

Fair Lotteries for Participatory Budgeting

Haris Aziz, Xinhang Lu, Mashbat Suzuki, Jeremy Vollen, Toby Walsh

UNSW Sydney, Australia

{haris.aziz, xinhang.lu, mashbat.suzuki, j.vollen, t.walsh}@unsw.edu.au

Abstract

In pursuit of participatory budgeting (PB) outcomes with broader fairness guarantees, we initiate the study of lotteries over discrete PB outcomes. As the projects have heterogeneous costs, the amount spent may not be equal ex ante and ex post. To address this, we develop a technique to bound the amount by which the ex-post spend differs from the ex-ante spend—the property is termed *budget balanced up to one project (BB1)*. With respect to fairness, we take a best-of-both-worlds perspective, seeking outcomes that are both ex-ante and ex-post fair. Towards this goal, we initiate a study of ex-ante fairness properties in PB, including Individual Fair Share (IFS), Unanimous Fair Share (UFS) and their stronger variants, as well as Group Fair Share (GFS). We show several incompatibility results between these ex-ante fairness notions and existing ex-post concepts based on justified representation. One of our main contributions is a randomized algorithm which simultaneously satisfies ex-ante Strong UFS, ex-post full justified representation (FJR) and ex-post BB1 for PB with binary utilities.

1 Introduction

Participatory Budgeting (PB) is one of the exciting democratic paradigms that facilitates members of a community, municipality, or town to collectively make public project funding decisions. Considering that PB generalizes classical voting and committee selection problems in social choice theory, it also poses interesting axiomatic and algorithmic research challenges (Aziz and Shah 2021; Rey and Maly 2023). A major effort underway in computational social choice is to design meaningful axioms that capture elusive properties such as fairness and representation, and to design computationally efficient algorithms that satisfy such axioms. There is also a movement to apply such rules in various countries (see, e.g., <https://equalshares.net>).

While existing algorithms for selecting discrete projects satisfy meaningful proportional representation properties, they do not provide any minimal representation guarantees to individual voters, even in an ex-ante sense. For example, existing rules allow for the possibility that certain voters do not like any of the projects selected by the PB process. Indeed, this is an inevitable reality for any deterministic algo-

rithm. One approach to address this issue is the use of randomization to achieve stronger fairness properties ex-ante. This approach has been successful in various contexts, such as resource allocation (Bogomolnaia and Moulin 2001), voting (Bogomolnaia, Moulin, and Stong 2005), and committee voting (Cheng et al. 2020).

In this paper, we initiate the study of randomization for PB. We begin by studying the problem of implementation in PB: *Given a marginal probability for each project, how can we compute a probability distribution (or lottery) over discrete outcomes that realizes these probabilities?* While this question has already been answered in the PB special case of committee voting (Cheng et al. 2020), we find that the PB setting with heterogeneous costs gives rise to significant obstacles on this front. Since, unlike committee voting, we cannot guarantee that the total spend is equal ex-ante and ex-post, we develop a technique which bounds the amount by which the ex-post spend differs from the ex-ante spend.

Given our new implementation approach, we then return to the objective of strong ex-ante fairness properties in PB. Our goal is to achieve ex-ante fairness while also ensuring that our lotteries only include ex-post fair outcomes, thus guaranteeing desirable fairness, both ex-ante and ex-post. This approach is known as the *best-of-both-worlds fairness* perspective, and has been employed in other adjacent contexts, such as fair division (see, e.g., Aziz et al. 2023a) and committee voting (Aziz et al. 2023b). In this paper, we define a hierarchy of ex-ante fairness properties for PB and then consider a hierarchy of PB settings, giving algorithms which guarantee best-of-both-worlds fairness and/or incompatibility results for each setting considered.

1.1 Our Results

In Section 3, we tackle the question of implementation in PB. We first show that, unlike the committee voting setting, fractional outcomes cannot always be implemented by a lottery over integral outcomes using the same amount of budget. Given this, we define a well-motivated axiom for lotteries — *budget balanced up to one (BB1)* — which enforces a natural bound on the amount by which the total ex-post cost can differ from the cost of the fractional outcome the lottery implements. We then demonstrate an approach which gives an implementation satisfying our axiom for any fractional outcome, and complement this result by showing that

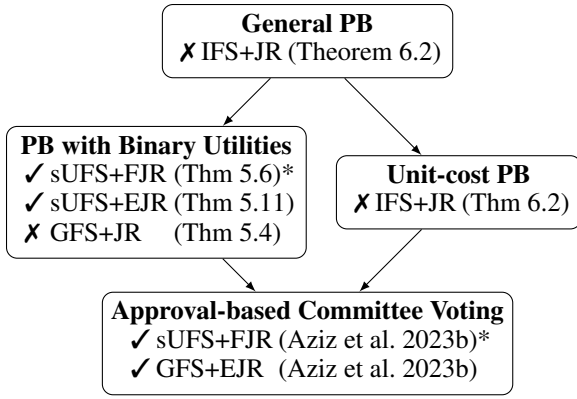


Figure 1: Summary of best-of-both-worlds fairness results in PB and special cases. Arrows point from generalizations to special cases. Compatibility results are represented by \checkmark and impossibility results by \times . (*) denotes exponential-time results. We abbreviate Strong UFS to sUFS.

a lottery satisfying a natural *up to any* strengthening of our axiom may not always exist.

In Section 4, we extend a hierarchy of well-studied properties capturing ex-ante fairness from the committee voting setting to our most general PB setting. The “fair share” hierarchy of fairness axioms begins with the basic notion that each voter should receive at least a $1/n$ fraction of their optimal utility. On the other hand, the “strong fair share” hierarchy starts with the stronger guarantee that each voter should be able to control their proportion of the budget.

We then turn to the goal of achieving best-of-both-worlds fairness in PB. In Section 5, we investigate the special case of PB with binary utilities. We show that our strongest ex-ante fairness notion cannot be guaranteed in tandem with any of our ex-post fairness notions (Section 5.1), notable since this is not the case in committee voting. We then give a strong, positive result using an exponential-time algorithm (Section 5.2) and a slightly weaker positive result using a polynomial-time algorithm (Section 5.3). Lastly, in Section 6, we show that in the general model with cardinal utilities, ex-ante and ex-post fairness are not compatible, even for the weakest pair of axioms and even in the restricted case where projects are of unit cost. Our best-of-both-worlds fairness results are summarized in Figure 1.

