortlisting rule is NP-hard.*

Proof. Note that for any $k' \in \mathbb{N}$, if all projects have cost B/k', computing the outcome of the k-equal-representation shortlisting rule amounts to finding a committee of size k' with a Thiele voting rule with weights $(1, 1/n, 1/n^2, \ldots)$ in a multiwinner election [Janson, 2016]. Interestingly, the reduction presented by Aziz et al. [2015] to show that the well-known rule of proportional approval voting (PAV) is NP-hard works for all Thiele rules with decreasing weights. Since this is the case here, their reduction applies as well and show NP-hardness for the k-equal-representation rule. □

3.2 Median-Based Shortlisting Rules

One criterion frequently used for excluding projects from shortlists in practice is the similarity between them. We now introduce a family of shortlisting rules that allows us to operationalise this intuition. The idea is to cluster the projects and only select one representative project per cluster.

We call *distance* any metric over $\mathbb{P}$. For a distance δ, the *geometric median* of $P \subseteq \mathbb{P}$ is defined as $med(P) = T(\text{argmin}_{p^\star \in P} \sum_{p' \in P} \delta(p^\star, p'))$. A partition of P, denoted by $V = \{V_1, \ldots, V_p\}$, is a (k, ℓ)-*Voronoï partition* w.r.t. δ if $\sum_{V_j \in V} c(med(V_j)) \leq kB$ and for every distinct $V_j, V_{j'} \in V$ and every $p \in V_j$, we have:

- $\delta(p, med(V_j)) \leq \delta(p, med(V_{j'}))$, i.e., every project is in the cluster of its closest geometric median; and
- $\delta(p, med(V_j)) \leq \ell$, i.e., every project is within distance ℓ of the geometric median of its cluster.

Let $\mathcal{V}_{\delta,k,\ell}(P)$ be the set of all (k, ℓ)-Voronoï partitions of P w.r.t. δ. We are now ready to define our rules.

Definition 7 (k-median shortlisting rules)**.** *Let $k \in \mathbb{N}$. The k-median shortlisting rule R w.r.t. distance δ returns, for any shortlisting instance I and profile $\boldsymbol{P}$, the following shortlist:*

$$R(I, \boldsymbol{P}) = T\left(\bigcup_{V \in \mathcal{V}_{\delta,k,\ell^\star}(\bigcup \boldsymbol{P})} \{med(V_j) \mid V_j \in V\}\right)$$

Here $\ell^\star$ is the smallest ℓ such that $\mathcal{V}_{\delta,k,\ell}(\bigcup \boldsymbol{P}) \neq \emptyset$.

Note that we chose to minimise ℓ in our definition; one could similarly try to minimise k, or both ℓ and k, instead.

For most natural choices of δ, computing outcomes for the k-median shortlisting rule will be NP-hard. For instance, for the Euclidean distance, our formulation coincides with the *k-median problem*, known to be NP-hard [Kariv and Hakimi, 1979]. Still, known results on approximation algorithms and fixed-parameter tractability can be exploited here.

3.3 End-to-End Example

Let us now clarify our whole setting with an example.

Example 1. Consider a shortlisting instance $I = \langle \mathbb{P}, c, B \rangle$ with projects $\mathbb{P} = \{p_1, \ldots p_{10}\}$, cost $c(p) = 1$ for every project $p \in \mathbb{P}$, and a budget limit of $B = 3$. Now consider five agents with the following characteristics:

	Preference Order	*Awareness Set*
Agent 1	$p_1 \rhd p_4 \rhd p_5 \rhd p_{10} \rhd \ \ldots$	$\{p_1, p_4, p_5, p_{10}\}$
Agent 2	$p_1 \rhd p_2 \rhd p_6 \rhd p_4 \rhd \ \ldots$	$\{p_2, p_6\}$
Agent 3	$p_1 \rhd p_2 \rhd p_7 \rhd p_4 \rhd \ \ldots$	$\{p_2, p_7\}$
Agent 4	$p_1 \rhd p_3 \rhd p_8 \rhd p_5 \rhd \ \ldots$	$\{p_3, p_8\}$
Agent 5	$p_1 \rhd p_3 \rhd p_9 \rhd p_5 \rhd \ \ldots$	$\{p_3, p_9\}$

At the shortlisting stage, if agents are truthful they will all submit their ideal set w.r.t. their awareness set, leading to the profile $(\{p_1, p_4, p_5\}, \{p_2, p_6\}, \{p_2, p_7\}, \{p_3, p_8\}, \{p_3, p_9\})$. The outcome would be $\mathbb{P} \setminus \{p_{10}\}$ for the nomination rule, and $\{p_1, p_2, p_3, p_4, p_6, p_7\}$ for the 2-equal-representation rule.

