The Fortieth AAAI Conference on Artificial Intelligence (AAAI-26)

Utilitarian Guarantees for the Method of Equal Shares

Anton Baychkov[1], **Markus Brill**[1], **Jannik Peters**[2]
[1]University of Warwick, Coventry, UK
[2]National University of Singapore, Singapore
anton.baychkov@warwick.ac.uk, markus.brill@warwick.ac.uk, peters@nus.edu.sg

Abstract

In recent years, research in Participatory Budgeting (PB) has put a greater emphasis on rules satisfying notions of fairness and proportionality, with the *Method of Equal Shares* (MES) being a prominent example. However, proportionality can come at a cost to the total utilitarian welfare. Our work formalizes this relationship, by deriving minimum utilitarian welfare guarantees for MES for a subclass of satisfaction functions called DNS functions, which includes two of the most popular ways of measuring a voter's utility in the PB setting: considering (1) the total cost of approved projects or (2) the total number of those projects. Our results are parameterized in terms of minimum and maximum project costs, which allows us to improve on the mostly negative results found in prior studies, and reduce to the existing multiwinner guarantee when project costs are equal. We show that our guarantees are asymptotically tight for rules satisfying *Extended Justified Representation up to one project*, showing that no proportional rule can achieve a better utilitarian guarantee than MES.

1 Introduction

Participatory budgeting (PB) is a key topic in computational social choice (Rey, Schmidt, and Maly 2025). PB enables voters to directly vote on how public money is spent. This is often accomplished by allowing the voters to express their preferences over a set of available projects. These preferences are then aggregated using a voting rule selecting a subset of the projects that fits within a pre-defined budget limit (Wampler, McNulty, and Touchton 2021).

In practice, the voting process is usually conducted via *approval ballots*, i.e., each voter can indicate their support for the subset of projects they like, without the need to specify a ranking among them. These ballots are usually aggregated using a simple GREEDY method: projects are processed in order of decreasing welfare-to-cost ratio (i.e., "bang for buck"), where a project's welfare is determined by the total satisfaction it provides to voters. In the most common formulation, where satisfaction is measured by project cost, this reduces to ordering projects by decreasing vote count. When processing projects in this order, each project that fits into the remaining budget is added to the outcome. This method is simple and easy to explain, and greedily maximizes the *utilitarian welfare* of the outcome.

However, GREEDY has a major downside: it is not *proportional*. To see this, consider a scenario where 60% of voters approve only infrastructure projects while 40% approve only leisure projects. GREEDY would spend all the budget on infrastructure projects, leaving the minority completely unrepresented despite their legitimate claim to a proportional share of the budget. Meanwhile, the recently introduced *Method of Equal Shares* (MES) (Peters, Pierczyński, and Skowron 2021) yields a fairer outcome by splitting the total budget equally among voters, and allowing groups of voters to "buy" projects using their allocated funds. While this leads to proportional outcomes — as formalized by variants of the *extended justified representation* (EJR) axiom — MES can produce outcomes with low utilitarian welfare.

This tradeoff between utilitarian and proportionality objectives leads to the question of how much utilitarian welfare can still be guaranteed by proportional rules in general, and by MES in particular. To answer this question, we study *utilitarian guarantees*, i.e., lower bounds on the ratio between the utilitarian welfare provided by a rule on a given instance and the maximal achievable utilitarian welfare for that instance.[1]

For PB, the definition of utilitarian guarantees immediately runs into two fundamental issues. First, it is unclear how voter satisfaction (or welfare) should be defined in this setting, with two differing definitions having emerged in the literature: (i) *cost satisfaction* assumes that a voter's satisfaction is measured in terms of the amount of money spent on projects they approve; whereas (ii) *cardinality satisfaction* assumes that the satisfaction of a voter is given by the number of approved projects in the outcome. Second, the meaningfulness of utilitarian guarantees depends critically on how they are parameterized. The PB setting offers a variety of potential parameters: budget size, project cost distributions, voter preferences, voter turnout, etc. However, the value of guarantees that depend on uncontrollable factors (such as voter turnout or preference diversity) is questionable, as these factors may grow large or vary unpredictably. Instead, meaningful guarantees should be expressed in terms of parameters

---

[1] The inverse of this ratio (for "fair" voting rules) is sometimes referred as the *price of fairness* (e.g., Elkind et al. 2024); a lower ratio corresponds to a higher price.

that PB organizers have control over, such as the total budget and/or the range of possible project costs.

To illustrate these issues, it is instructive to consider the simpler setting of approval-based multiwinner voting (Aziz et al. 2017), where all projects (referred to as candidates) have the same cost and k of them need to be selected. In this setting, the first issue disappears entirely, as cost and cardinality satisfaction become equivalent. Moreover, the *committee size k* is a natural parameter. The best possible utilitarian guarantee of any proportional voting rule is approximately $\frac{2}{\sqrt{k}}$ (Lackner and Skowron 2020; Elkind et al. 2024), with this bound also being achieved by a variant of MES (Brill and Peters 2024).