1.2 Related Work

There is a fast growing body of work investigating proportionally representative outcomes in indivisible PB (e.g., Aziz, Lee, and Talmon 2018; Peters, Pierczyński, and Skowron 2021; Los, Christoff, and Grossi 2022; Brill et al. 2023; Rey and Maly 2023). Since PB generalizes committee voting (Lackner and Skowron 2023), work on proportionality in PB often extends axioms and algorithms from the literature on proportional representation in committee voting (Aziz et al. 2017; Sánchez-Fernández et al. 2017; Elkind et al. 2022; Brill and Peters 2023). It is to this literature that we are adding the tool of randomization.

Best-of-both-worlds fairness has been examined in several papers in the context of resource allocation (Babaioff, Ezra, and Feige 2022; Aziz et al. 2023a; Aziz, Ganguly, and Micha 2023; Hoefer, Schmalhofer, and Varricchio 2023; Feldman et al. 2023). A couple of the earlier papers in this line of work are by Aziz et al. (2023a) and Aziz (2019). In recent work, Aziz et al. (2023b) initiated a best-of-both-worlds fairness perspective for the social choice setting of approval-based committee voting. We extend this approach to multiple PB settings, all of which generalize approval-based committee voting.

Indeed, we are the first to study lotteries in PB. In the single-winner voting setting, Bogomolnaia, Moulin, and Stong (2005) computed “mixtures” of public outcomes which satisfy natural axioms seeking to give groups of agents their fair share, the same set of axioms which have inspired the ex-ante axioms in this work.

Lottery implementation techniques have been studied in other social choice settings. For example, the dependent randomized rounding technique of Gandhi et al. (2006) has been employed to compute randomized outcomes which implement desirable fractional outcomes in various settings such as fair division Akbarpour and Nikzad (2020) and committee voting (Cheng et al. 2020). In the context of apportionment, Aziz et al. (2019) and Gözl, Peters, and Procaccia (2022) have created new rounding techniques to facilitate randomization when distributing legislative seats.

There are also several papers which study divisible PB, wherein projects can be funded fractionally. Fain, Goel, and Munagala (2016) showed that the Nash solution in this setting is in the *core*, a property which captures proportional representation. In this paper and most work in this area, projects do not have fixed costs, and thus any distribution of budget amongst projects is feasible, which contrasts significantly with our setting. Goel et al. (2019) incorporated project costs in the divisible PB setting and study strategic concerns. However, their outcomes still allow for fractional project funding. In this paper, we investigate the question of how such an outcome can be converted into a lottery over outcomes in which funding decisions are binary.

2 Preliminaries

For any positive integer $t \in \mathbb{N}$, let $[t] := \{1, 2, \dots, t\}$. A participatory budgeting (PB) *instance* is represented as a tuple $I = \langle N, C, \text{cost}, B, (u_i)_{i \in [n]} \rangle$, where:

- $N = [n]$ and $C = [m]$ are the set of *voters* and *projects*, respectively.
- $\text{cost}: C \rightarrow \mathbb{R}_{\geq 0}$ is the *cost function*, associating each project $c \in C$ with its cost that needs to be paid if c is selected. For any subset of projects $T \subseteq C$, denote by $\text{cost}(T) := \sum_{c \in T} \text{cost}(c)$ the total cost of T .

We say projects have *unit costs* if $\text{cost}(c) = 1$ for all $c \in C$, and refer to the setting as *unit-cost PB*.

- $B \in \mathbb{R}_{\geq 0}$ is the *budget limit*. We assume without loss of generality that $\text{cost}(c) \leq B \forall c \in C$ and $\text{cost}(C) \geq B$.
- For each $i \in N$, *utility function* $u_i: C \rightarrow \mathbb{R}_{\geq 0}$ expresses how voter i values each project. This most general formulation is referred to as *general PB*. We call the set of

projects for which voter i has non-zero utility their *approval set* and denote it by A_i .

We say voters have *binary utilities* if $u_i(c) \in \{0, 1\}$ for all $i \in N, c \in C$, and refer to the setting as *PB with binary utilities*.

Integral Outcomes An *integral outcome* (or simply *outcome*) is a set of projects $W \subseteq C$, and it is said to be *feasible* if $\text{cost}(W) \leq B$. We assume *additive utilities*, meaning that given a subset of projects $T \subseteq C$, $u_i(T) := \sum_{c \in T} u_i(c)$.

Fractional Outcomes A *fractional outcome* is an m -dimensional vector $\vec{p} \in [0, 1]^m$, where the component $p_c \in [0, 1]$ represents the fraction of project c funded. Given an integral outcome W , for notational convenience, let $\vec{1}_W \in \{0, 1\}^m$ be the binary vector whose j 'th component is 1 if and only if $j \in W$. Let $\text{cost}(\vec{p}) := \sum_{c \in C} p_c \cdot \text{cost}(c)$. A fractional outcome \vec{p} is said to be *feasible* if $\text{cost}(\vec{p}) = B$. Given budget B' , denote by $\mathcal{X}(B')$ the space of feasible fractional outcomes. For simplicity, we use \mathcal{X} to denote the space of feasible fractional outcomes given budget B . Given a fractional outcome \vec{p} , voter i 's utility is denoted by $u_i(\vec{p}) := \sum_{c \in C} p_c \cdot u_i(c)$.

Lotteries and Implementation A *lottery* is a probability distribution over integral outcomes. Formally, a lottery is specified by a set of $s \in \mathbb{N}$ tuples $\{(\lambda_j, W_j)\}_{j \in [s]}$, where $\lambda_j \in [0, 1]$, $\sum_j \lambda_j = 1$, and for every $j \in [s]$, the integral outcome $W_j \subseteq C$ is selected with probability λ_j . A lottery $\{(\lambda_j, W_j)\}_{j \in [s]}$ is called an *implementation* of (or, interchangeably, *implements*) a fractional outcome \vec{p} if $\vec{p} = \sum_{j \in [s]} \lambda_j \cdot \vec{1}_{W_j}$. In this paper, we only consider lotteries which implement feasible fractional outcomes. We say a lottery satisfies a property *ex ante* (resp., *ex post*) if the fractional outcome it implements (resp., every integral outcome in its support) satisfies the property.