Suppose the shortlist is $\mathcal{P} = \mathbb{P} \setminus \{p_{10}\}$. All agents now become aware of all the shortlisted projects. The truthful profile for the allocation stage would then be $\boldsymbol{A} = (\{p_1, p_4, p_5\}, \{p_1, p_2, p_6\}, \{p_1, p_2, p_7\}, \{p_1, p_3, p_8\}, \{p_1, p_3, p_9\})$. Both the greedy-approval and the approval-maximising allocation rules would select the budget allocation $A = \{p_1, p_2, p_3\}$. $\triangle$

3.4 Axioms for Shortlisting Rules

We now present some axioms for shortlisting rules. The first one, *non-wastefulness*, stipulates that it should be possible to exhaust the budget in the allocation stage.

Definition 8 (Non-wastefulness)**.** *A shortlisting rule R is* non-wasteful *if, for every shortlisting instance $I = \langle \mathbb{P}, c, B \rangle$ and profile $\boldsymbol{P}$, either $c(R(I, \boldsymbol{P})) \geq B$ or $R(I, \boldsymbol{P}) = \bigcup \boldsymbol{P}$.*

The second axiom we put forward here encodes the idea that every agent should be *represented* by the outcome.

Definition 9 (Representation efficiency)**.** *For shortlisting instance $I = \langle \mathbb{P}, c, B \rangle$ and a shortlisting profile $\boldsymbol{P}$, a shortlist $\mathcal{P} \subseteq \mathbb{P}$ is representatively dominated if there exists a set $\mathcal{P}' \subseteq \mathbb{P}$ with $c(\mathcal{P}') \leq c(\mathcal{P})$, and $|\mathcal{P}' \cap P_i| \geq |\mathcal{P} \cap P_i|$ for all $i \in \mathcal{N}$, with a strict inequality for at least one agent.*

A shortlisting rule R is representatively efficient *if its outcome is never representatively dominated.*

Let us now see how our shortlisting rules perform w.r.t. these axioms. The proof of the next result is immediate.

Fact 2. *The nomination shortlisting rule is both non-wasteful and representatively efficient.*

Proposition 3. *For $k \geq 2$, the k-equal-representation shortlisting rule R is non-wasteful. For $k \geq 1$, the k-equal-representation shortlisting rule is representatively efficient.*

Proof. First, assume that there exists a $k \geq 2$ such that R is wasteful. Then there must exist $I = \langle \mathbb{P}, c, B \rangle$ and $\boldsymbol{P}$ such that $c(R(I, \boldsymbol{P})) < B$ and $R(I, \boldsymbol{P}) \neq \bigcup \boldsymbol{P}$. There exists then a project $p \in \bigcup \boldsymbol{P} \setminus R(I, \boldsymbol{P})$ such that the representation score of the set $R(I, \boldsymbol{P}) \cup \{p\}$ is higher than that of $R(I, \boldsymbol{P})$. Moreover, since $c(p) \leq B$, we have $c(R(I, \boldsymbol{P}) \cup \{p\}) \leq 2B \leq kB$. Hence, $R(I, \boldsymbol{P}) \cup \{p\}$ would have been returned by R, yielding a contradiction.

The fact that, for every $k \geq 1$, the k-equal-representation shortlisting rule is representatively efficient is immediate from the choice of the weight $1/n$ in Definition 6. □

Proposition 4. *For $k \geq 2$, the k-median shortlisting rule is non-wasteful. But there exists no $k \in \mathbb{N}$ such that the k-median shortlisting rule is representatively efficient.*

Proof. The proof that, for $k \geq 2$, the k-median shortlisting rule is non-wasteful is similar to the corresponding part of the proof of Proposition 3. Indeed, for no shortlisting instance I and profile $\boldsymbol{P}$, can there be a $p \in \bigcup \boldsymbol{P} \setminus R(I, \boldsymbol{P})$ such that $c(R(I, \boldsymbol{P}) \cup \{p\}) \leq kB$, since selecting this p would lead to a smaller within-cluster distance.

Finally, a k-median shortlisting rule is not efficiently representative, since the agents are not taken into account. □

4 First-Stage Strategyproofness

We now turn to the analysis of strategic interaction in our end-to-end model, with a focus on the shortlisting stage.

Let us first discuss the information available to a manipulator. In the classical voting framework [Zwicker, 2016], it is assumed that the manipulator has access to all the other ballots before submitting her own. In our setting, when considering a manipulator choosing which proposal to submit during the first stage, the same assumption is reasonable w.r.t. the proposals about to be submitted by the other agents during the first stage—but not w.r.t. the ballots the other agents are going to submit during the second stage, *after* the shortlist will have been determined. Indeed, the set of actions for the second stage depends on the proposal of the manipulator in the first stage. We explore here three possibilities. In the first two cases, a manipulator in the first stage is unsure what will happen during the second stage, but assumes that either the worst scenario will be realised (*pessimistic manipulation*) or the best one (*optimistic manipulation*). In the third case, she knows the other agents' true preferences and trusts they will vote accordingly (*anticipative manipulation*).

