The Distortion of Public-Spirited Participatory Budgeting

Mark Bedaywi[1]**, Bailey Flanigan**[2]**, Mohamad Latifian**[3]**, Nisarg Shah**[1]

[1]University of Toronto
[2]Harvard University
[3]University of Edinburgh
mark.bedaywi@mail.utoronto.ca, bflaniga@andrew.cmu.edu, mohamad.latifian@ed.ac.uk, nisarg@cs.toronto.edu

Abstract

Participatory budgeting (PB) is an increasingly popular tool for democratically allocating limited budgets to public-good projects. In PB, constituents vote on their preferred projects via *ballots*, and then an *aggregation rule* selects a set of projects whose total cost fits within the budget. Recent work studies how to design PB ballots and aggregation rules that yield *low distortion* outcomes (informally, outcomes with high social welfare). Existing distortion bounds, however, rely on strong assumptions that restrict voters' latent utilities. We prove that low distortion PB outcomes can be achieved by dropping these assumptions and instead leveraging the established idea that voters can be *public-spirited*: they may consider others' interests alongside their own when voting.

Flanigan, Procaccia, and Wang (2023) prove that in public-spirited *single-winner* voting (the special case of PB where exactly one project can be funded) with ranking ballots, deterministic aggregation rules can achieve constant distortion. Our first contribution is to extend this analysis to PB; there, we prove that the best distortion permitted by deterministic rules with ranking ballots grows *linearly* in the number of projects. We find that this impossibility—a problem in practice, where m is often large—holds for other known ballots as well. Our second contribution is the design of a new PB ballot format that breaks this linear distortion barrier. This ballot asks voters to rank a predetermined set of *entire feasible bundles* of projects. We design multiple protocols for implementing these ballots, each striking a different trade-off between the number of bundles voters must rank and the distortion: with m bundles, we get sublinear distortion; with polynomial bundles, we get logarithmic distortion; and with pseudopolynomial bundles, we get constant distortion.

1 Introduction

Governments at all scales regularly face the question: Which public-good projects — e.g., building bike paths or installing streetlamps — should they fund with their limited budget? To make such decisions democratically, governments are increasingly using *participatory budgeting* (PB), where a group of constituents convenes to vote on which projects their government should fund. In PB, the government supplies a budget B and a list of m potential projects with corresponding costs. Voters submit their preferences via *ballots*, which are then aggregated via an *aggregation rule*. The output of this rule is a set of projects to be funded whose total cost is at most B. PB is now used all over the world to allocate public funds[1] (Participedia 2023; De Vries, Nemec, and Špaček 2022; Wampler, McNulty, and Touchton 2021).

When designing the PB process, one goal that many consider important is ensuring that the ultimate allocation of funds has high societal benefit. As have many others (e.g., Benadè et al. (2021)), we formalize the "societal benefit" of an allocation by its *utilitarian social welfare*: the total utility it gives to all voters. As such, we adopt the standard model of latent additive utilities: each voter i has *utility* $u_i(a) \in \mathbb{R}_{\geqslant 0}$ for each project a, and her total utility for a set of projects S being funded is $u_i(S) = \sum_{a \in S} u_i(a)$. Then, the *social welfare* of S is equal to $\mathsf{sw}(S) = \sum_{i \in N} u_i(S)$.

If voters' utilities were known, choosing the maximum-welfare allocation would amount to solving the knapsack problem. However, in practice voters' preferences can only be elicited more coarsely through *ballots*. One popular PB ballot format is *rankings by value*, where voters rank all individual projects. It is not hard to see that this ballot format loses far too much information about voters' utilities: suppose there are two projects a and b, and they both cost B so we can fund only one. If the utilities for (a, b) are $(1, 0)$ for half the population and $(0, X)$ for the other half (so, the welfare of b is X times that of a), the resulting ballots will be symmetric. Any deterministic aggregation rule will choose a without loss of generality, suffering unbounded welfare loss as X grows large; the best a randomized aggregation rule can do is to choose a project uniformly at random.

This example illustrates a prohibitive impossibility: in the worst case, *any* deterministic rule over ranking ballots will select an outcome with arbitrarily sub-optimal social welfare (and while randomized rules can do better, they do so trivially by ignoring voters' preferences). In fact, this impossibility holds for all widely-studied ballots in the PB literature due to the same example. Formally, this sub-optimality is captured with the *distortion*: the worst-case (over latent utilities) ratio between the best possible social welfare and that of the outcome. Existing work sidesteps this impossibility by assuming that each voter's utilities sum to 1 (Benadè et al.

[1]See https://en.wikipedia.org/wiki/List_of_participatory_budgeting_votes for a list of use cases.

2021). Although this permits bounded distortion in theory, it is unclear whether these bounds apply in practice: for example, this assumption may not hold in the likely case that all public-good projects on the ballot will more greatly benefit lower-income constituents.

Fortunately, recent work by Flanigan, Procaccia, and Wang (2023) offers a source of hope: under unrestricted utilities, they achieve low distortion in single-winner elections by leveraging the idea that voters may be *public-spirited*: when casting their ballots, voters consider others' interests in addition to their own. While it is not clear that such public-spirited voting behavior would be reliably present in the wild, as Flanigan et al. argue in-depth, research suggests that public spirit can be cultivated via democratic deliberation (Kinder and Kiewiet 1981; Wang, Fishkin, and Luskin 2020; Gastil, Bacci, and Dollinger 2010) — *a practice that is already commonplace in PB elections* (Participedia 2023; De Vries, Nemec, and Špaček 2022). The possibility of cultivating public spirit among PB participants motivates our main research question:

> **Question:** *If voters are public-spirited, can we design PB elections that achieve low (perhaps even constant) distortion with unrestricted voter utilities?*

An affirmative answer to this question would support deliberation as a practicable approach to achieving higher-welfare PB outcomes. While this question builds on Flanigan, Procaccia, and Wang (2023), answering it will require fundamentally new methods because Flanigan *et al*'s results apply only to single-winner voting, a substantially restricted version of the PB setting in which all projects cost B.

1.1 Our Contributions

In the public-spirited voting model of Flanigan, Procaccia, and Wang (2023), each voter i has some *public spirit level* $\gamma_i \in [0, 1]$. She then evaluates each alternative (project) $a \in [m]$ according to her *public-spirited (PS) value* $v_i(a) := (1-\gamma_i)u_i(a)+\gamma_i \mathsf{sw}(a)$, the convex combination of her own utility and the social welfare. Note that this generalizes the standard model in which γ_i is assumed to be 0 (that is, i evaluates a based on just her own utility). We extend Flanigan et al.'s model to PB by assuming additive valuations, so i's PS-value for a set of projects S is simply $v_i(S) = \sum_{a \in S} v_i(a)$. Like Flanigan, Procaccia, and Wang (2023), our distortion bounds are parameterized by $\gamma_{\min} = \min_i \gamma_i$, the minimum public spirit level of any voter. For simplicity, we summarize our main results below assuming $\gamma_{\min}$ is a constant.