This paper concerns the problem of designing (randomized) PB rules (or, interchangeably, algorithms) that simultaneously achieve desirable properties both *ex ante* and *ex post*. Our algorithms do not explicitly output the desired lotteries (which in principle can be exponential in size), but instead sample integral outcomes from them.

3 Implementing Fractional Outcomes in PB

In this section, we study how to implement a fractional outcome in the context of participatory budgeting. Decomposing a fractional outcome into a distribution over integral outcomes introduces novel challenges in the presence of costs.

Recall that, given marginal probabilities p_i of selecting project i , implementing a fractional outcome entails computing a probability distribution $\Delta = \{(\lambda_j, W_j)\}_{j \in [s]}$ over integral outcomes that realizes the prescribed marginal probabilities. As a result, any implementation of a feasible fractional outcome has the property that $\mathbb{E}_{W \sim \Delta}[\text{cost}(W)] = B$. It is now easy to see that unless each integral outcome in the support of the lottery has cost equal to B (not possible in general), there must exist an integral outcome in the support of the lottery that exceeds the budget.

The aforementioned issue raises the natural question of whether it is possible to implement a fractional outcome while bounding the *ex-post* budget violations. This is especially important in participatory budgeting since if the *ex-post* budget constraint is exceedingly violated, such an outcome is unlikely to be implemented in practice. We thus formalize an axiom which guarantees that an integral outcome is approximately within budget.

Definition 3.1 (BB1). An integral outcome W is said to be *budget balanced up to one project (BB1)* if either

- $\text{cost}(W) \leq B$ and there exists some project $c \in C \setminus W$ such that $\text{cost}(W \cup \{c\}) \geq B$, or
- $\text{cost}(W) \geq B$ and there exists some project $c \in W$ such that $\text{cost}(W \setminus \{c\}) \leq B$.

We now show, perhaps surprisingly, that *any* feasible fractional outcome \vec{p} can be implemented by a lottery over integral outcomes, each of which satisfies BB1.

Theorem 3.2. *For any feasible fractional outcome \vec{p} , there exists a random process running in polynomial time, that defines random variables $P_i \in \{0, 1\}$ for all $i \in C$ such that the following properties hold:*

- (P1) $\mathbb{E}[P_i] = p_i$ for each $i \in C$;
- (P2) Random integral committee $W = \{i \in C \mid P_i = 1\}$ satisfies BB1 with probability 1.

Note that there is a lottery associated with the random process described in Theorem 3.2, but we only return an integral outcome sampled from this underlying lottery as it may be exponential in size. By (P1), the underlying lottery implements \vec{p} , and by (P2), it satisfies *ex-post* BB1. We remark that Theorem 3.2 is a consequence of applying the dependent rounding scheme of Gandhi et al. (2006) to our setting.

One might wonder whether we can further strengthen the *ex-post* budget feasibility axiom BB1. A natural strengthening is the following:

Definition 3.3 (BFx). An integral outcome W is said to be *budget feasible up to any project (BFx)* if for all $c \in W$, $\text{cost}(W \setminus \{c\}) \leq B$.

We, however, show that not all fractional outcomes may be implemented by *ex-post* BFx lotteries.

Proposition 3.4. *There exists some fractional outcome \vec{p} that cannot be implemented by a lottery that is *ex-post* BFx.*

Handling Hard Constraints We note that, in addition to being well-suited to scenarios in which budget constraints have some flexibility, the implementation techniques introduced in this section are also relevant to settings with hard *ex-post* budget constraints. To see this, consider a problem wherein every *ex-post* outcome is restricted to having a cost of at most B . If we now apply Theorem 3.2 to any fractional outcome that spends $B' = B - \max_{g \in C} \text{cost}(g)$, the resulting implementation has the property that every integral outcome in its support has cost at most B .

4 Ex-ante Fairness Concepts

We now present *ex-ante* axioms for the PB setting inspired by the fair share axioms first introduced by Bogomolnaia,

Moulin, and Stong (2005). The fair share axioms were recently extended by Aziz et al. (2023b) to the committee voting setting, the special case of our own model in which each project is of unit cost and voters have binary utilities. In that work, Aziz et al. (2023b) highlighted two alternative interpretations of the idea behind fair share, and defined axiom hierarchies for each of these interpretations. These two interpretations are given as follows:

- (a) Fair share: each voter is given $1/n$ probability to choose their favourite outcome, or
- (b) Strong fair share: each voter can select $1/n$ of the outcome.

In this section, we generalize the (strong) fair share axiom hierarchies to the general PB setting, with the intention of formulating axioms which (i) collapse to those defined by Aziz et al. (2023b) in approval-based committee voting, and (ii) reflect their respective interpretations as detailed above. Each of the axioms given in this section provides lower bounds on utilities derived from fractional outcomes. We note that these utilities can also be interpreted as expected utilities from implementations of these fractional outcomes. In particular, if $\Delta = (\lambda_j, W_j)_{j \in [r]}$ is an implementation of a fractional outcome \vec{p} , then $\mathbb{E}_{W \sim \Delta} [u_i(W)] = u_i(\vec{p})$.

The first axiom in our hierarchy, *individual fair share (IFS)*, guarantees each agent receives utility which is at least a $1/n$ -fraction of the utility they receive from their favourite fractional outcome.

Definition 4.1 (IFS). A fractional outcome \vec{p} is said to provide *IFS* if for each $i \in N$,

$$u_i(\vec{p}) \geq \frac{1}{n} \cdot \max_{\vec{t} \in \mathcal{X}} u_i(\vec{t}).$$

Note that the quantity expressed by the max-operator is the utility-maximizing fractional outcome for the agent i , and hence it is immediately clear that Definition 4.1 follows interpretation (a). In general, this can be computed by selecting projects in order of descending utility per cost. For binary utilities, this means selecting approved projects in order of ascending cost. We can already observe that, in our setting, IFS seems quite a bit more demanding than its approval-based committee voting counterpart, in which a project/candidate can only take on a utility per cost value of one or zero. In contrast, in the PB setting, each voter-project pair could result in a unique utility per cost value.