For PB, however, the picture is less clear. Prior work on utilitarian guarantees for proportional PB rules is extremely limited, with Fairstein et al. (2022) providing the only study to date. However, their analysis yields only weak results: they derived bounds for various voting rules including rules satisfying EJR, but their guarantees are parameterized by the number of voters, making them vanishingly small for realistic PB instances. Moreover, their analysis was restricted to the cardinality satisfaction function (while most real-world instances use cost satisfaction) and produced lower and upper bounds on their utilitarian guarantees that are orders of magnitude apart. This led Rey, Schmidt, and Maly (2025) to comment that "the existing analysis of the price of fairness in PB (Fairstein et al. 2022) is still at a preliminary stage and there is a lot of room for improvement" (page 85).

1.1 Our Contribution

We derive bounds on the utilitarian guarantee for MES, and proportional rules in general, for a whole class of satisfaction functions. To quantify our utilitarian guarantees, we use three parameters: the budget b, the cost of the cheapest project $c_{\min}$, and the cost of the most expensive project $c_{\max}$. In particular, we express our guarantees in terms of $\frac{b}{c_{\min}}$ and $\frac{b}{c_{\max}}$, which bound the number of projects that can be chosen in an exhaustive outcome, and are thus closely related to the committee size k in multiwinner voting. To motivate our parametrization, we show that if project costs are allowed to vary arbitrarily, then we cannot get a meaningful utilitarian guarantee for MES or even GREEDY, for any satisfaction function (Section 3). From a practical perspective, the budget and bounds on permissible project costs can be controlled by the PB organizer, while this is less true for other variables like the number of voters.

We derive positive results for a broad class of satisfaction functions known as *DNS functions* (see Definition 1), introduced by Brill et al. (2023). This class contains cost satisfaction and cardinality satisfaction as special cases. We also show that functions outside this class can lead to arbitrarily bad utilitarian ratios for MES, no matter how much project costs are constrained (Proposition 3).

Our main result is a utilitarian guarantee for MES (for any DNS satisfaction function). This guarantee is derived by comparing the welfare of the outputs of MES and GREEDY (Section 4), and combining that with a knapsack-inspired utilitarian guarantee for GREEDY to find a guarantee for MES (Section 5). We show that this guarantee is asymptotically tight, not just for MES, but for any proportional PB voting rule. Our bounds reduce to the multiwinner guarantees of Brill and Peters (2024) when all project costs are equal.

We also derive a utilitarian guarantee for the GREEDY rule for the curious case that the satisfaction function used to implement the rule differs from the satisfaction function used to measure voter welfare (Section 6).

1.2 Related Work

Utilitarian Guarantees in Multiwinner Voting. Besides the aforementioned work of Fairstein et al. (2022), our work is most closely related to the paper of Brill and Peters (2024), who consider utilitarian guarantees in the multiwinner voting setting. They derive utilitarian guarantees for two proportional rules, showing that MES achieves a guarantee of $\frac{2}{\sqrt{k}} - \frac{2}{k}$, and the Greedy Justified Candidate Rule (GJCR) achieves a guarantee of $\frac{2}{\sqrt{k}} - \frac{1}{k}$. The latter matches the upper bound on the utilitarian guarantee that is achievable by any proportional rule, as proven by Lackner and Skowron (2020) in the paper that initiated the study of utilitarian (and representation) guarantees in multiwinner voting. Moreover, Elkind et al. (2024) studied the impact of requiring the axiom of justified representation (a weakening of EJR) on these guarantees. Utilitarian guarantees were also studied by Dong et al. (2025) and Revel et al. (2025), who analyzed how well one can approximate them together with other objectives.

Utilitarian Guarantees in Participatory Budgeting. Extending these multiwinner results to PB, however, proves challenging. Fairstein et al. (2022) established bounds for various voting rules including those satisfying EJR. However, their analysis suffers from several limitations. First, their lower bound of $\Omega(\frac{1}{n \cdot c_{\max}})$ (with $c_{\min}$ normalized to 1) applies to all rules that exhaust the budget, regardless of proportionality. Second, their construction to prove the upper bound of $O(\frac{1}{\sqrt{n}})$ directly adapts an existing construction for multiwinner voting (Lackner and Skowron 2020, Theorem 9) without exploiting the richer structure available in the PB setting. Notably, both bounds are parameterized by the number of voters n, making them practically meaningless for realistic PB instances: even with $n = 1{,}000$ voters,[2] their lower bound guarantees less than 0.1% of the optimal welfare.

Participatory Budgeting. PB has become a very active research area in computational social choice. We refer to Rey, Schmidt, and Maly (2025) for an up-to-date survey on the topic. In particular, we are closely related to recent papers on proportionality in PB (see, e.g., Aziz, Lee, and Talmon 2018; Peters, Pierczyński, and Skowron 2021; Los, Christoff, and Grossi 2022; Aziz and Lee 2021; Maly 2025; Aziz et al. 2024). Specifically, through our usage of the satisfaction function framework (and DNS functions in particular) to parameterize voting rules and welfare measures, we build on the work of Brill et al. (2023). Finally, we want to highlight the recent work of Papasotiropoulos et al. (2025) who were

---

[2] At time of writing, 68% of the instances in the PB data library *Pabulib* (Faliszewski et al. 2023) have $n \geq 1000$.

motivated by the sometimes poor utilitarian performance of MES to design alternative voting rules, allowing for a better balance between proportionality and other desiderata.