We first consider the canonical *rankings* ballot format also studied by Flanigan, Procaccia, and Wang (2023), where voters rank all individual projects by their value. For our first main contribution, we show that the best distortion achievable by any deterministic PB rule using ranking ballots is $\Theta(m)$; for randomized rules, it is $\Theta(\log m)$ (Section 3). Our upper bounds are proven via general reductions from PB to single-winner elections, which may be of independent interest; proving and applying these reductions also leads to new results for the single-winner setting. Our lower bounds imply a fundamental separation between the single-winner voting and PB under public spirit: in single-winner voting with ranking ballots, there are deterministic rules that achieve *constant* distortion (Flanigan, Procaccia, and Wang 2023).

Our linear lower bound of $\Omega(m)$ for deterministic rules is especially bad news: deterministic rules are typically used in practice, and in many PB elections, m can be in the hundreds or thousands (see Footnote 1). Focusing henceforth on deterministic rules, we pursue sublinear distortion by broadening our consideration to other ballot formats. Unfortunately, in Section 4 we find that none of the main PB ballot formats studied in past work (*rankings by value-for-money*, *k-approval*, *knapsack*, and *threshold approval* (Benadè et al. 2021)) permit sublinear distortion (in m). Establishing this linear distortion barrier faced by existing PB ballot formats under public spirit constitutes our next main contribution.

Motivated by this impossibility, in Section 5 we introduce *rankings of predefined bundles*, a novel PB ballot format that asks voters to rank *entire bundles* of projects rather than individual projects. We show that with carefully-chosen bundles, this ballots format *does* permit sublinear distortion in PB. Further, with sufficient (but still fairly limited) information about voters' preferences elicited ahead of time, they can even drop the distortion to *constant*. We study three protocols for using this ballot format, each eliciting more information than the last in exchange for lower distortion:

1. *Protocol 1* permits $O(\sqrt{m})$ distortion and asks voters to rank at most m feasible bundles.
2. *Protocol 2* permits $O(\log m)$ distortion and is a two-round protocol: in round 1 voters rank individual projects, then in round 2 they rank at most $\log m$ bundles crafted based on their votes in round 1.
3. *Protocol 3* permits $O(1)$ distortion. It is similar to Protocol 2 except that in round 2, voters rank $O(m^{1+\log \log m})$ feasible bundles in the worst case.

While Protocol 3 may be impractical in the worst case, we provide empirical evidence that in realistic PB elections, even Protocol 3 would require voters to rank less than m bundles (Section 6). From a theoretical standpoint, Protocol 3 demonstrates the possibility of constant distortion with only pseudo-polynomial bits of information, raising the tantalizing open question of whether constant distortion can be achieved with only *polynomially* many bits.

1.2 Related Work

Our work builds most directly on Flanigan, Procaccia, and Wang (2023), who introduced the public-spirited model. We generalize their work from single-winner elections to the more general PB setting; and while they consider only deterministic aggregation rules, we additionally consider randomized rules. In the process, we prove new insights for the single-winner case. Our work also directly builds on the works of Benadè et al. (2021), who analyzed distortion in PB under the unit-sum utilities assumption. We contrast our bounds to those achievable in their model in the full version[2].

Procaccia and Rosenschein (2006) introduce the distortion framework in single-winner elections under the unit-

[2]https://www.cs.toronto.edu/~nisarg/papers/ps-pb.pdf

sum assumption. We now know that the best distortion achievable by deterministic and randomized rules for this special case are $\Theta(m^2)$ (Caragiannis and Procaccia 2011; Caragiannis et al. 2017) and $\Theta(\sqrt{m})$ (Boutilier et al. 2015; Ebadian et al. 2022), respectively. Optimal distortion bounds have also been identified for k-committee selection (Caragiannis et al. 2017; Borodin et al. 2022), which is still a special case of PB. Some have studied unit-range utilities or metric costs in place of unit-sum utilities (Filos-Ratsikas and Miltersen 2014; Anshelevich et al. 2018), but all these models directly restrict voters' cardinal preferences. For further details, see the survey of Anshelevich et al. (2021).

Multiple approaches other than distortion have been studied for PB. The axiomatic approach has been used to identify aggregation rules satisfying desirable axioms such as various monotonicity properties Talmon and Faliszewski (2019); Baumeister, Boes, and Seeger (2020); Rey, Endriss, and de Haan (2020). Another important consideration in PB is whether the allocation of funds is fair with respect to (groups of) voters (Fain, Munagala, and Shah 2018; Peters, Pierczyński, and Skowron 2021; Brill et al. 2023). For further details, we suggest the survey of Rey and Maly (2023) and the book chapter of Aziz and Shah (2021).

2 Model and Preliminaries

We let $[k] = \{1, \dots, k\}$ for any $k \in \mathbb{N}$, and for a finite set S, let $\Delta(S)$ denote the set of probability distributions over S. We introduce the general framework of participatory budgeting (PB) first, and later introduce single-winner and multiwinner voting as its special cases.

Alternatives A, budget B, and costs c. In a PB instance, there is a set of n *voters* $N = [n]$ and a set of m *alternatives* (projects) A. We denote voters by i, j and alternatives by a, b. There is a total budget of B, which is normalized to 1 without loss of generality, and a cost function $c : A \to [0, 1]$, where $c(a)$ is the *cost* of a. Slightly abusing notation, denote by $c(S) = \sum_{a \in S} c_a$ the total cost of alternatives in S.

Utilities U. Each voter $i \in N$ has a *utility* for each alternative $a \in A$ denoted by $u_i(a) \in \mathbb{R}_{\geqslant 0}$. Together, these utilities form a *utility matrix* $U \in \mathbb{R}^{n \times m}_{\geqslant 0}$. The *social welfare* of $a \in A$ w.r.t. utility matrix U is $\mathsf{sw}(a, U) = \sum_{i \in N} u_i(a)$; for any *set* of alternatives $S \subseteq A$, $\mathsf{sw}(S, U) = \sum_{a \in S} \mathsf{sw}(a, U)$. We use $\mathsf{sw}(a)$ or $\mathsf{sw}(S)$ when U is clear from context.