Let us fix some further notation. For a given allocation rule F, allocation instance $I = \langle \mathcal{P}, c, B\rangle$, profile $\boldsymbol{A}$, and agent $i \in \mathcal{N}$, let $A_i^\star(I, \boldsymbol{A}, F)$ be defined as the ballot $T(\mathit{best}(\succ_{top_i(\mathcal{P})}, \{F(I, (\boldsymbol{A}_{-i}, A_i')) \mid A_i' \subseteq \mathcal{P}\}))$, the best response of i to $\boldsymbol{A}$. When clear from the context, we omit I, $\boldsymbol{A}$, and/or F. Also recall that every agent $i \in \mathcal{N}$ can determine an ideal set $top_i(\mathcal{P})$ for any given set $\mathcal{P}$, which induces a preference relation $\succeq_{top_i(\mathcal{P})}$ on budget allocations.

Definition 10 (Successful manipulation)**.** *Let R be a shortlisting rule, F an allocation rule, $I_1 = \langle \mathbb{P}, c, B\rangle$ a shortlisting instance, $\boldsymbol{P}$ a shortlisting profile, and $P_i' \subseteq \mathbb{P}$ an alternative proposal for agent $i \in \mathcal{N}$. Consider the shortlists $\mathcal{P} = R(I_1, \boldsymbol{P})$ and $\mathcal{P}' = R(I_1, (\boldsymbol{P}_{-i}, P_i'))$, determining the allocation instances $I_2 = \langle \mathcal{P}, c, B\rangle$ and $I_2' = \langle \mathcal{P}', c, B\rangle$, and abbreviate $F(I_2, (\boldsymbol{A}_{-i}, A_i^\star(I_2, \boldsymbol{A})))$ as $F^\star(I_2, \boldsymbol{A})$ and $F(I_2', (\boldsymbol{A}'_{-i}, A_i^\star(I_2', \boldsymbol{A}')))$ as $F^\star(I_2', \boldsymbol{A}')$, for any two approval profiles $\boldsymbol{A}$ on $\mathcal{P}$ and $\boldsymbol{A}'$ on $\mathcal{P}'$. Then we say that:*

- *P_i' is a successful pessimistic manipulation if, for all profiles $\boldsymbol{A}$ on $\mathcal{P}$ and $\boldsymbol{A}'$ on $\mathcal{P}'$, it is the case that $F^\star(I_2', \boldsymbol{A}') \succeq_{top_i(\mathcal{P}\cup\mathcal{P}')} F^\star(I_2, \boldsymbol{A})$, with a strict preference for at least one pair $(\boldsymbol{A}, \boldsymbol{A}')$.*
- *P_i' is a successful optimistic manipulation if, for at least one profile $\boldsymbol{A}$ on $\mathcal{P}$ and one profile $\boldsymbol{A}'$ on $\mathcal{P}'$, it is the case that $F^\star(I_2', \boldsymbol{A}') \succ_{top_i(\mathcal{P}\cup\mathcal{P}')} F^\star(I_2, \boldsymbol{A})$.*
- *P_i' is a successful anticipative manipulation if, for the two profiles $\boldsymbol{A} = \boldsymbol{top}(\mathcal{P})$ and $\boldsymbol{A}' = \boldsymbol{top}(\mathcal{P}')$, it is the case that $F^\star(I_2', \boldsymbol{A}') \succ_{top_i(\mathcal{P}\cup\mathcal{P}')} F^\star(I_2, \boldsymbol{A})$.*

Thus, a pessimist is pessimistic w.r.t. the advantages she can gain from manipulating. For optimists it is the other way round. Finally, an anticipative manipulator knows everyone's preferences on both $\mathcal{P}$ and $\mathcal{P}'$ and uses them to predict their votes during the second stage for both scenarios. Note that this definition applies seamlessly under the overlap preference model and the cost preference model.

We are looking for rules that do not allow for successful manipulation, i.e., that are first-stage strategyproof (FSSP). We distinguish two cases: either the manipulator is *restricted to* her awareness set (R-FSSP) or she can also propose any of the projects proposed by others (*unrestricted*, U-FSSP).

Definition 11 (FSSP)**.** *For a given preference model, a pair $\langle R, F\rangle$ consisting of a shortlisting and an allocation rule is R-FSSP w.r.t. a given type of manipulation, if for every shortlisting instance $\langle \mathbb{P}, c, B\rangle$, every awareness profile $\boldsymbol{C} = (C_1, \ldots, C_n)$, every shortlisting profile $\boldsymbol{P} = (P_1, \ldots, P_n)$ where $P_{i'} \subseteq C_{i'}$ for all $i' \in \mathcal{N}$, and every agent $i \in \mathcal{N}$, there is no $P_i' \subseteq C_i$ such that submitting P_i' instead of $top_i(C_i)$ is a successful manipulation for i.*

When $P_i' \subseteq C_i \cup \bigcup \boldsymbol{P}$ and we consider $top_i(C_i \cup \bigcup \boldsymbol{P})$ instead of $top_i(C_i)$, we say that $\langle R, F\rangle$ is U-FSSP.