2 Preliminaries

Let P be a set of *projects* and $N = \{1, \dots, n\}$ a set of *voters*. For each voter $i \in N$, we let $A_i \subseteq P$ denote the set of projects approved by this voter. Together, $A = (A_i)_{i \in N}$ forms an *approval profile*. For a project $p \in P$, we let $N_p = \{i \in N : p \in A_i\}$ denote the *supporters of* p, i.e., the set of voters approving p. Finally, for each project $p \in P$ we are given a *cost* $c(p) \in \mathbb{R}_{>0}$. Together with a budget limit $b \in \mathbb{R}_{>0}$, the tuple $I = (P, A, b, c)$ forms an approval-based participatory budgeting (PB) *instance*. We will use the PB instance depicted in Table 1 as a running example in this section. Throughout the paper we will assume that $c(p) \leq b$ for any project $p \in P$. We let $\mathcal{I}$ be the set of all possible PB instances. A *feasible outcome* for a given instance I is any subset $P' \subseteq P$ with $c(P') = \sum_{p \in P'} c(p) \leq b$. We call a feasible outcome P' *exhaustive* if there is no other feasible outcome $P'' \supset P'$.

Multiwinner instances are a special class of PB instances in which projects all have the same cost, $c(p) = c$ for all $p \in P$, and the total budget is $b = kc$. Here, k is also referred to as the *committee size* (i.e., the number of projects to be chosen).

Satisfaction Functions. Throughout the paper, we model voters' utilities using *additive approval-based satisfaction functions* (Talmon and Faliszewski 2019; Brill et al. 2023). In order to argue about satisfaction and voting rules across different instances, we define satisfaction functions in an instance-agnostic way. In particular, we assume that there is a universe $\mathcal{P}$ of all possible projects, and we let $\mathcal{F}(\mathcal{P})$ be the set of all finite subsets of $\mathcal{P}$. A *(global) satisfaction function* $\mu\colon \mathcal{F}(\mathcal{P}) \to \mathbb{R}_{\geq 0}$ maps subsets of projects to the satisfaction received from this subset. In line with Brill et al. (2023), we assume that satisfaction functions are *additive*, i.e., it holds that $\mu(P') = \sum_{p \in P'} \mu(p)$ (with $\mu(p)$ being shorthand for $\mu(\{p\})$). We also assume that the satisfaction of any project p is strictly positive, so that $\mu(P') = 0$ if and only if $P' = \emptyset$. For a given voter i, we let $\mu_i(P') = \mu(P' \cap A_i)$, i.e., the voter only receives satisfaction from the projects they approve. The two most common satisfaction functions considered in the PB literature are the *cost satisfaction function* μ^c, with $\mu^c(p) = c(p)$ and $\mu^c(P') = c(P')$, and the *cardinality satisfaction function* $\mu^\#$, with $\mu^\#(p) = 1$ and $\mu^\#(P') = |P'|$. Note that these are already implicitly defined for all possible projects in an instance-agnostic manner.

Both μ^c and $\mu^\#$ belong to the larger class of *weakly decreasing normalized satisfaction* functions (Brill et al. 2023).

Definition 1. *An additive satisfaction function* μ *is a* weakly decreasing normalized satisfaction (DNS) *function if for any two projects* $p, p' \in P$ *with* $c(p) \leq c(p')$*, the following hold:*

(1) $\mu(p) \leq \mu(p')$*, i.e., the satisfaction of more expensive projects is weakly greater than that of cheaper projects;*

(2) $\frac{\mu(p)}{c(p)} \geq \frac{\mu(p')}{c(p')}$*, i.e., the satisfaction per unit cost of more expensive projects is weakly smaller than that of cheaper projects.*

Project p	$c(p)$	A_1	A_2	A_3	A_4	A_5	A_6	A_7	A_8	A_9	A_{10}
p_1	65	✓	✓	✓	✓	✓	✓				
p_2	60	✓	✓	✓	✓	✓					
p_3	40	✓	✓	✓				✓			
p_4	20						✓	✓	✓		
p_5	20									✓	✓

Table 1: Costs and approval sets for Example 1.

The cost and cardinality satisfaction functions can be viewed as the two extreme examples of DNS functions: condition (1) holds with equality for $\mu^\#$, while condition (2) holds with equality for μ^c. A different example of a DNS function would be $\mu(p) = \sqrt{c(p)}$.

Any DNS function is not just additive, but also *cost-neutral* (Brill et al. 2023): for any two projects p, p' with $c(p) = c(p')$, it holds that $\mu(p) = \mu(p')$. Hence, the satisfaction of a project depends solely on the cost of the project. Note that in multiwinner voting instances, all cost-neutral satisfaction functions are equivalent to each other (up to normalization).

For a given satisfaction function μ, we can define the *utilitarian welfare* of a project p as $uw_\mu(p) = \sum_{i \in N} \mu_i(p) = |N_p| \cdot \mu(p)$. We further let $uw_\mu(P') = \sum_{p \in P'} uw_\mu(p)$ for any $P' \subseteq P$. For a project p and satisfaction function μ, we define the *value* $v_\mu(p)$ as the ratio of p's utilitarian welfare to its cost ("bang per buck"), i.e., $v_\mu(p) = \frac{uw_\mu(p)}{c(p)} = \frac{|N_p| \mu(p)}{c(p)}$. When μ is clear from the context, we will omit the subscript.