PS-levels $\vec{\gamma}$ and PS-values V. Following Flanigan, Procaccia, and Wang (2023), we assume that each voter $i \in N$ has a *public spirit (PS) level* $\gamma_i \in [0, 1]$, and, together, these PS-levels form the PS-vector $\vec{\gamma} \in [0, 1]^n$. For a given $\vec{\gamma}$, we let $\gamma_{\min} := \min_{i \in N} \gamma_i$ be the minimum level of public spirit among voters. Each voter i evaluates each alternative a by her *PS-value* $v_i(a)$, a convex combination of her personal utility $u_i(a)$ and $\mathsf{sw}(a)/n$, the average voter's utility for a:

$$v_i(a) = (1 - \gamma_i) \cdot u_i(a) + \gamma_i \cdot \mathsf{sw}(a)/n.$$

Note that this model does not restrict voters' utilities; rather, it assumes something about how they translate their utilities into votes. These PS-values form the *PS-value matrix* $V_{\vec{\gamma}, U} \in \mathbb{R}^{n \times m}_{\geqslant 0}$. For each $S \subseteq A$, let $v_i(S) := \sum_{a \in S} v_i(a)$.

Instances and special cases. An *instance* of the PB problem is composed of the elements defined so far: $I = (A, B, c, U, \vec{\gamma})$. Let $\mathcal{I}$ be the set of all PB instances. Let $\mathcal{F}(I) = \{S \subseteq A \colon c(S) \leqslant 1\}$ be the set of *budget-feasible* subsets of A in instance I. $\mathcal{F}$ will be a generic such set.

We will sometimes build our results using ideas from k*-committee selection* and *single-winner voting* — two restrictions of the PB setting. Formally, all instances of k*-committee selection* are captured when $\mathcal{I}$ is exclusively restricted to instances I with $c(\cdot) = 1/k$ (i.e., all alternatives have cost $1/k$), so $\mathcal{F}(I)$ consists of all subsets of alternatives of size k. Single-winner voting is the further restriction in which $c(\cdot) = 1$. We will let $\mathcal{I}^{\text{single-win}} := \{I | c(\cdot) = 1\}$ denote the set of all single-winner voting instances.

Ballot formats. Since it is cognitively burdensome for voters to report cardinal preferences, preferences are often elicited using discrete ballots. We denote a generic ballot format as X, and let $\rho_i(I, \mathsf{X})$ be the ballot submitted by voter i in instance I. Correspondingly, let $\vec{\rho}(I, \mathsf{X}) = (\rho_1(I, \mathsf{X}), \dots, \rho_n(I, \mathsf{X}))$ be the *vote profile*. When I, X are clear, we will drop these from the notation. In Section 5, we will design multi-round elicitation protocols; when there are multiple rounds, a "vote profile" will refer to the profile of votes collected in the final round of elicitation.

We primarily consider ordinal ballot formats, of which we study two types. First, we consider canonical *ranking* ballots ($\mathsf{X} = \mathsf{rank}$), which ask voters to rank alternatives. Then, $\rho_i(I, \mathsf{rank})$ is the permutation of A implied by the ordering of i's PS-values, so $v_i(a) > v_i(b) \Rightarrow a \succ_{\rho_i(I, \mathsf{rank})} b$ for all $a, b \in A$ (ties are broken arbitrarily, and $a \succ_\rho b$ denotes that a is ranked ahead of b in ballot ρ). We then introduce a novel ordinal ballot format, *ranking of predefined bundles* ($\mathsf{X} = \mathsf{rank\text{-}b}$), which asks voters to rank *entire bundles* of alternatives. Formally, the $\mathsf{rank\text{-}b}$ ballot format accepts an argument of a collection of predefined bundles $\mathcal{P} \subseteq \mathcal{F}(I)$; then, $\rho_i(I, \mathsf{rank\text{-}b}(\mathcal{P}))$ is a permutation of the elements of $\mathcal{P}$ such that $v_i(S) > v_i(S') \Rightarrow S \succ_{\rho_i(I, \mathsf{rank\text{-}b}(\mathcal{P}))} S'$ for all $S, S' \in \mathcal{P}$. To paint a more complete picture, we also briefly consider other non-ordinal ballot formats in Section 4, but their definitions and proofs can be found in the full version.

Aggregation rules. A (randomized) *aggregation rule* f takes as input the vote profile $\vec{\rho}$ (from the final round of elicitation, if there are multiple rounds) and returns a distribution over feasible bundles (an element of $\Delta(\mathcal{F})$). We say that f is deterministic if its output always has singleton support. We will sometimes talk about *single-winner rules* versus *PB rules*. Formally, a single-winner rule must output an element of $\Delta([m])$ while a PB rule can output any element of $\Delta(\mathcal{F})$.

We will frequently use the rule *Copeland*, so we define it here. *Copeland* is traditionally defined for the single-winner case with rank ballots. We make the natural extension here to define *Copeland* also for $\mathsf{rank\text{-}b}(\mathcal{P})$ ballot formats. All other rules we consider are defined as needed.

Definition 1 (*Copeland*)**.** Each alternative has a *score*, equal to the number of alternatives it defeats in pairwise elections. The *Copeland* winner is the one with the highest score.

When we want to choose multiple winners W, we often use the rule *Iterative Copeland*: *Copeland* is used, the winner is added to W and removed from the election, and then *Copeland* is run again on the remaining instance, and so on.

Distortion. The *distortion* measures the efficiency of a combination of a ballot format and an aggregation rule (if there are multiple rounds, the rule applies to the final round). Formally, it is the *worst-case* over all instances of the ratio between the best achievable social welfare and the output of the aggregation rule. Our bounds will depend explicitly on m and $\gamma_{\min}$, so we denote the subset of $\mathcal{I}$ with $m, \gamma_{\min}$ as

$$\mathcal{I}_{m,\gamma_{\min}} := \{I \in \mathcal{I} : |A| = m \wedge \min_{i \in N} \gamma_i = \gamma_{\min}\}.$$

Then, the distortion is defined as

$$\text{dist}_{\mathsf{X}}(f) = \sup_{n \geqslant 1} \sup_{I \in \mathcal{I}_{m,\gamma_{\min}}} \frac{\max_{S \in \mathcal{F}(I)} \mathsf{sw}(S, U)}{\mathbb{E}_{S' \sim f(\vec{\rho}(I, \mathsf{X}))} \mathsf{sw}(S', U)}.$$

We sometimes study a rule's distortion in the *single-winner* case, where the set of instances is restricted to $\mathcal{I}^{\text{single-win}}$. The *single-winner distortion* $\text{dist}_{\mathsf{X}}^{\text{single-win}}(f)$ is therefore defined identically to $\text{dist}_{\mathsf{X}}(f)$ except the second supremum is taken over $\mathcal{I}_{m,\gamma_{\min}}^{\text{single-win}}$ (analogous to $\mathcal{I}_{m,\gamma_{\min}}$).