Thus, under U, agents are assumed to *first* gain access to everyone's proposals and *then* decide whether or not to vote truthfully. We are going to use FSSP-P to denote FSSP w.r.t. pessimistic manipulation attempts, FSSP-O for optimistic manipulation, and FSSP-A for anticipative manipulation.

The following result summarises how the different notions introduced relate to each other, where $\mathfrak{X}$ implying $\mathfrak{X}'$ means that any pair $\langle R, F\rangle$ satisfying $\mathfrak{X}$ also satisfies $\mathfrak{X}'$.

Proposition 5. *The following implications hold for both the overlap preference model and the cost preference model:*

- *R-FSSP-O implies R-FSSP-A and R-FSSP-P.*
- *U-FSSP-O implies U-FSSP-A and U-FSSP-P.*
- *R-FSSP implies U-FSSP for all types of manipulation.*

Proof. To see that the last of these claims is true, observe that U-FSSP is a special case of R-FSSP, namely when the manipulator can conceive of all the proposed projects, i.e., when $C_i = \bigcup \boldsymbol{P}$.[2] The other claims are immediate. □

All results in this section hold for both the overlap and the cost pref. models, so we will omit them in our statements.

[2] This may be counter-intuitive at first, but as explained at the end of Section 4.1, U-FSSP does not imply R-FSSP.

4.1 Awareness-Restricted Manipulation

We start by proving an impossibility theorem, stating that, when manipulators are restricted to their awareness sets, no pair of reasonable rules can be first-stage strategyproof.

Theorem 6. *Any pair $\langle R, F\rangle$ of a non-wasteful shortlisting rule R and an exhaustive allocation rule F is neither R-FSSP-P nor R-FSSP-A (and thus also not R-FSSP-O).*

Proof. We provide a proof for R-FSSP-P, but the same proof also goes through for R-FSSP-A. The claim for R-FSSP-O then follows from Proposition 5.

Let $I = \langle \mathbb{P}, c, B\rangle$ be the shortlisting instance with $\mathbb{P} = \{p_1, p_2\}$, $c(p_1) = c(p_2) = 1$, and $B = 1$. Suppose there are two agents, with $p_2 \rhd_1 p_1$ and $C_1 = \{p_1\}$ as well as $p_1 \rhd_2 p_2$ and $C_2 = \{p_2\}$. So each of them is aware only of the project they like less. Thus, the truthful shortlisting profile is $\boldsymbol{P} = (\{p_1\}, \{p_2\})$. Since R is non-wasteful, we know that $|R(I, \boldsymbol{P})| \geq 1$. Let us then consider the three possible outcomes of R on I and $\boldsymbol{P}$.

In case $R(I, \boldsymbol{P}) = \{p_1\}$, whichever way the two agents vote in the allocation stage, as F is exhaustive, the final budget allocation must be $\{p_1\}$. If agent 1 manipulates by not proposing any project for the shortlist, then $\{p_2\}$ will get shortlisted, since R is non-wasteful. In that case, $\{p_2\}$ will also be the final budget allocation, since F is exhaustive. It is clear that for either preference model, agent 1 prefers $\{p_2\}$ over $\{p_1\}$. So agent 1 has an incentive to pessimistically manipulate in this case.

The case of $R(I, \boldsymbol{P}) = \{p_2\}$ is symmetric to the previous one, except that now agent 2 can manipulate.

Finally, consider the case $R(I, \boldsymbol{P}) = \{p_1, p_2\}$. W.l.o.g., suppose the final budget allocation is $\{p_1\}$ in case both agents vote truthfully. Then, just as in the first case, agent 1 has an incentive to submit an empty set of proposals instead, as that guarantees a final budget allocation of $\{p_2\}$. □

Note that the scenario used in the proof shows that strategyproofness under U does *not* imply strategyproofness under R. Indeed, under U no agent would have an incentive to manipulate in this scenario, as they would have all the information they need to submit an optimal truthful proposal.

4.2 Unrestricted Manipulation

For the case of U, let us start with the nomination rule. We first prove that it is immune to pessimistic manipulation when used with a strongly unanimous allocation rule.