Example 1. Consider the PB instance with budget $b = 100$ and $P = \{p_1, p_2, p_3, p_4, p_5\}$, with costs and approval sets given in Table 1. For the cost satisfaction function, we have $\mu^c(p_1) = c(p_1) = 65$, $uw_{\mu^c}(p_1) = 65 \cdot 6 = 390$, and $v_{\mu^c} = 6$. For the cardinality satisfaction function, we have $\mu^\#(p_1) = 1$, $uw_{\mu^\#}(p_1) = |N_{p_1}| = 6$, and $v_{\mu^\#} = 6/65 \approx 0.09$.

Voting Rules. A *voting rule* R is a function mapping each instance I to a nonempty set of outcomes $\mathrm{R}(I)$. Here, we do not just introduce individual voting rules, but families of rules, parameterized by the satisfaction function μ they use. The first two voting rules we consider are *sequential*: each begins with an empty bundle $P_0 = \emptyset$ and then iteratively adds projects to it, selecting project p_k during *stage* k. Let $P_k = \{p_1, \dots, p_k\}$ be the set of projects already chosen by some sequential rule R after some stage k in its execution. We say that at this point an unchosen project $p \in P \setminus P_k$ is *affordable* if $c(P_k) + c(p) \leq b$.

As our first family of sequential voting rules, we consider *greedy voting rules*. Let μ be a satisfaction function. For a given instance, the rule $\textsc{Greedy}_\mu$ iteratively chooses the unchosen affordable project p with the highest value $v_\mu(p)$ until no further project is affordable.

As our second family of rules, we study the *Method of Equal Shares* (Peters, Pierczyński, and Skowron 2021). For a given satisfaction function μ, the rule MES_μ works as follows. Each voter $i \in N$ is assigned a budget b_i, which is initialized to $\frac{b}{n}$. Then, at any given stage, for each $p \in P \setminus W$ compute $\rho(p)$ such that $\sum_{i \in N_p} \min(b_i, \rho(p) \cdot \mu_i(p)) = c(p)$; if no such value exists, we set $\rho(p) = \infty$. If $\rho(p) = \infty$

for every project, MES$_\mu$ stops. Otherwise, it adds a project $p \in \arg\min_{p \in C \setminus W} \rho(p)$ to the outcome and updates the voter budgets to $b_i = b_i - \min(b_i, \rho(p)\mu_i(p))$ for all $i \in N_p$.

Finally, we also consider the *maximum satisfaction rule*, which serves as the benchmark for our utilitarian guarantees. For a given satisfaction function μ, the rule $\textsc{MaxSat}_\mu$ selects all feasible outcomes P' maximizing $\sum_{i \in N} \mu_i(P')$. In multiwinner instances, $\textsc{MaxSat}$ and $\textsc{Greedy}$ produce the same outcome.

One well-known shortcoming of MES is that it is not exhaustive. To make it exhaustive, we consider *completion rules*. A completion rule takes a feasible outcome P' and supplements it with additional projects, outputting a feasible and exhaustive superset $P'' \supseteq P'$. We let $\overline{\text{MES}_\mu}$ be MES$_\mu$ *completed* by the $\textsc{Greedy}_\mu$ rule. This refers to running MES$_\mu$ on the full instance, and then, once MES$_\mu$ terminates, running $\textsc{Greedy}_\mu$ with the remaining projects and budget.

For Example 1, the outcomes of our voting rules are as follows (employing cost satisfaction): $\textsc{Greedy}_{\mu^c}$ selects $P_G = \{p_1, p_4\}$ with $uw(P_G) = 450$; $\textsc{MaxSat}_{\mu^c}$ selects $P^* = \{p_2, p_3\}$ with $uw(P^*) = 460$; and MES$_{\mu^c}$ selects $P_M = \{p_3, p_4\}$ with $uw(P_M) = 220$. $\overline{\text{MES}_{\mu^c}}$ additionally selects p_5 and achieves $uw(\{p_3, p_4, p_5\}) = 260$.

Proportionality. MES is a rule that produces a proportional outcome, in the sense that it provides an appropriate level of representation to sufficiently cohesive groups. Following Peters, Pierczyński, and Skowron (2021), we define the notion of cohesiveness below.

Definition 2. *Let $T \subseteq P$. A group $N' \subseteq N$ of voters is* T-cohesive *if and only if $T \subseteq \bigcap_{i \in N'} A_i$ and $c(T) \leq \frac{|N'|}{n} b$.*

Using the notion of cohesiveness, we can define the following proportionality axiom. Just like our voting rules, the axiom is parameterized by a satisfaction function.

Definition 3. *An outcome $P' \subseteq P$ satisfies* Extended Justified Representation up to one project with respect to μ *(EJR1$_\mu$) if for every T-cohesive group N', either $T \subseteq P'$ or there exists a voter $i \in N'$ and a project $p \in A_i \cap (P \setminus P')$ such that $\mu_i(P') + \mu_i(p) > \mu_i(T)$.*

We say that a rule R satisfies EJR1$_\mu$ if the outcome of R satisfies EJR1$_\mu$ for every instance $I \in \mathcal{I}$. MES$_\mu$ satisfies EJR1$_\mu$ for any satisfaction function μ (Peters, Pierczyński, and Skowron 2021).[3] This is not the case for $\textsc{Greedy}_\mu$; for example, $\textsc{Greedy}_{\mu^c}$ violates EJR1$_{\mu^c}$ in Example 1 due to voter group $\{9, 10\}$, which is $\{p_5\}$-cohesive.