Our distortion bounds will assume $m \geqslant 2$ and $\gamma_{\min} \in (0, 1]$. We are interested in the lowest distortion possible by *any* aggregation rule using a given ballot format; this is a measure of the usefulness of the information contained in the ballot format for social welfare maximization.

Preliminaries. For comparison, in the full version we prove that with no public spirit and unrestricted utilities, for all ballot formats we consider, all deterministic rules have unbounded distortion and the randomized rules have at least m distortion. We also contrast many of our bounds to those achievable under unit-sum utilities.

Our upper bounds will often use the following lemma, which is a simple generalization of Lemma 3.1 of Flanigan, Procaccia, and Wang (2023).

Lemma 1. *Let $A_1, A_2 \subseteq A$ be any two subsets of alternatives. Fix any $\alpha \geqslant 0$ and define $N_{A_1 \succ A_2} = \{i \in N : \alpha \cdot v_i(A_1) \geqslant v_i(A_2)\}$. Then:*

$$\frac{\mathsf{sw}(A_2)}{\mathsf{sw}(A_1)} \leqslant \alpha \cdot \left(\frac{1 - \gamma_{\min}}{\gamma_{\min}} \frac{n}{|N_{A_1 \succ A_2}|} + 1 \right).$$

In the full version, we also prove a *robust* version of Lemma 1 showing that its guarantee degrades smoothly as an increasing number of voters have $\gamma_i = 0$. We can replace Lemma 1 with its robust version in all our upper bound proofs, meaning that our upper bounds degrade smoothly.

3 PB with Rankings over Projects

We begin by studying *ranking* ballot format rank, the canonical ballot format in single-winner election and the one studied by Flanigan, Procaccia, and Wang (2023) (for only deterministic aggregation rules). Here we extend their results to PB for deterministic *and* randomized rules.

3.1 Deterministic Rules

We begin by upper-bounding the distortion of *Copeland* in PB, due to its strong performance in the single-winner case. Note that applying a single-winner rule in the PB setting leads to selecting a single project

Theorem 1. $\text{dist}_{\mathsf{rank}}(\textit{Copeland}) \in \mathcal{O}\left(m/\gamma_{\min}^2\right)$.

Proof. First, we prove a general reduction that converts any deterministic single-winner rule to a deterministic PB rule.

Lemma 2 (PB → Single-Winner). *For any $d \geqslant 1$, any deterministic rule f with distortion d in single-winner voting has distortion $\text{dist}_{\mathsf{rank}}(f) \leqslant m \cdot d$ in participatory budgeting.*

Proof. The proof of this lemma is straightforward: fix any instance and let f return the singleton set $\{a\}$. Let A^* be an optimal budget-feasible set. Then,

$$\frac{\mathsf{sw}(A^*)}{\mathsf{sw}(a)} = \sum_{a^* \in A^*} \frac{\mathsf{sw}(a^*)}{\mathsf{sw}(a)} \leqslant m \cdot \max_{a^* \in A^*} \frac{\mathsf{sw}(a^*)}{\mathsf{sw}(a)} \leqslant m \cdot d. \quad \square$$

Intuitively, a factor of m should be incurred from single-winner to PB: unlike in the single-winner case, in PB even when $\gamma_{\min} = 1$ (so all voters vote unanimously for the highest-welfare alternative), it can remain unclear how to make cost trade-offs without cardinal information, and deterministic rules still incur $\Omega(m)$ distortion. To conclude the proof, we apply this reduction via the following bound on *Copeland*'s distortion in the single-winner case, proven in Thm. 3.3 of Flanigan, Procaccia, and Wang (2023):

$$\text{dist}_{\mathsf{rank}}^{\text{single-win}}(\textit{Copeland}) \in O(1/\gamma_{\min}^2). \quad \square \quad (1)$$

Now, we prove a lower bound on the distortion achievable by *any* deterministic aggregation rule in the PB setting.

Theorem 2. *Every deterministic rule f has distortion*

$$\text{dist}_{\mathsf{rank}}(f) \in \Omega\left(m/\gamma_{\min}\right).$$

Proof sketch. The lower bound instance has only two (maximal) feasible sets, one containing a single alternative a and the other containing the remaining alternatives. A few voters rank a above the other alternatives, while all other voters do the opposite. We show that there exist utilities for which either choice can be sub-optimal by $\Omega(m/\gamma_{\min})$. $\square$

Together, Theorems 1 and 2 imply that *Copeland* achieves optimal dependency on m and is within a $1/\gamma_{\min}$ factor of optimal overall. There are two possible sources of this remaining gap: our use of (or analysis of) a single-winner rule via the reduction in Lemma 2, and our choice to apply the reduction specifically *Copeland*. To shed light on the role of each of these, we now prove a universal lower bound showing that at least for large m, *Copeland* is the optimal single-winner rule. This is a novel finding of independent interest for the single-winner case, given that Flanigan, Procaccia, and Wang (2023) do not give any universal lower bounds.

Theorem 3. *For all deterministic single-winner rules f,*

$$\text{dist}_{\mathsf{rank}}^{\text{single-win}}(f) \in \Omega(\min\{m/\gamma_{\min},\ 1/\gamma_{\min}^2\}).$$

The proof of this theorem alongside all other missing proofs can be found in the full version. The idea is to use a cyclic profile, where an equal number of voters submit each of m cyclically shifted permutations. The contribution is in the intricate derivation of the piecewise bound.

Comparing Equation (1) and Theorem 3, when $m \in \Omega(1/\gamma_{\min})$, *Copeland*'s distortion matches this lower bound. When $m \in o(1/\gamma_{\min})$, *Plurality*, which selects the most common first-choice alternative, provides a matching distortion upper bound of $O(m/\gamma_{\min})$ (Flanigan, Procaccia, and Wang 2023, Proposition 3.5). Hence, this lower bound "resolves" the deterministic single-winner case in that for every regime of m and $\gamma_{\min}$, there is some voting rule that asymptotically matches it. Whether a single, $\gamma_{\min}$-oblivious rule can do so remains open for future work.