Proposition 7. *For every allocation rule F that is exhaustive and strongly unanimous, the pair $\langle R, F\rangle$, where R is the nomination rule, is U-FSSP-P.*

Proof. Let $I = \langle \mathbb{P}, c, B\rangle$ be a shortlisting instance, $\boldsymbol{C}$ an awareness profile, and $\boldsymbol{P}$ the truthful shortlisting profile. Consider an agent $i \in \mathcal{N}$. Let $P_i = top_i(C_i \cup \bigcup \boldsymbol{P})$. From Definition 5, we know that if i submits P_i' instead of P_i, the set of shortlisted projects will become $P_i' \cup (\bigcup_{i' \in \mathcal{N} \setminus \{i\}} P_{i'})$. Since now weakly fewer projects from $top_i(C_i \cup \bigcup \boldsymbol{P})$ are shortlisted, none of the budget allocations newly reachable will be strictly better for i. Moreover, strong unanimity entails that for every exhaustive budget allocation A, there is a profile realising it, namely, the one where every agent except i submits A. This directly implies that i cannot be better off by pessimistically manipulating. □

On the other hand, we can show that the nomination shortlisting rule paired with either one of the allocation rules we introduced is not U-FSSP-A (and thus not U-FSSP-O).

Example 2. Recall Example 1, where for both the greedy-approval and the approval-maximising rule the outcome was $A = \{p_1, p_2, p_3\}$. Assume now that agent 1 submits $\{p_4, p_5, p_{10}\}$ instead of $\{p_1, p_4, p_5\}$ in the shortlisting stage. The shortlist then becomes $\mathcal{P} = \mathbb{P} \setminus \{p_1\}$. In the second stage, all agents now approve of their second, third, and fourth most preferred projects, leading to the budget allocation $\{p_2, p_4, p_5\}$. It is clear that under both of our preference models, this is better than A for agent 1. △

Unfortunately, also the other shortlisting rules we defined turn out to not be first-stage strategyproof.

Proposition 8. *For all $k \in \mathbb{N}$, the pair $\langle R, F\rangle$, where R is the k-equal-representation shortlisting rule and F is a unanimous allocation rule, is neither U-FSSP-P nor U-FSSP-O.*

Proof. We first prove the claim for $k = 1$ and then explain how to generalise to any $k \in \mathbb{N}$. Let $I = \langle \mathbb{P}, c, B\rangle$ be a shortlisting instance with $\mathbb{P} = \{p_1, \ldots p_4\}$, $c(p_2) = 2$, $c(p) = 1$ for all $p \in \mathbb{P} \setminus \{p_2\}$, and $B = 2$. Consider this scenario:

	Preference Order	*Awareness Set*
Agent 1	$p_3 \rhd p_4 \rhd p_2 \rhd p_1$	$\{p_1, p_2, p_3, p_4\}$
Agent 2	$p_1 \rhd p_2 \rhd p_3 \rhd p_4$	$\{p_1, p_2\}$
Agent 3	$p_2 \rhd p_1 \rhd p_3 \rhd p_4$	$\{p_2\}$

Consider the 1-equal-representation shortlisting rule. Under the truthful profile $\boldsymbol{P} = \{\{p_3, p_4\}, \{p_1, p_2\}, \{p_2\}\}$, the shortlist would be $\mathcal{P} = \{p_2\}$. Note that $\mathcal{P} \cap top_1(\bigcup \boldsymbol{P}) = \emptyset$. Assume now that agent 1 submits $\{p_1, p_3\}$ instead of $\{p_3, p_4\}$. Then the outcome of the first stage becomes $\mathcal{P}' = \{p_1, p_3\}$. It is clear that every $A' \in \mathcal{A}(\langle \mathcal{P}', c, B\rangle)$ is weakly preferred by agent 1 to every $A \in \mathcal{A}(\langle \mathcal{P}, c, B\rangle)$ and some are strictly preferred (the ones in which p_3 appears). Since the allocation rule is unanimous, all these budget allocations can be reached (when every agent submit the budget allocation) so agent 1's manipulation is pessimistically successful.

To generalise to $k > 1$, add $3(k-1)$ agents in groups of 3. Each group can conceive and approve of two new projects. It is easy to check that all the new projects will always be shortlisted, so we are back to the scenario above. □

Proposition 9. *For all $k \in \mathbb{N}$, the pair $\langle R, F\rangle$, where R is the k-median shortlisting rule based on the Euclidean distance over $\mathbb{R}^2$ and F is an unanimous allocation rule, is neither U-FSSP-P nor U-FSSP-O.*

Proof. We first prove the claim for $k = 1$ and then explain how to generalise it to all $k \in \mathbb{N}$. Consider the shortlisting instance $I = \langle \mathbb{P}, c, B\rangle$ with $\mathbb{P} = \{p_1, \ldots, p_6\}$, all projects have cost 1, and $B = 3$. Suppose the distance δ is the usual distance in the plane, with the projects as in the figure below:

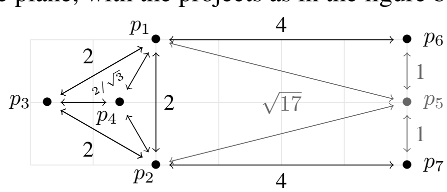

Consider two agents such that $p_1 \rhd_1 p_2 \rhd_1 p_3 \rhd_1 \ldots$ and $C_1 = \{p_1, p_2, p_3, p_5\}$, and $p_4 \rhd_2 p_6 \rhd_2 p_7 \rhd_2 \ldots$ and $C_2 = \{p_4, p_6, p_7\}$. The truthful profile is $(C_1 \setminus \{p_5\}, C_2)$ which will lead to the clusters $\{p_1, p_2, p_3, p_4\}$, $\{p_5\}$, and $\{p_6\}$ for the 1-median shortlisting rule. The set of shortlisted projects then is $\mathcal{P} = \{p_4, p_6, p_7\}$. Note that $\mathcal{P} \cap top_1(\mathbb{P}) = \emptyset$.