Definition 4.2 (Strong IFS). A fractional outcome \vec{p} is said to provide *Strong IFS* if for each $i \in N$,

$$u_i(\vec{p}) \geq \max_{\vec{t} \in \mathcal{X}(\frac{B}{n})} u_i(\vec{t}).$$

For Strong IFS, keeping with interpretation (b) above, an agent's utility lower bound is given by the optimal utility they could achieve if given their proportion of the budget. Next, we strengthen IFS to *unanimous fair share (UFS)*, which strengthens the fair share utility guarantee for groups of voters with identical preferences.

Definition 4.3 (UFS). A fractional outcome \vec{p} is said to provide *UFS* if for any $S \subseteq N$ where $u_i = u_j$ for any $i, j \in S$,

then the following holds for each $i \in S$:

$$u_i(\vec{p}) \geq \frac{|S|}{n} \cdot \max_{\vec{t} \in \mathcal{X}} u_i(\vec{t}).$$

Definition 4.4 (Strong UFS). A fractional outcome \vec{p} is said to provide *Strong UFS* if for any $S \subseteq N$ where $u_i = u_j$ for any $i, j \in S$, the following holds for each $i \in S$:

$$u_i(\vec{p}) \geq \max_{\vec{t} \in \mathcal{X}(|S| \cdot \frac{B}{n})} u_i(\vec{t}). \quad (1)$$

As its name suggests, Strong UFS implies UFS (and hence Strong IFS implies IFS).¹ While (Strong) UFS gives a utility guarantee to groups of voters with identical preferences, our next axiom — *group fair share (GFS)* — extends a non-trivial representation guarantee to all groups of voters.

Definition 4.5 (GFS). A fractional outcome \vec{p} is said to provide *GFS* if the following holds for any $S \subseteq N$:

$$\sum_{j \in C} (p_j \cdot \max_{i \in S} u_{ij}) \geq \frac{1}{n} \cdot \sum_{i \in S} \max_{\vec{t} \in \mathcal{X}} u_i(\vec{t}).$$

In committee voting, the LHS of GFS is simply the probability allocated to candidates in the union of the group of voters' approval sets. Thus, while it is clear that our definition collapses to the GFS of Aziz et al. (2023b) in committee voting, ours is not the only formulation of the LHS of GFS which does so. For example, instead of taking the maximum utility for each project j over all agents in S , we could have instead taken the average or median (or minimum) utility over all agents in S with non-zero utility for project j . Of these options, our formulation results in the weakest definition. Since, as we will see, this definition of GFS is not compatible with any of the ex-post fairness notions we consider in any PB setting, each of the results considering GFS in this paper would also hold for any stronger notions of GFS.

Fractional Random Dictator We now extend the well-known Random Dictator algorithm (Bogomolnaia, Moulin, and Stong 2005) to the computation of fractional PB outcomes. The high-level idea of the algorithm is to compute the fractional outcome resulting from giving each agent $1/n$ probability to select their own favourite fractional outcome. For an agent $i \in N$, let X_i be the maximal set of projects which can be funded fully in order of maximum utility per cost and let g_i be the project with highest utility per cost in the remaining set of projects. Also, denote the indicator function by $\mathbb{1}_{\{\cdot\}}$. For each $j \in C$, the *Fractional Random Dictator* algorithm outputs the fractional outcome \vec{p} defined as follows:

$$p_j = \frac{1}{n} \sum_{i \in N} \mathbb{1}_{\{j \in X_i\}} + \mathbb{1}_{\{j = g_i\}} \frac{B - \text{cost}(X_i)}{\text{cost}(j)}.$$

Theorem 4.6. *The Fractional Random Dictator algorithm computes an ex-ante GFS fractional outcome.*

¹To see this, let $\vec{q} = \arg \max_{\vec{t} \in \mathcal{X}} u_i(\vec{t})$. Now simply note that $\frac{|S|}{n} \vec{q} \in \mathcal{X}(|S| \cdot \frac{B}{n})$.

5 BOBW Fairness in PB with Binary Utilities

In this section, we consider the setting of PB with binary utilities, in which the voters have binary utilities (while the projects have arbitrary costs). Our main focus is to investigate whether the ex-ante fair share notions defined in Section 4 can be achieved simultaneously with ex-post fairness properties like *justified representation (JR)*, *extended justified representation (EJR)* and *full justified representation (FJR)* (Peters, Pierczyński, and Skowron 2021).

Definition 5.1 (*T-cohesive group for PB with binary utilities*). A group of voters $S \subseteq N$ is said to be *T-cohesive* for $T \subseteq C$ if $|S| \cdot \frac{B}{n} \geq \text{cost}(T)$ and $T \subseteq \bigcap_{i \in S} A_i$.

Definition 5.2 (*JR & EJR for PB with binary utilities*). An outcome W is said to satisfy

- *justified representation (JR)* if for each $j \in C$ and every $\{j\}$ -cohesive group of voters $S \subseteq N$, it holds that $u_i(W) = |A_i \cap W| \geq 1$ for some $i \in S$; and
- *extended justified representation (EJR)* if for each $T \subseteq C$ and every *T-cohesive* group of voters $S \subseteq N$, it holds that $u_i(W) = |A_i \cap W| \geq |T|$ for some $i \in S$.

By definition, EJR implies JR. We next introduce a notion stronger than EJR.

Definition 5.3 (*FJR for PB with binary utilities*). Given a positive integer β and a set of projects $T \subseteq C$, a group of voters $S \subseteq N$ is said to be *weakly (β, T) -cohesive* if $|S| \cdot \frac{B}{n} \geq \text{cost}(T)$ and $|A_i \cap T| \geq \beta$ for all $i \in S$.

An outcome W is said to satisfy *full justified representation (FJR)* if for every weakly (β, T) -cohesive group of voters $S \subseteq N$, it holds that $|A_i \cap W| \geq \beta$ for some $i \in S$.

The remainder of this section is organized as follows:

- In Section 5.1, we show that it is impossible to simultaneously achieve ex-ante GFS and ex-post JR.
- In Section 5.2, we show constructively that ex-ante Strong UFS and ex-post FJR are compatible, though our randomized algorithm is not polynomial time.
- In Section 5.3, we devise a polynomial-time randomized algorithm which simultaneously achieves ex-ante Strong UFS and ex-post EJR.