3.2 Randomized Rules

We take a parallel approach to analyze what distortion is achievable with rank ballots and *randomized* rules. Because Flanigan, Procaccia, and Wang (2023) did not study randomized single-winner rules, we must first identify and analyze a low-distortion single-winner rule anew—a result that is of independent interest for the single-winner case. The rule we select is *Maximal Lottery*, a single-winner rule originally proposed by Kreweras (1965).[3]

Definition 2 (*Maximal Lottery*)**.** The (directed) domination graph G consists of a vertex corresponding to each alternative $a \in A$, and an edge from a to b whenever a defeats b in a pairwise election (ties can be broken arbitrarily). The maximal lottery rule returns a distribution p over the vertices such that for any vertex $b \in A$, the probability of picking b or a vertex a with an edge to b is at least $1/2$. The existence of such a distribution can be inferred from, e.g., Farkas' lemma (see Thm. 2.4 of Harutyunyan et al. (2017)).

We now upper-bound *Maximal Lottery*'s distortion in PB:

Theorem 4. $\mathrm{dist}_{\mathsf{rank}}(\textit{Maximal Lottery}) \in O({}^{\log(m)}/_{\gamma_{\min}})$.

Proof. We again begin by proving a general-purpose reduction to convert single-winner rules to PB rules. The reduction in the randomized case is more involved, and we do it in two steps: we first reduce PB to committee selection (Lemma 3), and then reduce that to single-winner voting (Lemma 4).

Lemma 3 (PB → Committee). *Fix any $d \geqslant 1$. If there exists a randomized k-committee selection rule $f_{m',k}$ with distortion at most d for each $m' \leqslant m$ and $k \in [m']$, then there exists a randomized participatory budgeting rule f with distortion at most $\mathrm{dist}_{\mathsf{rank}}(f) \leqslant 2d \cdot (\lceil \log_2(m) \rceil + 1)$.*

Lemma 4 (Committee → Single-Winner). *Fix any $k \in [m]$ and $d \geqslant 1$. If there exists a single-winner rule $f_{m'}$ with distortion at most d for each $m' \leqslant m$, then there exists a k-committee selection rule f with distortion at most d. If $f_{m'}$ is deterministic then so is f.*

[3]This rule has been rediscovered numerous times (Laffond, Laslier, and Le Breton 1993; Fishburn 1984; Fisher and Ryan 1995; Rivest and Shen 2010). To the best of our knowledge, this is the first analysis of this rule's utilitarian distortion.

The $\log(m)$ overhead in Lemma 3 comes from partitioning the alternatives into $O(\log m)$ buckets and then applying a k-committee selection rule to a random bucket (similar approaches appear in other work, e.g. Benadè et al. (2021)). The proof of Lemma 4 generalizes ideas from an analogous reduction for deterministic rules by Goel, Hulett, and Krishnaswamy (2018).

Next, to bound the distortion of *Maximal Lottery* via Lemmas 3 and 4, we must first upper-bound its distortion in the public-spirited *single-winner* setting.

Theorem 5. $\mathrm{dist}^{\text{single-win}}_{\mathsf{rank}}(\textit{Maximal Lottery}) \in \mathcal{O}\left(1/\gamma_{\min}\right).$

The approach is to apply Lemma 1 using the insight that *Maximal Lottery* picks either the optimal alternative or an alternative that pairwise-defeats it with probability at least $1/2$.

Finally, applying Theorem 5 along with our reductions, we conclude that in the PB setting, *Maximal Lottery* has distortion at most $O(\log(m)/\gamma_{\min})$, as needed. ☐

Now, we lower bound the distortion achievable by *any* randomized aggregation rule in the PB setting.

Theorem 6. *For all randomized rules f,*

$$\mathrm{dist}_{\mathsf{rank}}(f) \in \Omega(\log(m)).$$

Proof sketch. To prove this we use the following novel construction: the alternatives are partitioned into roughly $\sqrt{m}$ buckets, and the ℓ-th bucket consists of ℓ alternatives with cost $1/\ell$ each. We use budget-feasibility to bound the average marginal probability with which an alternative from each bucket can be chosen, and construct a set of latent utility matrices such that any outcome is suboptimal with respect to at least one of them. ☐

Together, Theorems 4 and 6 imply that *Maximal Lottery* achieves optimal dependency on m and is within a $1/\gamma_{\min}$ factor of optimal overall. As before, we explore the source of this $1/\gamma_{\min}$ gap by showing that at least for large m, *Maximal Lottery* is the optimal randomized single-winner rule:

Theorem 7. *For all randomized single-winner rules f,*

$$\mathrm{dist}^{\text{single-win}}_{\mathsf{rank}}(f) \in \Omega\left(\min\left\{m, 1/\gamma_{\min}\right\}\right).$$

This lower bound is proven by the same construction as in Theorem 3, the analogous lower bound for the deterministic case. As with *Copeland* in the deterministic case, this bound shows that *Maximal Lottery* is optimal when $m \geqslant \Omega(1/\gamma_{\min})$. When $m \in o(1/\gamma_{\min})$, this bound is matched by the rule that chooses a uniformly random alternative. As before, whether a *single* $\gamma_{\min}$-oblivious randomized rule can match this lower bound remains open.

4 PB with Other Known Ballot Formats

In Section 3, we found something that may initially seem strange: in both the randomized and deterministic cases, voting rules designed for the single-winner setting—which output just a single alternative, even in PB—achieved optimal dependence on m. This is due to the weakness of the rank ballot format: designed for single-winner voting, rank fails to capture fundamental aspects of the PB setting.

Others have tried to address this problem in the unit-sum utilities model by designing better ballot formats. We now pursue the same approach in the public spirit model: focusing henceforth on deterministic rules for their practicality, we aim to identify a ballot format that achieves sublinear distortion in m, thereby surpassing our lower bound of $\Omega(m)$ in Theorem 2. To this end, we examine four other known PB ballot formats (see Benadè et al. (2021)): *rankings by value for money*, where each voter i ranks alternatives by $v_i(a)/c(a)$; *k-approvals*, where each voter i submits a set of k alternatives with the highest $v_i(a)$; *knapsack*, where each voter i submits the budget-feasible set with the highest value $v_i(S)$; and *threshold approvals*, where the ballot format specifies a threshold t and each voter i submits the set of alternatives with $v_i(a) \geqslant t$. Unfortunately, the answer is resoundingly negative for all these ballot formats:

Theorem 8. *All deterministic rules have distortion...*

- ∞ *using* ***rankings by value for money*** *ballots;*
- $\Omega(m^2/\gamma_{\min})$ *using* **1-*approval*** *ballots and* ∞ *using* ***k-approval*** *ballots with* $k > 1$*;*
- $\Omega(m/\gamma_{\min})$ *using* ***knapsack*** *ballots; and*
- $\Omega(m)$ *using* ***threshold approval*** *ballots (any threshold).*

While none of these ballot formats yield better distortion than rank ballots, the worst-case distortion of *knapsack* and *approval* ballots under public-spirited voting may still be of interest, given they are used in real-world PB elections (see Footnote 1). We thus also derive *upper* bounds on the distortion possible with these ballot formats (in some cases, for randomized rules, too) in the full version. Most notably, we show that with *1-approval* ballots there is a deterministic rule with $O(m^2/\gamma_{\min})$ distortion, perfectly matching our lower bound; for *knapsack* ballots, there is a deterministic rule with $O(m^3/\gamma_{\min}^2)$ distortion—a striking improvement over the unit-sum case;[4] and for *threshold approval* ballots, there is a deterministic rule with $O(m^2/\gamma_{\min})$ distortion.