Now assume that agent 1 submits $\{p_1, p_2, p_5\}$. Then there will be the clusters $\{p_1\}$, $\{p_2, p_4\}$, and $\{p_5, p_6, p_7\}$. The shortlist would then be $\mathcal{P}' = \{p_1, p_2, p_5\}$. For the same reason as in the proof of Proposition 8, since F is unanimous, agent 1's manipulation is pessimistically successful.

To extend this to all $k > 1$, add $k - 1$ agents, all knowing and approving of three new projects of cost 1. All the new projects are placed uniformly on a circle with centre p_4 and a radius large enough so that all new projects will be in their own cluster. Then all the new projects will be shortlisted and we are back to the original case for $k = 1$. □

Note that the previous statement can be generalised to other distances; in particular, it holds for any surjective distance.

To conclude this section, note that the proofs of Propositions 8 and 9 do not extend to U-FSSP-A as that notion is about a specific profile for which we have no relevant information. However, the following facts hold for U-FSSP-A.

Fact 10. *For no $k \in \mathbb{N}$ is the pair $\langle R, F\rangle$ U-FSSP-A, when R is the k-equal-representation shortlisting rule and F is either the greedy-approval or the approval-maximising rule.*

Fact 11. *For no $k \in \mathbb{N}$ is the pair $\langle R, F \rangle$ U-FSSP-A, when R is the k-median shortlisting rule and F is either the greedy-approval or the approval-maximising rule.*

The counterexample used in the proof of Proposition 9 also works for Fact 11. That of Proposition 8 can be made to work for Fact 10 by slightly changing the agent preferences (making p_1 and p_3 the most preferred projects of many agents).

5 Second-Stage Strategyproofness

To round off the picture of strategic behaviour in PB, we now discuss strategyproofness for the second stage alone, once the shortlist has been settled. We first recall the notion of strategyproofness (SP) as well as a relaxation, *approximate strategyproofness*, first sketched by Goel et al. [2019].

Definition 12 (Strategyproofness)**.** *For a given preference model, an allocation rule F is strategyproof over a class $\mathcal{I}$ of allocation instances, if for every $I = \langle \mathcal{P}, c, B \rangle \in \mathcal{I}$, every profile $\boldsymbol{A}$, and every agent $i \in \mathcal{N}$, we have:*

$$F(I, (\boldsymbol{A}_{-i}, top_i(\mathcal{P}))) \succeq_{top_i(\mathcal{P})} F(I, \boldsymbol{A})$$

Moreover, F is SP if it is SP over the class of all instances.

Definition 13 (Approximate strategyproofness)**.** *For a given preference model, an allocation rule F is approximately strategyproof over a class of allocation instances $\mathcal{I}$, if for every $I = \langle \mathcal{P}, c, B \rangle \in \mathcal{I}$, every profile $\boldsymbol{A}$, and every agent $i \in \mathcal{N}$, there exists a project $p \in \mathcal{P}$ such that:*

$$F(I, (\boldsymbol{A}_{-i}, top_i(\mathcal{P}))) \cup \{p\} \succeq_{top_i(\mathcal{P})} F(I, \boldsymbol{A}).$$

Moreover, F is approximately strategyproof if it is approximately strategyproof over the class of all instances.

Let us now introduce the class of allocation instances with *unit costs:* $I = \langle \mathcal{P}, c, B \rangle$ belongs to $\mathcal{I}_{unit}$ if and only if, for any two projects $p, p' \in \mathcal{P}$, we have $c(p) = c(p')$. Observe that the instances in $\mathcal{I}_{unit}$ are equivalent to approval-based multiwinner elections [Faliszewski et al., 2017], with a set of candidates $\mathcal{P}$ and a committee size of B.

5.1 The Greedy-Approval Rule

Recall that the greedy-approval rule for multiwinner voting is SP for overlap preferences [Peters, 2018]. On the PB side, it immediately implies that the greedy-approval rule is SP over $\mathcal{I}_{unit}$ for both overlap and cost preferences. However, this result does not extend to arbitrary instances, as already noted by Goel et al. [2019], who also showed that the greedy-approval rule nevertheless is approximately strategyproof. Next, we sketch what we consider to be a much simpler proof of this latter insight.