5.1 Impossibility: Ex-ante GFS + Ex-post JR

Our first main result in PB with binary utilities states that it is impossible to simultaneously achieve ex-ante GFS and ex-post JR. Note that in the more restricted setting with unit-cost projects (i.e., approval-based committee voting), Aziz et al. (2023b) showed that ex-ante GFS is compatible even with ex-post EJR. Our impossibility result demonstrates a clear and strong separation between PB with binary utilities and approval-based committee voting.

Theorem 5.4. *In PB with binary utilities, ex-ante GFS and ex-post JR are incompatible.*

Proof. Consider an instance with $n \geq 6$ and the following approval sets and project costs:

- each voter $i \in N$ approves $A_i = \{g^*, a_i, b_i, c_i\}$ with $\text{cost}(g^*) = \frac{B}{2}$ and $\text{cost}(a_i) = \text{cost}(b_i) = \text{cost}(c_i) = \frac{B}{2} - \varepsilon$, where $\varepsilon < \frac{B}{2} - \frac{2B}{n}$;
- note that g^* is the common project approved by every voter, and for any pair of voters $i \neq j$, $a_i, b_i, c_i \notin A_j$.

We establish the incompatibility using this instance by showing that *any* feasible fractional outcome satisfying GFS cannot be implemented by *any* lottery that is ex-post JR, even without imposing BB1. Suppose, for the sake of contradiction, that $\{(\lambda_j, W_j)\}_{j \in [s]}$ is an ex-post JR lottery implementing GFS fractional outcome \vec{p} .

We first point out that some integral outcome in the lottery includes g^* , and hence $p_{g^*} > 0$.

Claim 5.5. *There exists an outcome W_j such that $g^* \in W_j$.*

Feasibility of \vec{p} means that $B = \sum_{c \in C} p_c \cdot \text{cost}(c) = \sum_{c \in C \setminus \{g^*\}} p_c \cdot \left(\frac{B}{2} - \varepsilon\right) + p_{g^*} \cdot \frac{B}{2}$. Thus,

$$\sum_{c \in C} p_c = \frac{B - B/2 \cdot p_{g^*}}{B/2 - \varepsilon} + p_{g^*} = \frac{B - \varepsilon \cdot p_{g^*}}{B/2 - \varepsilon}.$$

Since \vec{p} satisfies GFS with respect to voters N , we thus have

$$\begin{aligned} \frac{B - \varepsilon \cdot p_{g^*}}{B/2 - \varepsilon} &= \sum_{c \in C} p_c \\ &\geq \frac{1}{n} \cdot \sum_{i \in N} \max_{\vec{t} \in \mathcal{X}} u_i(\vec{t}) \\ &= \frac{1}{n} \cdot \sum_{i \in N} \frac{B}{B/2 - \varepsilon} = \frac{B}{B/2 - \varepsilon}, \end{aligned}$$

a contradiction because $p_{g^*} > 0$. \square

As demonstrated by Aziz et al. (2023b) in the approval-based committee voting, there is no logical dependence between GFS and Strong UFS. It is thus unclear whether ex-ante Strong UFS can be compatible with any ex-post fairness properties. We answer the question in the affirmative below.

5.2 Ex-ante Strong UFS + Ex-post FJR

We now show that if we only focus on giving ex-ante fair share guarantees to *unanimous* (instead of any) groups, ex-ante Strong UFS is compatible even with ex-post FJR.

Theorem 5.6. *In PB with binary utilities, Algorithm 1 computes an integral outcome sampled from a lottery that is ex-ante Strong UFS, ex-post BB1 and ex-post FJR.*

The Algorithm: BW-GCR-PB Our algorithm, whose pseudocode can be found in Algorithm 1, starts by feeding the given PB instance into the *Greedy Cohesive Rule (GCR)* of Peters, Pierczyński, and Skowron (2021) and obtains an FJR outcome. More specifically, GCR begins by making all voters as *active* and initializing $W = \emptyset$. In each step, GCR searches for a set of voters $N' \subseteq N$ who are all active and a set of projects $T \subseteq C \setminus W$ such that N' is weakly (β, T) -cohesive, breaking ties in favour of larger β , next smaller $\text{cost}(T)$, and then larger $|N'|$. GCR then includes projects T to W and labels voters N' as inactive. If, at any

Algorithm 1: BW-GCR-PB: Strong UFS and FJR

Input: Voters $N = [n]$, projects $C = [m]$, cost function cost , budget B , and utilities $(\bar{u}_i)_{i \in N}$.

- 1 $W_{\text{GCR}} \leftarrow \text{GCR}(N, C, \text{cost}, B, (u_i)_{i \in N})$
- 2 $\vec{p} \leftarrow \vec{1}_{W_{\text{GCR}}}$
- 3 $\tilde{N} \leftarrow \emptyset$
- 4 $b_i \leftarrow 0$ for all $i \in N$
- 5 Let $\{N^1, \dots, N^\eta\}$ be the unanimous groups of N .
- 6 **foreach** $z \in [\eta]$ **do**
- 7 **if** $|A_{N^z} \cap W_{\text{GCR}}| = |G_{N^z}|$ **then**
- 8 $\tilde{N} \leftarrow \tilde{N} \cup N^z$
- 9 $b_i \leftarrow \frac{B}{n} - \frac{1}{|N^z|} \cdot \text{cost}(G_{N^z})$ for all $i \in N^z$
- 10 Let voters N^z spend their total budget $|N^z| \cdot \frac{B}{n} - \text{cost}(G_{N^z})$ on project $c \in A_{N^z}$ with the smallest cost, provided the updated $p_c \leq 1$.
- 11 Increase \vec{p} arbitrarily such that for all $c \in C$, $p_c \leq 1$ and $\text{cost}(\vec{p}) = B$.
- 12 Obtain an outcome W sampled from the lottery implementing \vec{p} by applying Theorem 3.2.
- 13 **return** \vec{p} and W

step, no weakly (β, T) -cohesive group exists for any positive integer β , then GCR returns W and terminates. Denote by r the number of steps GCR executes before terminating. For each $j \in [r]$, we refer to β_j , T_j and N_j as the values of β , T and N' for the weakly cohesive group selected in the j -th step of GCR. Denote by $W_{\text{GCR}} := \bigcup_{j \in [r]} T_j$ and initialize \vec{p} as $\vec{1}_{W_{\text{GCR}}}$.