Utilitarian Guarantees. Fix a satisfaction function μ and let R be a voting rule. Let $P_R = \text{R}(I)$ be the outcome of R for a particular PB instance $I = (P, A, b, c)$, and let $P^* = \textsc{MaxSat}_\mu(I)$. We say that R has a *utilitarian ratio* (w.r.t. μ) of $r \in [0, 1]$ for instance I if $\frac{uw_\mu(P_R)}{uw_\mu(P^*)} = r$.

In Example 1, the utilitarian ratio (w.r.t. μ^c) of $\textsc{Greedy}_{\mu^c}$ is $\frac{450}{460} \approx 0.98$ and that of $\overline{\text{MES}_{\mu^c}}$ is $\frac{260}{460} \approx 0.57$.

Now let $g\colon \mathcal{I} \to [0, 1]$ be a function that can depend on properties of a particular PB instance $I \in \mathcal{I}$. We say that R has a *utilitarian guarantee* of g if, for any $I \in \mathcal{I}$, the outcome $P_R = \text{R}(I)$ has a utilitarian ratio of at least $g(I)$.

In principle, g may depend on any properties of the instance. However, the more parameters it depends on, the less meaningful it will be. At one extreme, we could set the guarantee of rule R to be the utilitarian ratio of R for every instance. At the other extreme, if we do not allow g to depend on any parameters then we cannot do better than $g = 0$.[4]

We will formulate our utilitarian guarantees using only the following properties of a PB instance $I = (P, A, b, c)$:

- the instance budget $b(I) = b$;
- the *minimum project cost* $c_{\min}(I) = \min_{p \in P} c(p)$; and
- the *maximum project cost* $c_{\max}(I) = \max_{p \in P} c(p)$.

In fact, all of our results will actually depend on just two cost-based parameters, the *normalized minimum* and *maximum project costs*: $\frac{c_{\min}(I)}{b(I)}$ and $\frac{c_{\max}(I)}{b(I)}$. We can think of the reciprocals of these, $k_1(I) = \frac{b(I)}{c_{\max}(I)}$ and $k_2(I) = \frac{b(I)}{c_{\min}(I)}$, as providing lower and upper bounds (up to rounding) on the number of projects that can be contained in any exhaustive outcome. These generalize the multiwinner concept of committee size to the PB setting, with k_1 and k_2 corresponding to *minimum* and *maximum committee size*, respectively.

Similarly, we define $\mu_{\min}(I) = \min_{p \in P} \mu(p)$ and $\mu_{\max}(I) = \max_{p \in P} \mu(p)$ as the minimum and maximum satisfaction of any project in a given instance I. In our notation, we often omit I when the instance is clear from context.

3 Negative Results

As a warm up, we demonstrate that certain natural assumptions are necessary to obtain meaningful utilitarian guarantees. In particular, the negative results in this section motivate both (i) our focus on costs as parameters for the guarantees and (ii) our restriction to DNS satisfaction functions.

In this section, we let PROP$_\mu$ be a rule satisfying EJR1$_\mu$. For instance, PROP$_\mu$ could be MES$_\mu$ with any completion rule. We begin by showing that, for any DNS satisfaction function μ, we can construct instances such that PROP$_\mu$, and even $\textsc{Greedy}_\mu$, achieve an arbitrarily low utilitarian ratio.

We first consider the case of satisfaction functions that are non-trivially bounded from below, like the cardinality satisfaction function $\mu^{\#}$.

Proposition 1. *Let μ be a DNS function and $\varepsilon > 0$ a constant such that $\mu(p) > \varepsilon$ for all $p \in \mathcal{P}$. Then,* PROP$_\mu$ *and* $\textsc{Greedy}_\mu$ *have a utilitarian guarantee of at most $\frac{1}{n-1}$.*

Proof. Consider a PB instance with $n \geq 3$ voters and a project set P containing two projects: project p_1 with $N_{p_1} = N \setminus \{1\}$ and $c(p_1) = b$, and project p_2 with $N_{p_2} = \{1\}$

---

[3] For $\mu = \mu^{\#}$, EJR1$_\mu$ is equivalent to the EJR axiom used by Fairstein et al. (2022). If μ is a DNS function, the outcome of MES$_\mu$ even satisfies the stronger *EJR$_\mu$ up to any project* (Brill et al. 2023). For the cost satisfaction function μ^c, MES$_{\mu^c}$ satisfies the even stronger *EJR+$_{\mu^c}$ up to any project* (Brill and Peters 2023).