Proposition 12 (Goel et al., 2019)**.** *The greedy-approval rule is approximately SP for the cost preference model.*

Proof. Let F be the greedy-approval rule. For any given $I = \langle \mathcal{P}, c, B \rangle$, consider $I' = \langle \mathcal{P}', c', B \rangle$, where projects in $\mathcal{P}$ have been split into sets of subprojects, each of cost 1. We thus have $I' \in \mathcal{I}_{unit}$. Any given profile $\boldsymbol{A}$ can be split in the same manner to obtain $\boldsymbol{A}'$. Now, it is clear that the approval scores of the projects in $\boldsymbol{A}'$ are equal to those of the projects in $\mathcal{P}$ they come from in $\boldsymbol{A}$. Assume that the tie-breaking rule is extended in a consistent way from projects in $\mathcal{P}$ to projects in $\mathcal{P}'$. Then we know that there exists at most one project $p \in \mathcal{P}$ such that $F(I', \boldsymbol{A}')$ contains a proper subset of its subprojects. Let $A \subseteq \mathcal{P}$ be the budget allocation corresponding to $F(I', \boldsymbol{A}')$, where subprojects have been merged back into the original projects. Since F is SP over $\mathcal{I}_{unit}$, a potential manipulator would not prefer A to $F(I, \boldsymbol{A}) \cup \{p\}$ under the cost preference model. □

Observe that the proof crucially relies on the cost model, which ensures that an agent will appreciate a project as much as the sum of its subprojects. This is not the case under the overlap preference model, and indeed it is easy to find counterexamples showing that approximate strategyproofness of the greedy-approval rule fails for that model.

5.2 The Approval-Maximising Rule

We now turn to the approval-maximising rule. As we are going to see, rather surprisingly, this rule is not approximately strategyproof. First, consider the following example, which shows that when we do not use the canonical tie-breaking rule, then we do not even get strategyproofness for $\mathcal{I}_{unit}$.

Example 3. Consider an allocation instance $I = \langle \mathcal{P}, c, B \rangle$ in $\mathcal{I}_{unit}$ with $\mathcal{P} = \{p_1, \dots, p_7\}$ and $B = 4$. Consider two agents with $p_1 \rhd_1 p_2 \rhd_1 p_3 \rhd_1 p_4 \rhd_1 \dots$ and $p_4 \rhd_2 p_5 \rhd_2 p_6 \rhd_2 p_7 \rhd_2 \dots$ Let us apply the approval-maximising rule to the resulting profile. But suppose that rather than using T, ties are broken such that $\{p_1, p_2, p_3, p_5\}$ comes first, then $\{p_4, p_5, p_6, p_7\}$, and then all the other subsets of $\mathcal{P}$, in any order. When agents truthfully report their ideal sets, the outcome is $A = \{p_4, p_5, p_6, p_7\}$. Now, if agent 1 submits $\{p_1, p_2, p_3\}$ instead of $\{p_1, p_2, p_3, p_4\}$, the outcome would be $A' = \{p_1, p_2, p_3, p_5\}$. Both under the overlap preference model and the cost preference model, agent 1 prefers A' over $A \cup \{p\}$, for any $p \in \{p_1, p_2, p_3\}$. $\triangle$

Example 3 relies on the specific tie-breaking rule employed. Next, we show that this problem does not occur for the canonical (and thus any lexicographic) tie-breaking rule.

Proposition 13. *The approval-maximising rule is SP over $\mathcal{I}_{unit}$ under both the cost and the overlap preference models.*

Proof. First, note that on $\mathcal{I}_{unit}$ the two preference models coincide. Then, since the greedy-approval rule is SP over $\mathcal{I}_{unit}$, we are done if we can show that the two rules coincide for $\mathcal{I}_{unit}$. This clearly is the case: By greedily selecting unit-cost projects with the highest approval score, you also maximise total approval. To rule out the possibility that the approval-maximising rule might select a different budget allocation with maximal total approval, observe that the lexicographic definition of the canonical tie-breaking rule T ensures that (i) breaking ties at the project level between projects of equal approval score and (ii) breaking ties at the budget allocation level between budget allocations of equal total approval score lead to the same outcome. $\square$

This positive result not withstanding, for arbitrary allocation instances the approval-maximising rule ceases to be strategyproof—or even approximately strategyproof. So in this sense the greedy-approval rule really is the superior rule.

Example 4. Consider the instance $I = \langle \mathcal{P}, c, B \rangle$ with $\mathcal{P} = \{p_1, \dots p_8\}$, $c(p_1) = c(p_2) = c(p_3) = 4$, $c(p_4) = c(p_5) = 6$ and $c(p_6) = c(p_7) = c(p_8) = 3$, and $B = 12$. Consider three agents with the following preferences: $p_6 \rhd_1 p_7 \rhd_1 p_8 \rhd_1 \dots$; $p_4 \rhd_2 p_5 \rhd_2 \dots$; and $p_1 \rhd_3 p_2 \rhd_3 p_3 \rhd_3 \dots$.