5 PB with Ranking of Predefined Bundles

We have now shown that for all commonly-studied PB ballot formats, all deterministic rules incur $\Omega(m)$ distortion — an issue in the practical case where m is large. Motivated by this, we now study the distortion possible with our novel ballot format, *ranking of predefined bundles* (rank-b($\mathcal{P}$)).

In Sections 5.1-5.3, we will explore various protocols for using rank-b($\mathcal{P}$) ballots, which differ in how $\mathcal{P}$ is chosen. For intuition, note that the lowest-distortion choice of $\mathcal{P}$ is simply $\mathcal{F}$; then, we are effectively in the single-winner setting and *Copeland* guarantees constant (in m) distortion. However, this choice of $\mathcal{P}$ comes at a steep elicitation cost, requiring voters to rank exponentially many bundles.

Our refined goal, therefore, is to design $\mathcal{P}$ to permit low distortion *while containing at most polynomial (or pseudopolynomial) bundles*. After designing and analyzing various such choices of $\mathcal{P}$, in Section 6 we explore the practicality of the resulting elictation protocols.

5.1 Sublinear Distortion

We first propose rank-b with *high-low bundles* ($\mathcal{P} = \text{HLB}$), which we show permits sublinear distortion.

High-low bundles (HLB): Let $L = \{a \in A : c(a) \leqslant 1/\lceil\sqrt{m}\rceil\}$ be the set of *low-cost* alternatives and $H = A \backslash L$ be the set of *high-cost* alternatives. The *high-low bundling* rule (HLB) partitions L into at most $\lceil\sqrt{m}\rceil$ feasible bundles,[5] and H into an arbitrary partition of feasible bundles. Then it defines $\mathcal{P}$ to be the union of these partitions.

The rank-b(HLB) ballot asks voters to rank $|\mathcal{P}| \leqslant |H| + |L| = m$ bundles (and in fact, $|\mathcal{P}| \leqslant m - |L| \cdot (1 - 1/\lceil\sqrt{m}\rceil)$, so if there are many low-cost projects, $|\mathcal{P}| \ll m$). Finally, in Theorem 9 we show that if *Copeland* rule is applied to the voting profile elicited via rank-b(HLB) on $\mathcal{P}$, the rank-b ballot format dominates all the previous ballot formats by a factor of $O(\sqrt{m})$.

Theorem 9. $\text{dist}_{\text{rank-b(HLB)}}(\text{Copeland}) = O(\sqrt{m}/\gamma_{\min}^2)$.

The main idea is that if A^* is the optimal bundle, then either $\mathsf{sw}(L \cap A^*)$ or $\mathsf{sw}(H \cap A^*)$ must have welfare at least $\mathsf{sw}(A^*)/2$. Then, there must be a bundle in $\mathcal{P}$ with welfare at least $(1/\lceil\sqrt{m}\rceil) \cdot \mathsf{sw}(L \cap A^*)$ because L was partitioned into at most $\lceil\sqrt{m}\rceil$ bundles, and also one with welfare at least $(1/\lceil\sqrt{m}\rceil) \cdot \mathsf{sw}(H \cap A^*)$ because $|H \cap A^*| \leqslant \lceil\sqrt{m}\rceil$.

While this is already a significant improvement on previous results, there is room for more: crafting predefined bundles with no information about (and thus no regard for) voters' preferences can be both theoretically lossy and practically unappealing. Thus, we next explore: *what distortion is possible when our bundling rule has some knowledge of voters' preferences?* We explore this question in Sections 5.2 and 5.3 by defining a two-round elicitation protocol: in Round 1, we elicit voter preferences using the canonical rank ballot format; then, in Round 2, we use this preference information to craft $\mathcal{P}$ and deploy ballot rank-b($\mathcal{P}$). We denote this two-round ballot format as rank $\rightarrow$ rank-b($\mathcal{P}$).

5.2 Logarithmic Distortion in Two Rounds

We now propose rank-b with *tiered-cost bundles* ($\mathcal{P} = \text{TCB}$). At a high level, TCB partitions alternatives into $O(\log m)$ tiers by cost, and then uses *Iterative Copeland* to select a feasible bundle of $m/2^\ell$ alternatives from the tier containing alternatives with costs between $2^{\ell-1}/m$ and $2^\ell/m$.

Tiered-cost bundles (TCB): Set $L = \lceil\log_2 m\rceil$. For each $\ell \in [L]$, define the tier T_ℓ such that

$$T_\ell = \{a \in A : 2^{\ell-1}/m < c(a) \leqslant 2^\ell/m\} \quad \text{for all } \ell > 0;$$

let $T_0 = \{a \in A : c(a) \leqslant 1/m\}$. Then, use *Iterative Copeland* to pick a bundle $P_\ell \subseteq T_\ell$ of size $t_\ell = \lfloor\min(|T_\ell|, \max(1, m/2^\ell))\rfloor$. Since $c(a) \leqslant 2^\ell/m$ for each $a \in T_\ell$, P_ℓ is budget-feasible. Set $\mathcal{P} = \{P_0, P_1, \ldots, P_L\}$. Then, TCB asks voters to rank $L \leqslant 1 + \lceil\log_2 m\rceil$ bundles.

[4] It is striking how beneficial public spirit is to *knapsack* ballots: without it, even with unit-sum utilities, any deterministic rule over knapsack ballots has $\Omega(2^m/\sqrt{m})$ distortion (Benadè et al. 2021).

[5] This is possible because $|L| \leqslant m$ and any subset of $\lceil\sqrt{m}\rceil$ alternatives from L is feasible.