[4] To see this, consider the multiwinner guarantee upper bound of $\frac{2}{\sqrt{k}} - \frac{1}{k}$ (Lackner and Skowron 2020) for arbitrarily large k.

and $c(p_2) = \frac{\varepsilon b}{n\mu(p_1)}$. This instance has two exhaustive outcomes, $\{p_1\}$ and $\{p_2\}$, as $c(p_1) + c(p_2) > b$. Since μ is a DNS function, we know that p_1 has higher utilitarian welfare than p_2. Hence, MaxSat$_\mu$ chooses the outcome $\{p_1\}$. Observe that $v(p_1) = \frac{(n-1)\mu(p_1)}{c(p_1)} = \frac{(n-1)\mu(p_1)}{b} < \frac{n\mu(p_1)}{b}$ and $v(p_2) = \frac{\mu(p_2)}{c(p_2)} > \frac{\varepsilon}{\frac{\varepsilon b}{n\mu(p_1)}} = \frac{n\mu(p_1)}{b}$. Thus, $v(p_1) < v(p_2)$ and Greedy$_\mu$ selects the outcome $\{p_2\}$. As $\frac{\varepsilon}{\mu(p_1)} < 1$, we know that $c(p_2) < \frac{b}{n}$ and can thus be afforded by voter 1. Thus, $\{1\}$ is a $\{p_2\}$-cohesive group, and as voter 1 approves no other projects, we get that PROP$_\mu$ must select $\{p_2\}$. Using DNS assumption (1), this implies that the utilitarian ratio of Greedy$_\mu$ and PROP$_\mu$ is $\frac{uw(p_2)}{uw(p_1)} = \frac{\mu(p_2)}{(n-1)\mu(p_1)} \leq \frac{1}{n-1}$. □

Therefore, we can choose an instance with a sufficiently large number of voters to make the utilitarian ratio of Greedy and PROP arbitrarily bad. Notably, Proposition 1 tightens the upper bound on the utilitarian guarantee (when parameterized only by n only) for proportional rules with $\mu = \mu^\#$ from $O(1/\sqrt{n})$ (Fairstein et al. 2022) to $O(1/n)$.

Next, we consider the case of satisfaction functions that can get arbitrarily small, like the cost satisfaction function μ^c.

Proposition 2. *Let μ be a DNS satisfaction function such that for any $\varepsilon > 0$ there exists a project $p \in \mathcal{P}$ with $\mu(p) < \varepsilon$. Then, for any $\delta > 0$,* Greedy$_\mu$ *and* PROP$_\mu$ *do not achieve a utilitarian guarantee of δ.*

Propositions 1 and 2 combine to show that utilitarian guarantees for DNS satisfaction functions for both Greedy and PROP can be arbitrarily bad. It's worth emphasizing the fact that these results hold for Greedy — a rule whose explicit goal is to (greedily) maximize utilitarian welfare. This mirrors similar results for the greedy knapsack algorithm (see e.g. Kellerer, Pferschy, and Pisinger (2004)). The instances constructed in both proofs involve a very small project preventing the purchase of a larger one, which is reminiscent of the "Helenka Paradox" discussed by Papasotiropoulos et al. (2025). This suggests that we could potentially obtain better results for instances where project costs are closer together.

We now motivate the idea of constraining ourselves to DNS satisfaction functions by showing that no matter how much we limit project costs (without making them all equal), we can find a non-DNS satisfaction function μ that would make any proportional rule achieve an arbitrarily low utilitarian ratio.

Proposition 3. *Fix k_1, k_2 with $1 \leq k_1 < k_2$ and $\varepsilon > 0$. There exists a PB instance I with $c_{\max}(I) = \frac{b}{k_1}$ and $c_{\min}(I) = \frac{b}{k_2}$, and a (non-DNS) satisfaction function μ, such that the utilitarian ratio achieved by* PROP$_\mu$ *is at most ε.*

Together, Propositions 1 to 3 motivate the need for using costs as utilitarian guarantee parameters, and for constraining ourselves to DNS satisfaction functions. In the next sections, we will see that combining these allows us to obtain meaningful utilitarian guarantees for proportional rules.

4 Comparing the Welfare of MES and Greedy

In this section, we compare the utilitarian welfare generated by MES$_\mu$ to that generated by Greedy$_\mu$, and bound the ratio of these from below. We can think of this as a *comparative* utilitarian guarantee. We will use our results here to derive a utilitarian guarantee for MES$_\mu$ in Section 5.

Theorem 1. *Let μ be a DNS satisfaction function, and consider a PB instance I. Let P_G and P_M be the outcomes of* Greedy$_\mu$ *and* $\overline{\text{MES}}_\mu$*, respectively, for this instance. Then,*

$$\frac{uw(P_M)}{uw(P_G)} \geq 2\sqrt{\frac{c_{\min}(I)}{b(I)}} - \frac{c_{\min}(I) + c_{\max}(I)}{b(I)}.$$

This guarantee mirrors the $\frac{2}{\sqrt{k}} - \frac{2}{k}$ guarantee for multiwinner voting (Brill and Peters 2024). Similarly to the proof of Brill and Peters (2024), our proof utilizes the fact that Greedy$_\mu$ and MES$_\mu$ will select the same projects for a while, before diverging once voters start running out of their budgets during the execution of MES$_\mu$. In order to determine how early this divergence might occur, we introduce the concept of *budget-limited* voters.

Definition 4. *Let P_k be the set of projects selected by* MES *after stage k of its execution. We say that voter i is* budget-limited *for an unchosen project $p \in A_i \setminus P_k$, if p is affordable, but splitting its cost equally among its supporters would result in i exceeding their personal budget. Formally, if b_i is the remaining budget of voter i after stage k, voter i is budget-limited for p if p is affordable but $c(p)/|N_p| > b_i$.*

We show that MES$_\mu$ selects the same projects as Greedy$_\mu$ until the stage in its execution where, for the first time, a voter is budget-limited for the affordable, unchosen project with highest value.