Algorithm 1 next loops over each unanimous group one by one and set budgets for the voters. For ease of expression, let $\{N^1, N^2, \dots, N^\eta\}$ be the partition of the (maximal) unanimous groups of voters N , i.e., for each $z \in [\eta]$, voters N^z are unanimous and for any $i \in N^z$ and $i' \in N^{z'}$ with $z \neq z'$, $A_i \neq A_{i'}$. Fix any $z \in [\eta]$. If voter $i \in N^z$ gets utility exactly $|G_{N^z}|$ from W_{GCR} , i.e., the if-condition holds, we set budget $b_i := \frac{B}{n} - \frac{1}{|N^z|} \cdot \text{cost}(G_{N^z})$. The unanimous group of voters N^z then spend their total budget of $|N^z| \cdot \frac{B}{n} - \text{cost}(G_{N^z})$ on project $c \in A_{N^z}$ with the smallest cost, provided the updated $p_c \leq 1$. We will show shortly that the budget set-up is valid (Lemma 5.7) and would make each unanimous group satisfy Strong UFS (Lemma 5.9). Line 11 then increases \vec{p} in an arbitrary way so that \vec{p} is feasible, that is, for all $c \in C$, $p_c \leq 1$ and $\text{cost}(\vec{p}) = B$.

Finally, given the feasible fractional outcome \vec{p} , we apply the randomized rounding scheme (Theorem 3.2) and sample an outcome from the lottery implementing \vec{p} .

The Analysis of BW-GCR-PB Below, we highlight two important components in the proof of Theorem 5.6. We start with the fact that the total budgets we give the voters in line 9 is upper bounded by the leftover budget limit after selecting W_{GCR} .

Lemma 5.7. $\sum_{i \in N} b_i \leq B - \text{cost}(W_{\text{GCR}})$.

Proof. For ease of exposition, in this proof, we re-order the r weakly cohesive groups encountered by GCR in line 1. Let $r' \in \{0, 1, 2, \dots, r\}$ be an index such that for all $j \in [r']$, $N_j \cap \tilde{N} \neq \emptyset$, i.e., there exists a unanimous group in N_j such that the if-condition holds. For each $j \in [r']$, let $\{N_j^1, N_j^2, \dots, N_j^{\eta_j}\}$ be the partition of the (maximal) unanimous groups of voters N_j . We also assume without loss of generality that the first $\eta_j' \in \{0, 1, 2, \dots, \eta_j\}$ unanimous groups are the ones such that the if-condition holds. Note that $\eta_j' \geq 1$ for all $j \in [r']$. Denote by $N_{\text{GCR}} := \bigcup_{j \in [r']} N_j$ the set of inactive voters due to GCR. Note that for all $i \in N$, $b_i \leq \frac{B}{n}$. We will also make use of the following claim:

Claim 5.8. $\forall j \in [r'], z \in [\eta_j'], \text{cost}(T_j) \leq \text{cost}(G_{N_j^z})$.

We are now ready to establish the statement:

$$\begin{aligned} \sum_{i \in N} b_i &= \sum_{i \in N_{\text{GCR}}} b_i + \sum_{i \in N \setminus N_{\text{GCR}}} b_i \\ &= \sum_{j=1}^{r'} \sum_{z=1}^{\eta_j'} \left(|N_j^z| \cdot \frac{B}{n} - \text{cost}(G_{N_j^z}) \right) + \sum_{i \in N \setminus N_{\text{GCR}}} b_i \\ &\leq \sum_{j=1}^{r'} \left(|N_j| \cdot \frac{B}{n} - \sum_{z=1}^{\eta_j'} \text{cost}(G_{N_j^z}) \right) + \sum_{i \in N \setminus N_{\text{GCR}}} b_i \\ &\leq \sum_{j=1}^{r'} \left(|N_j| \cdot \frac{B}{n} - \text{cost}(T_j) \right) + \sum_{i \in N \setminus N_{\text{GCR}}} b_i \\ &\leq \sum_{j=1}^r \left(|N_j| \cdot \frac{B}{n} - \text{cost}(T_j) \right) + \sum_{i \in N \setminus N_{\text{GCR}}} b_i \\ &\leq |N_{\text{GCR}}| \cdot \frac{B}{n} - \text{cost}(W_{\text{GCR}}) + |N \setminus N_{\text{GCR}}| \cdot \frac{B}{n} \\ &= B - \text{cost}(W_{\text{GCR}}), \end{aligned}$$

where the fourth transition is due to Claim 5.8 and $\eta_j' \geq 1$, and the fifth transition is due to weak cohesiveness. \square

Our next result establishes that \vec{p} satisfies Strong UFS.

Lemma 5.9. Algorithm 1 outputs a fractional outcome \vec{p} that satisfies Strong UFS.

We provide here some observations and intuition for Lemma 5.9. Let us first establish connections between unanimous groups considered by Strong UFS and cohesive groups considered by EJR.² Fix any unanimous group $S \subseteq N$. Denote by A_S the approval set of the unanimous group S (i.e., for all $i \in S$, $A_i = A_S$). Let us rename the projects in A_S in a non-decreasing order of cost with arbitrary tie-breaking, i.e., $\text{cost}(g_1) \leq \text{cost}(g_2) \leq \dots \leq \text{cost}(g_{|A_S|})$. Denote by $G_S := \{g_1, g_2, \dots, g_{\kappa_S}\}$ the maximal set of projects such that $\text{cost}(G_S) \leq |S| \cdot \frac{B}{n}$. Put differently, if $A_S \setminus G_S \neq \emptyset$, $\text{cost}(G_S \cup \{g_{\kappa_S+1}\}) > |S| \cdot \frac{B}{n}$. Since $|S| \cdot \frac{B}{n} \geq \text{cost}(G_S)$ and $G_S \subseteq A_S$, we have the following observation:

²Since FJR implies EJR, our discussion is carried over to weakly cohesive groups considered by FJR.