After asking voters to rank a total of at most $m + 1 + \lceil \log_2 m \rceil$ objects over both rounds, aggregating via *Copeland* achieves distortion $O(\log(m)/\gamma_{\min}^4)$:

Theorem 10.

$$\text{dist}_{\mathsf{rank}\to\mathsf{rank\text{-}b}(\text{TCB})}(\text{Copeland}) = O(\log(m)/\gamma_{\min}^4).$$

The key insight is that for the optimal bundle A^*, $\mathsf{sw}(A^*) = \sum_\ell \mathsf{sw}(A^* \cap T_\ell)$, so the best of the $1 + \lceil \log_2 m \rceil$ feasible bundles in the sum (call it $A^* \cap T_{\ell'}$) must be an $O(\log m)$ approximation of A^*. Then, the welfare of the best t_ℓ-sized subset $P_\ell^* \subseteq T_{\ell'}$ 2-approximates that of $A^* \cap T_{\ell'}$; P_ℓ constant-approximates the welfare of P_ℓ^* (by the distortion of *Iterative Copeland*); and the chosen bundle constant-approximates the welfare of P_ℓ (by the distortion of the final *Copeland* aggregation).

5.3 Constant Distortion in Two Rounds

Finally, we propose rank-b with *exhaustive bundles* ($\mathcal{P}$ = EB). EB uses the same tiers as TCB, but instead of having each bundle consist of alternatives from the same tier, it crafts bundles by using *Iterative Copeland* to choose a subset $S_\ell \subseteq T_\ell$ of size t_ℓ from *every* tier and putting them together as $\cup_\ell S_\ell$. Ideally, we want to explore every possible combination of values of $t_0, \ldots, t_{\lceil \log_2 m \rceil}$, so long as the resulting bundle is feasible, but a slight optimization is achieved by only choosing values that are powers of 2.

Exhaustive bundles (EB): Let $L = \lceil \log_2 m \rceil$ and define tiers $T_0, \ldots, T_L$ as in TCB. Fix $R = \lfloor \log_2 m \rfloor$. For each $\ell \in [L]$ and $r \in \{0, 2^0, \ldots, 2^R\}$ such that $|T_\ell| \geqslant r$, choose $P_{\ell,r} \subseteq T_\ell$ of size r using *Iterative Copeland* applied to the rank ballots from Round 1 (if $p = 0$, simply choose $\emptyset$). Call a sequence $\vec{t} = (t_0, \ldots, t_L)$ *valid* if $t_\ell \in \{0, 2^0, \ldots, 2^R\}$ for each $\ell \in [L]$, and for such a sequence define $P_{\vec{t}} = \cup_{\ell=0}^{L} P_{\ell,t_\ell}$. In other words, in a valid sequence, the ℓ-th element represents the number of alternatives that are selected from T_ℓ, and with this definition each valid sequence leads to a potential bundle. Finally, let $\mathcal{P} = \{P_{\vec{t}} : \vec{t} \text{ is valid} \wedge P_{\vec{t}} \in \mathcal{F}\}$, so all bundles in $\mathcal{P}$ are feasible. Note that $|\mathcal{P}|$ is at most the number of valid sequences, which is $(1 + R)^{1+L} = O((\log m)^{O(\log m)}) = O(m^{O(\log \log m)})$.[6]

In exchange for asking voters to rank quasipolynomially many objects across two rounds, using *Copeland* to aggregate, we achieve constant distortion:

Theorem 11. $\text{dist}_{\mathsf{rank}\to\mathsf{rank\text{-}b}(\text{EB})}(\text{Copeland}) = O(1/\gamma_{\min}^4)$.

The proof uses the following approach. Let A^* be an optimal bundle. Then, consider the sequence $\vec{t}$, where $t_\ell = 2^{\lfloor \log_2 |A^* \cap T_\ell| \rfloor - 1}$ for each $\ell \in \{0, 1, \ldots, L\}$. Since $t_\ell \leqslant |A^* \cap T_\ell|/2$, $c(P_{\ell,t_\ell}) \leqslant c(A^* \cap T_\ell)$, $P_{\vec{t}}$ must be feasible. We then show that its welfare, constant-approximates that of A^*. The only thing is that when $|A^* \cap T_\ell| = 1$, we cannot set $t_\ell = 2^{\lfloor \log_2 |A^* \cap T_\ell| \rfloor - 1}$. This is addressed by taking two cases, depending on whether much of the welfare of A^* is contributed by those $A^* \cap T_\ell$ that have size 1 or the rest.

[6] Although this is many bundles, they are similar, potentially decreasing cognitive load of ranking them: they consist of combinations of at most $(1 + R) \cdot (1 + L) = O(m)$ many bundles $P_{\ell,r}$.

6 Discussion

On the practicality of rank-b ballots. We propose that our *ranking of predefined bundles* ballot format has three main practical advantages in addition to low distortion. (1) It is fully ordinal in contrast to, e.g., threshold approval votes, which ask voters to compare projects via precise numeric utility values. (2) Comparing entire feasible bundles may provide voters more context about cost trade-offs than ballots where they compare individual projects. (3) The aggregation rules always select a bundle that is on the ballot, allowing every vote favorably ranking the winning bundle to be interpreted as a direct endorsement of the final outcome.

While one may worry that doing two rounds of elicitation is impracticable, PB participants often meet several times, so doing so is likely feasible. In fact, the flexibility in the PB process may permit a variety multi-round protocols with even more favorable trade-offs. A second potential worry is that low-distortion rank-b ballots may ask voters to rank too many bundles. We now show that in 1244 real PB elections from https://pabulib.org (with some randomized imputation of incomplete preferences), even our rank $\to$ rank-b(EB) ballot (with some minor heuristic tweaks that maintain constant distortion) typically requires voters to rank *far less than* m *bundles*. Details on data and implementation, plus some supplemental results, are found in the full version.

Future directions. rank-b ballots represent an exciting future direction in PB: there is an expansive design space of bundles and multi-round protocols that can potentially drive down query complexity, guarantee low distortion, satisfy desirable axioms, and be well-received in real-world experiments. Beyond rank-b ballots, there are many questions remaining about what public spirit looks like in real democratic contexts, and how this can be incorporated into social choice theory. For instance: (1) In what other social choice contexts — such as matching (Filos-Ratsikas, Frederiksen, and Zhang 2014) and fair division (Halpern and Shah 2021) — is public spirit a reasonable assumption?, (2) Can we measure the degree and nature of public spirit resulting from different democratic processes, deliberative or otherwise?, and (3) If voters account for a welfare notion other than utilitarian social welfare (or likewise, we care about other objectives like Nash welfare or proportional fairness (Ebadian et al. 2022)), can one prove similar guarantees?