Lemma 1. *Let p_j be the project chosen in stage j by* Greedy$_\mu$ *and p'_j be the project chosen in stage j by* MES$_\mu$*. If for each $j \in \{1, \ldots, k\}$, no voter is budget-limited for p_j at stage j of* MES$_\mu$*, then $p'_k = p_k$.*

Using this, we can find the first MES$_\mu$ stage during which a voter would be budget-limited for the project chosen by Greedy at that stage. We call this stage $i + 1$.[5] Using this stage as a threshold, we will consider three sets of projects and compare the utilitarian welfare generated by these:

(1) the commonly chosen projects until stage i;

(2) the projects chosen by Greedy$_\mu$ after stage i; and

(3) the projects chosen by $\overline{\text{MES}_\mu}$ after stage i.

Together, (1) and (2) form P_G and (1) and (3) form P_M.

In order to help bound the utilitarian welfare of set (3), we can find a lower bound for the value of any project selected by MES$_\mu$. Note that this does not restrict the value of projects selected by any completion rule we use to supplement the outcome of MES$_\mu$ — we will consider those separately.

Lemma 2. *Any project p chosen by* MES$_\mu$ *has a value of at least $v_\mu(p) \geq \frac{n}{b}\mu_{\min}$.*

---

[5] By defining this threshold stage in a more nuanced way than in the proof of Theorem 10 of Brill and Peters (2024), we fix a minor error in their proof.

Moreover, we use the DNS assumption on μ to bound the ratio of satisfaction to cost in our instance.

Lemma 3. *Let μ be a DNS satisfaction function and I be a PB instance with project set P. Then, for any project $p \in P$, it holds that $\frac{\mu_{\min}(I)}{c_{\min}(I)} \geq \frac{\mu(p)}{c(p)} \geq \frac{\mu_{\max}(I)}{c_{\max}(I)}$.*

We are now ready to prove Theorem 1.

Proof of Theorem 1. We compare the projects chosen by GREEDY_μ to the projects chosen by $\overline{\text{MES}}_\mu$. Let

$$P_G = \{p_1, \dots, p_g\} \quad \text{and} \quad P_M = \{p'_1, \dots, p'_m\}$$

be the sets of projects chosen by GREEDY_μ and $\overline{\text{MES}}_\mu$, respectively, indexed in the order they were chosen by those methods. If $P_G = P_M$, then $\frac{uw(P_M)}{uw(P_G)} = 1$ and we are done, so we will assume that $P_G \neq P_M$. In particular, this means that p'_1 was chosen by MES_μ rather than by its completion.

Using Lemma 1, we know that as long as no voters are budget-limited for the affordable, unchosen project with highest value, the MES_μ part of $\overline{\text{MES}}_\mu$ selects the same projects as GREEDY_μ. As $P_G \neq P_M$ we can let $p_{i+1} \in P_G$ be the first project chosen by GREEDY_μ where a voter would be budget limited for it after $\{p_1, \dots, p_i\}$ were selected by MES_μ. Note that $i = 0$ is possible, i.e., p_{i+1} could be the very first project chosen. For notational convenience, let $v^* = v(p_{i+1})$ and $\mu^* = \mu(p_{i+1})$. We observe that $p_j = p'_j$ for $1 \leq j \leq i$ and let the total cost of those projects be $\zeta = \sum_{j=1}^{i} c(p_j)$.

In order to compare the utilitarian welfare generated by GREEDY_μ and $\overline{\text{MES}}_\mu$, we will consider the welfare of the following subsets of projects: (1) The set $\{p_1, \dots, p_i\}$ of projects that are initially chosen by both GREEDY_μ and $\overline{\text{MES}}_\mu$; (2) the set $\{p_{i+1}, \dots, p_g\}$ of projects chosen by GREEDY_μ after p_i; and (3) the set $\{p'_{i+1}, \dots, p'_m\}$ of projects chosen by $\overline{\text{MES}}_\mu$ after p_i.

We will prove the following bounds on the utilitarian welfare provided by these three projects sets:

$$uw(\{p_1, \dots, p_i\}) \geq \zeta v^* \tag{1}$$

$$uw(\{p_{i+1}, \dots, p_g\}) \leq (b - \zeta) v^* \tag{2}$$

$$uw(\{p'_{i+1}, \dots, p'_m\}) \geq [b - \zeta - c_{\max}] \frac{n}{b} \mu_{\min} \tag{3}$$

From our definition of value, we know that for any set of projects $P' \subseteq P$, we have $uw(P') = \sum_{p \in P'} c(p) v(p)$.

For (1), we observe that, by definition of GREEDY_μ, it holds that $v(p) \geq v^*$ for any $p \in \{p_1, \dots, p_i\}$ as any such project was picked before p_i. Hence, get $uw(\{p_1, \dots, p_i\}) = \sum_{j=1}^{i} c(p_j) v(p_j) \geq \sum_{j=1}^{i} c(p_j) v^* = \zeta v^*$.