Suppose that ties are broken by the canonical tie-breaking rule. Then, when agents are truthful, the outcome of the approval-maximising rule is $A = \{p_1, p_2, p_3\}$. Now, if agent 1 were to submit $\{p_5, p_6, p_7\}$ instead of $\{p_6, p_7, p_8\}$, the outcome would be $A' = \{p_5, p_6, p_7\}$. Is it clear that for every project $p \in \mathcal{P}$, agent 1 prefers A' over $A \cup \{p\}$ under both the overlap and the cost preference models. $\triangle$

6 Conclusion

We have initiated the study of PB in a way that accounts not only for the allocation stage but also for the shortlisting stage preceding it. This has prompted several proposals for concrete shortlisting rules, and allowed us to analyse the incentives of agents to manipulate the shortlisting stage, in view of how their actions affect the ultimate outcome of the allocation stage. We have complemented this analysis with a study of strategic behaviour during the allocation stage.

This paper is a first step towards the principled investigation of the full PB process. There are still many other features deserving attention. For instance, it would be interesting to consider other allocation rules, not based on approval scores, e.g., proportional rules [Aziz et al., 2018]. More generally, other types of interaction between the two stages can be investigated, such as devising allocation rules taking into account not only the outcome of the shortlisting stage but also the shortlisting profile itself, to enforce some kind of fairness across the two stages.

References

G. Allegretti and S. Antunes. The Lisbon participatory budget: Results and perspectives on an experience in slow but continuous transformation. *Field Actions Science Reports*, (Special Issue 11), 2014.

H. Aziz and N. Shah. Participatory budgeting: Models and approaches. In *Pathways between Social Science and Computational Social Science: Theories, Methods and Interpretations*. Springer-Verlag, 2020.

H. Aziz, S. Gaspers, J. Gudmundsson, S. Mackenzie, N. Mattei, and T. Walsh. Computational aspects of multi-winner approval voting. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 107–115, 2015.

H. Aziz, B. E. Lee, and N. Talmon. Proportionally representative participatory budgeting: Axioms and algorithms. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 23–31, 2018.

G. Benade, S. Nath, A. D. Procaccia, and N. Shah. Preference elicitation for participatory budgeting. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, pages 376–382, 2017.

Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.

Y. Cabannes. Participatory budgeting: A significant contribution to participatory democracy. *Environment and Urbanization*, 16(1):27–46, 2004.

B. Fain, A. Goel, and K. Munagala. The core of the participatory budgeting problem. In *Proceedings of the 12th International Workshop on Internet and Network Economics (WINE)*, pages 384–399. Springer-Verlag, 2016.

P. Faliszewski, P. Skowron, A. Slinko, and N. Talmon. Multiwinner voting: A new challenge for social choice theory. In U. Endriss, editor, *Trends in Computational Social Choice*. AI Access, 2017. Chapter 2.

P. Faliszewski, A. Slinko, and N. Talmon. Multiwinner rules with variable number of winners. In *Proceedings of the 24th European Conference on Artificial Intelligence (ECAI)*, 2020.

T. Fluschnik, P. Skowron, M. Triphaus, and K. Wilker. Fair knapsack. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, pages 1941–1948, 2019.

A. Goel, A. K. Krishnaswamy, S. Sakshuwong, and T. Aitamurto. Knapsack voting: Voting mechanisms for participatory budgeting. *ACM Transaction on Economics and Computation*, pages 8:1–8:27, 2019.

A. K. Jain and R. C. Dubes. *Algorithms for Clustering Data*. Prentice-Hall, 1988.

Pallavi Jain, Krzysztof Sornat, and Nimrod Talmon. Participatory budgeting with project interactions. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 386–392, 2020. doi: 10.24963/ijcai.2020/54.

S. Janson. Phragmén's and Thiele's election methods. *arXiv preprint arXiv:1611.08826*, 2016.

O. Kariv and S. L. Hakimi. An algorithmic approach to network location problems. II: The p-medians. *SIAM Journal on Applied Mathematics*, 37(3):539–560, 1979.

D. M. Kilgour. Approval elections with a variable number of winners. *Theory and Decision*, 81(2):199–211, 2016.

M. Lackner and J. Maly. Approval-based shortlisting. *arXiv preprint arXiv:2005.07094*, 2020.

D. Peters. Proportionality and strategyproofness in multiwinner elections. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1549–1557, 2018.

Simon Rey, Ulle Endriss, and Ronald de Haan. Designing participatory budgeting mechanisms grounded in judgment aggregation. In *Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning (KR)*, 2020.

A. Shah, editor. *Participatory budgeting*. Public Sector Governance and Accountability Series. The World Bank, Washington, DC, 2007.

N. Talmon and P. Faliszewski. A framework for approval-based budgeting methods. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, pages 2181–2188, 2019.

W. S. Zwicker. Introduction to the theory of voting. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*. Cambridge University Press, 2016. Chapter 2.