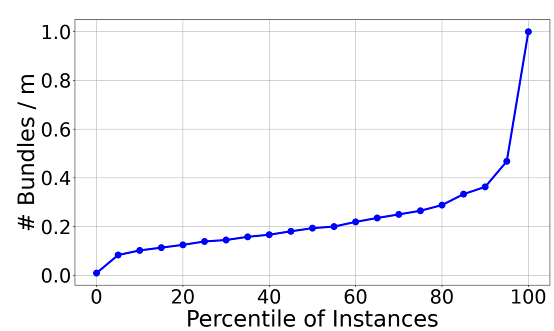

Figure 1: The number of bundles rank-b(EB) asks voters to rank, per alternative, in 1244 instances ordered by quantile.

Acknowledgments

Nisarg Shah gratefully acknowledges support from an NSERC Discovery Grant. Bailey Flanigan is thankful to be supported by the Harvard Data Science Initiative Fellowship.

References

Anshelevich, E.; Bhardwaj, O.; Elkind, E.; Postl, J.; and Skowron, P. 2018. Approximating Optimal Social Choice under Metric Preferences. *Artificial Intelligence*, 264: 27–51.

Anshelevich, E.; Filos-Ratsikas, A.; Shah, N.; and Voudouris, A. A. 2021. Distortion in Social Choice Problems: The First 15 Years and Beyond. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, 4294–4301. Survey Track.

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.

Baumeister, D.; Boes, L.; and Seeger, T. 2020. Irresolute approval-based budgeting. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*, 1774–1776.

Benadè, G.; Nath, S.; Procaccia, A. D.; and Shah, N. 2021. Preference Elicitation For Participatory Budgeting. *Management Science*, 65(5): 2813–2827.

Borodin, A.; Halpern, D.; Latifian, M.; and Shah, N. 2022. Distortion in voting with top-t preferences. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*, 116–122.

Boutilier, C.; Caragiannis, I.; Haber, S.; Lu, T.; Procaccia, A. D.; and Sheffet, O. 2015. Optimal Social Choice Functions: A Utilitarian View. *Artificial Intelligence*, 227: 190–213.

Brill, M.; Forster, S.; Lackner, M.; Maly, J.; and Peters, J. 2023. Proportionality in Approval-Based Participatory Budgeting. In *Thirty-Seventh AAAI Conference on Artificial Intelligence*, 5524–5531.

Caragiannis, I.; Nath, S.; Procaccia, A. D.; and Shah, N. 2017. Subset Selection Via Implicit Utilitarian Voting. *Journal of Artificial Intelligence Research*, 58: 123–152.

Caragiannis, I.; and Procaccia, A. D. 2011. Voting Almost Maximizes Social Welfare Despite Limited Communication. *Artificial Intelligence*, 175(9–10): 1655–1671.

De Vries, M. S.; Nemec, J.; and Špaček, D. 2022. International Trends in Participatory Budgeting. *Cham: Palgrave Macmillan*.

Ebadian, S.; Kahng, A.; Peters, D.; and Shah, N. 2022. Optimized Distortion and Proportional Fairness in Voting. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC)*, 563–600.

Fain, B.; Munagala, K.; and Shah, N. 2018. Fair Allocation of Indivisible Public Goods. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, 575–592.

Filos-Ratsikas, A.; Frederiksen, S. K. S.; and Zhang, J. 2014. Social welfare in one-sided matchings: Random priority and beyond. In *International Symposium on Algorithmic Game Theory*, 1–12.

Filos-Ratsikas, A.; and Miltersen, P. B. 2014. Truthful approximations to range voting. In *Proceedings of the 10th Conference on Web and Internet Economics (WINE)*, 175–188.

Fishburn, P. C. 1984. Probabilistic social choice based on simple voting comparisons. *The Review of Economic Studies*, 51(4): 683–692.

Fisher, D. C.; and Ryan, J. 1995. Tournament games and positive tournaments. *Journal of Graph Theory*, 19(2): 217–236.

Flanigan, B.; Procaccia, A. D.; and Wang, S. 2023. Distortion Under Public-Spirited Voting. In *Proceedings of the 24th ACM Conference on Economics and Computation*, 700.

Gastil, J.; Bacci, C.; and Dollinger, M. 2010. Is Deliberation Neutral? Patterns of Attitude Change During the Deliberative Polls. *Journal of public deliberation*, 6(2).

Goel, A.; Hulett, R.; and Krishnaswamy, A. K. 2018. Relating Metric Distortion and Fairness of Social Choice Rules. In *Proceedings of the 13th Workshop on Economics of Networks, Systems and Computation*, NetEcon '18. New York, NY, USA: Association for Computing Machinery. ISBN 9781450359160.

Halpern, D.; and Shah, N. 2021. Fair and Efficient Resource Allocation with Partial Information. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, 224–230.

Harutyunyan, A.; Le, T.-N.; Newman, A.; and Thomassé, S. 2017. Domination and fractional domination in digraphs. *arXiv preprint arXiv:1708.00423*.

Kinder, D. R.; and Kiewiet, D. R. 1981. Sociotropic politics: the American case. *British Journal of Political Science*, 11(2).

Kreweras, G. 1965. Aggregation of preference orderings. In *Mathematics and Social Sciences I: Proceedings of the seminars of Menthon-Saint-Bernard*, 73–79.

Laffond, G.; Laslier, J.-F.; and Le Breton, M. 1993. The bipartisan set of a tournament game. *Games and Economic Behavior*, 5(1): 182–201.

Participedia. 2023. https://participedia.net/search?selectedCategory=all\&query=participatory\%20budgeting. Accessed: 2025-01-27.

Peters, D.; Pierczyński, G.; and Skowron, P. 2021. Proportional participatory budgeting with additive utilities. *Advances in Neural Information Processing Systems*, 34: 12726–12737.

Procaccia, A. D.; and Rosenschein, J. S. 2006. The Distortion of Cardinal Preferences in Voting. In *Proceedings of the 10th International Workshop on Cooperative Information Agents (CIA)*, 317–331.

Rey, S.; Endriss, U.; and de Haan, R. 2020. Designing participatory budgeting mechanisms grounded in judgment aggregation. In *Proceedings of the 17th International Confer-*

ence on Principles of Knowledge Representation and Reasoning, 692–702.

Rey, S.; and Maly, J. 2023. The (computational) social choice take on indivisible participatory budgeting. arXiv:2303.00621.

Rivest, R. L.; and Shen, E. 2010. An optimal single-winner preferential voting system based on game theory. In *Proc. of 3rd International Workshop on Computational Social Choice*, 399–410.

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

Wampler, B.; McNulty, S.; and Touchton, M. 2021. *Participatory budgeting in global perspective*. Oxford University Press.

Wang, R.; Fishkin, J. S.; and Luskin, R. C. 2020. Does Deliberation Increase Public-Spiritedness? *Social Science Quarterly*, 101(6): 2163–2182.