Observation. *The unanimous group S is G_S -cohesive.*

It follows that given an EJR (or FJR) outcome W , for all $i \in S$, $|A_i \cap W| \geq |G_S|$. We thus conclude:

Claim 5.10. *Given an instance of PB with binary utilities, fix any unanimous group of voters $S \subseteq N$ and any EJR (or FJR) outcome W , then, for all $i \in S$, $|A_i \cap W| \geq |G_S|$.*

We now provide an alternative description of Equation (1) in the definition of Strong UFS, and focus on the case where $A_S \setminus G_S \neq \emptyset$, that is, project g_{κ_S+1} exists. According to the RHS of Equation (1), voters S is endowed with a budget of $|S| \cdot \frac{B}{n}$ to select a fractional outcome. An optimal fractional outcome can be achieved by fully funding G_S and next funding a δ_S fraction of project g_{κ_S+1} , where $\delta_S := \frac{|S| \cdot \frac{B}{n} - \text{cost}(G_S)}{\text{cost}(g_{\kappa_S+1})} < 1$. Thus, in this case, we can rewrite the RHS of Equation (1):

$$\max_{\vec{t} \in \mathcal{X}(|S| \cdot \frac{B}{n})} u_i(\vec{t}) = |G_S| + \delta_S. \quad (2)$$

Fix any $i \in S$ and an EJR (or FJR) outcome W . If $|A_i \cap W| > |G_S|$, i.e., $|A_i \cap W| \geq |G_S| + 1$, voter i is already satisfied with respect to Strong UFS under fractional outcome $\vec{1}_W$. In the case where $|A_i \cap W| = |G_S|$, we show in the proof of Lemma 5.9 that $\vec{1}_W$ is extended to a feasible fractional outcome in a way that voter i gains an additional utility of at least δ_S and meets the Strong UFS guarantee.

5.3 Ex-ante Strong UFS + Ex-post EJR (in Polynomial Time)

Despite providing strong ex-ante and ex-post fairness guarantees, BW-GCR-PB does not run in polynomial time. We present here a polynomial-time algorithm that is ex-ante Strong UFS, at the cost of weakening ex-post fairness guarantee to EJR.

Theorem 5.11. *In PB with binary utilities, Algorithm 2 computes an integral outcome sampled from a lottery that is ex-ante Strong UFS, ex-post BBI, and ex-post EJR in polynomial time.*

At a high level, our algorithm BW-MES-PB (Algorithm 2) gives each voter an initial budget of B/n and uses the *Method of Equal Shares (MES)* of Peters, Pierczyński, and Skowron (2021) as a subroutine to obtain an EJR outcome W_{MES} . Denote by y_{ic} the payment each voter $i \in N$ made for each project $c \in C$. A key step here is to show that for each unanimous group $N^z \subseteq N$, the remaining budget of the group $\sum_{i \in N^z} (\frac{B}{n} - \sum_{c \in W_{\text{MES}}} y_{ic})$ is at least $|N^z| \cdot \frac{B}{n} - \text{cost}(G_{N^z})$. As a result, the group together can use their remaining budget to fund the project with the smallest cost and be satisfied with respect to Strong UFS. Finally, as MES and all other steps run in polynomial time, we conclude that Algorithm 2 is a polynomial-time algorithm.

6 BOBW Fairness in General PB

We now move on to the setting of general PB, in which we show a strong impossibility that ex-ante IFS and ex-post JR are not compatible, even in the unit-cost PB setting. This impossibility is striking as, even in the general

Algorithm 2: BW-MES-PB: Strong UFS and EJR

Input: Voters $N = [n]$, projects $C = [m]$, budget B , cost function cost , and utilities $(\vec{u}_i)_{i \in N}$.

- 1 $W_{\text{MES}} \leftarrow \text{MES}(N, C, \text{cost}, B, (u_i)_{i \in N})$
- 2 $\vec{p} \leftarrow \vec{1}_{W_{\text{MES}}}$
- 3 Let y_{ij} for each $i \in N$ and $j \in W_{\text{MES}}$ be the amount voter i spent on project j during MES.
- 4 $b_i \leftarrow \frac{B}{n} - \sum_{j \in W_{\text{MES}}} y_{ij}$ for all $i \in N$, which is the remaining budget of voter i after MES
- 5 $N' \leftarrow \{i \in N \mid A_i \setminus W_{\text{MES}} \neq \emptyset\}$
- 6 **foreach** $i \in N'$ **do**
- 7 Let $\kappa_i \in \arg \min_{c \in A_i \setminus W_{\text{MES}}} \text{cost}(c)$
- 8 $y_{i\kappa_i} \leftarrow b_i$
- 9 **foreach** $i \in N \setminus N'$ **do**
- 10 Voter i spends b_i arbitrarily provided $\sum_{i \in N} y_{ij} \leq \text{cost}(j)$ for all $j \in C$.
- 11 **foreach** $j \in C$ **do** $p_j \leftarrow \frac{\sum_{i \in N} y_{ij}}{\text{cost}(j)}$.
- 12 Obtain an outcome W sampled from the lottery implementing \vec{p} by applying Theorem 3.2.
- 13 **return** \vec{p} and W

setting, the much stronger properties of ex-ante GFS and ex-post FJR are independently achievable via Fractional Random Dictator (Theorem 4.6) and GCR (Peters, Pierczyński, and Skowron 2021), respectively.

We first define the notion of *justified representation* (Peters, Pierczyński, and Skowron 2021).

Definition 6.1 (JR). In the general PB setting, a group of voters S is (α, T) -cohesive, where $\alpha : C \rightarrow [0, 1]$ and $T \subseteq C$, if $|S|/n \geq \text{cost}(T)/B$ and if $u_{ij} \geq \alpha(j)$ for all $i \in S$ and $j \in T$. An integral outcome W satisfies JR, if for each $\alpha : C \rightarrow [0, 1]$, $j \in C$, and each $(\alpha, \{j\})$ -cohesive group of agents, there exists an agent $i \in S$ such that $u_i(\vec{1}_W) \geq \alpha(j)$.

We show that ex-ante IFS and ex-post JR are incompatible in the unit-cost setting with cardinal utilities. Intuitively, in situations where voters have high utilities for distinct projects, the outcomes that guarantee the highest expected utility may not include a project which gives every voter non-zero utility.

Theorem 6.2. *In unit-cost PB, ex-ante IFS and ex-post JR are incompatible.*

7 Conclusion

In this paper, we initiated the study of PB lotteries and used this approach to study best-of-both-worlds fairness in PB. We provided a complete set of results for two natural restrictions of PB with cardinal utilities. Specifically, we gave algorithms which compute a lottery that guarantees each voter certain expected utility while maintaining the strongest indivisible PB fairness notions ex post. While we focused on fairness, it is an interesting future direction to seek best-of-both-worlds results for other desiderata, such as efficiency.

Acknowledgments

This work was partially supported by the NSF-CSIRO grant on “Fair Sequential Collective Decision-Making” and the ARC Laureate Project FL200100204 on “Trustworthy AI”.

References

- Akbarpour, M.; and Nikzad, A. 2020. Approximate Random Allocation Mechanisms. *The Review of Economic Studies*. Rdz066.
- Aziz, H. 2019. A Probabilistic Approach to Voting, Allocation, Matching, and Coalition Formation. In Laslier, J.-F.; Moulin, H.; Sanver, R.; and Zwicker, W. S., eds., *The Future of Economic Design*. Springer Cham.
- Aziz, H.; Brill, M.; Conitzer, V.; Elkind, E.; Freeman, R.; and Walsh, T. 2017. Justified Representation in Approval-based Committee Voting. *Social Choice and Welfare*, 48(2): 461–485.
- Aziz, H.; Freeman, R.; Shah, N.; and Vaish, R. 2023a. Best of Both Worlds: Ex Ante and Ex Post Fairness in Resource Allocation. *Operations Research*. Forthcoming.
- Aziz, H.; Ganguly, A.; and Micha, E. 2023. Best of Both Worlds Fairness under Entitlements. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 941–948.
- Aziz, H.; Lee, B. E.; and Talmon, N. 2018. Proportionally Representative Participatory Budgeting: Axioms and Algorithms. In *Proceedings of the 17th International Conference on Autonomous Agents and Multi Agent Systems (AAMAS)*, 23–31.
- Aziz, H.; Lev, O.; Mattei, N.; Rosenchein, J. S.; and Walsh, T. 2019. Strategyproof Peer Selection Using Randomization, Partitioning, and Apportionment. *Artificial Intelligence*, 275: 295–309.
- Aziz, H.; Lu, X.; Suzuki, M.; Vollen, J.; and Walsh, T. 2023b. Best-of-Both-Worlds Fairness in Committee Voting. In *Proceedings of the 19th Conference on Web and Internet Economics (WINE)*. Forthcoming. Extended version available at <https://arxiv.org/abs/2303.03642>.
- Aziz, H.; and Shah, N. 2021. Participatory Budgeting: Models and Approaches. In Rudas, T.; and Péli, G., eds., *Pathways Between Social Science and Computational Social Science: Theories, Methods, and Interpretations*, 215–236. Springer International Publishing.
- Babaioff, M.; Ezra, T.; and Feige, U. 2022. On Best-of-Both-Worlds Fair-Share Allocations. In *Proceedings of the 18th Conference on Web and Internet Economics (WINE)*, 237–255.
- Bogomolnaia, A.; and Moulin, H. 2001. A New Solution to the Random Assignment Problem. *Journal of Economic Theory*, 100(2): 295–328.
- Bogomolnaia, A.; Moulin, H.; and Stong, R. 2005. Collective choice under dichotomous preferences. *Journal of Economic Theory*, 122(2): 165–184.
- Brill, M.; Forster, S.; Lackner, M.; Maly, J.; and Peters, J. 2023. Proportionality in Approval-Based Participatory Budgeting. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, 5524–5531.
- Brill, M.; and Peters, J. 2023. Robust and Verifiable Proportionality Axioms for Multiwinner Voting. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC)*, 301.
- Cheng, Y.; Jiang, Z.; Munagala, K.; and Wang, K. 2020. Group Fairness in Committee Selection. *ACM Transactions on Economics and Computation*, 8(4): 1–18.
- Elkind, E.; Faliszewski, P.; Igarashi, A.; Manurangsi, P.; Schmidt-Kraepelin, U.; and Suksompong, W. 2022. The Price of Justified Representation. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*, 4983–4990.
- Fain, B.; Goel, A.; and Munagala, K. 2016. The Core of the Participatory Budgeting Problem. In *Proceedings of the 12th International Conference on Web and Internet Economics (WINE)*, 384–399.
- Feldman, M.; Mauras, S.; Narayan, V. V.; and Ponitka, T. 2023. Breaking the Envy Cycle: Best-of-Both-Worlds Guarantees for Subadditive Valuations. *CoRR*, abs/2304.03706v1.
- Gandhi, R.; Khuller, S.; Parthasarathy, S.; and Srinivasan, A. 2006. Dependent Rounding and Its Applications to Approximation Algorithms. *Journal of the ACM*, 53(3): 324–360.
- Goel, A.; Krishnaswamy, A. K.; Sakshuwong, S.; and Aitamurto, T. 2019. Knapsack Voting for Participatory Budgeting. *ACM Transactions on Economics and Computation*, 7(2): 1–27.
- Gölz, P.; Peters, D.; and Procaccia, A. D. 2022. In This Apportionment Lottery, the House Always Wins. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC)*, 562.
- Hoefler, M.; Schmalhofer, M.; and Varricchio, G. 2023. Best of Both Worlds: Agents with Entitlements. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 564–572.
- Lackner, M.; and Skowron, P. 2023. *Multi-Winner Voting with Approval Preferences*. Springer.
- Los, M.; Christoff, Z.; and Grossi, D. 2022. Proportional Budget Allocations: Towards a Systematization. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*, 398–404.
- Peters, D.; Pierczyński, G.; and Skowron, P. 2021. Proportional Participatory Budgeting with Additive Utilities. In *Proceedings of the 35th Conference on Neural Information Processing Systems (NeurIPS)*, 12726–12737. Extended version available at <https://arxiv.org/abs/2008.13276v2>.
- Rey, S.; and Maly, J. 2023. The (Computational) Social Choice Take on Indivisible Participatory Budgeting. *CoRR*, abs/2303.00621.
- Sánchez-Fernández, L.; Elkind, E.; Lackner, M.; Fernández, N.; Fisteus, J. A.; Basanta-Val, P.; and Skowron, P. 2017. Proportional Justified Representation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, 670–676.