For (2), we observe that any project $p \in \{p_{i+1}, \dots, p_g\}$ must have a weakly lower value than p_{i+1}. Since, moreover, $c(\{p_{i+1}, \dots, p_g\})$ is upper bounded by $b - \zeta$, we get that $uw(\{p_{i+1}, \dots, p_n\}) = \sum_{j=i+1}^{n} c(p_j) v(p_j) \leq \sum_{j=i+1}^{n} c(p_j) v^* \leq (b - \zeta) v^*$.

For (3), we construct a subset $P'_M \subseteq \{p'_{i+1}, \dots, p'_m\}$ s.t.

$$c(P'_M) \geq b - \zeta - c_{\max}, \text{ and} \tag{3a}$$

$$v(p) \geq \frac{n}{b} \mu_{\min} \text{ for each project } p \in P'_M. \tag{3b}$$

Together, (3a) and (3b) imply $uw(\{p'_{i+1}, \dots, p'_m\}) \geq uw(P'_M) \geq [b - \zeta - c_{\max}] \frac{n}{b} \mu_{\min}$.

The completion step of $\overline{\text{MES}}_\mu$ considers projects that were not already chosen by MES_μ, in order of decreasing value, and selects them as long as they are affordable. Let p^* be the first unaffordable project that the completion step considers, and call the project chosen by $\overline{\text{MES}}_\mu$ immediately before this point p'_j. Clearly $c(\{p'_1, \dots, p'_j\}) + c(p^*) > b$, else p^* would have been affordable. We define $P'_M = \{p'_{i+1}, \dots, p'_j\}$, noting that $P'_M = \emptyset$ in the case $j = i$. Then, $c(P'_M) > b - \zeta - c(p^*) \geq b - \zeta - c_{\max}$, showing (3a).

For (3b), we consider $v(p)$ for any $p \in P'_M$. We have two cases: either p was chosen by MES_μ or the greedy completion. If p was chosen by MES_μ then we know from Lemma 2 that $v(p) \geq \frac{n}{b} \mu_{\min}$.

Let $P' \subseteq P$ be the set of projects with $v(p) \geq \frac{n}{b} \mu_{\min}$. Suppose that p was chosen by the greedy completion, and, for a contradiction, assume that $v(p^*) < \frac{n}{b} \mu_{\min}$, which means that $\overline{\text{MES}}_\mu$ selected every project p with $v(p) \geq \frac{n}{b} \mu_{\min}$. Thus, $P' \subseteq \{p'_1, \dots, p'_j\} \subseteq P_M$. Furthermore, every project in $P_M \setminus P'$ must have been chosen by the greedy completion by Lemma 2. As $c(P') \leq c(P_M) \leq b$, the outcome of GREEDY_μ, P_G, must also contain P', as GREEDY_μ selects projects in decreasing order of value. Thus $P_G = P_M$, which contradicts our earlier assumption that $P_G \neq P_M$, and we can conclude that $v(p) \geq v(p^*) \geq \frac{n}{b} \mu_{\min}$, showing (3b).

In order to compute the tradeoff between GREEDY_μ and $\overline{\text{MES}}_\mu$ using (1), (2), and (3), we determine v^*, i.e., the value of p_{i+1}. To do so, we first define the parameter $\alpha = \frac{|N_{p_{i+1}}| b}{n c(p_{i+1})}$, representing how many times over the supporters of p_{i+1} could buy p_{i+1} with the budgets they are provided at the start of the execution of MES_μ. Using the definition of α and v^*, we can rewrite the value of p_{i+1} as $v^* = \frac{|N_{p_{i+1}}| \mu^*}{c(p_{i+1})} = \alpha \mu^* \frac{n}{b}$.

We distinguish three cases, based on the value of α. In each of the three cases, we show that the bound $\frac{uw(P_M)}{uw(P_G)} \geq 2\sqrt{\frac{c_{\min}}{b}} - \frac{c_{\min} + c_{\max}}{b}$ holds.

Case 1: $\alpha < \frac{\mu_{\min}}{\mu^*} \leq 1$. We can find that every affordable project chosen by $\overline{\text{MES}}_\mu$ after p_i was chosen by the GREEDY_μ completion and thus $P_G = P_M$.

Case 2: $\alpha \geq 1$. Using the fact that some voter was budget-limited for p_{i+1} in step $i + 1$ of MES_μ we can bound the cost of $\{p_1, \dots, p_i\}$ with $\zeta \geq (\alpha - 1) c_{\min} \frac{\mu^*}{\mu_{\min}}$. We then combine this with our earlier results, and our assumption that μ is a DNS function, to derive the bound above.

Case 3: $\frac{\mu_{\min}}{\mu^*} \leq \alpha < 1$. For this case we again combine our earlier results with the DNS assumptions. ☐

5 The Utilitarian Guarantee of MES

In this section, we combine a knapsack-inspired utilitarian guarantee for GREEDY with a modified version of Theorem 1 to derive a utilitarian guarantee for MES.

Using existing results from the knapsack literature (e.g., Kellerer, Pferschy, and Pisinger 2004), we can derive the following utilitarian guarantee for the greedy rule.

Proposition 4. *Let μ be a (possibly non-DNS) satisfaction function. Then* $\textsc{Greedy}_\mu$ *has a utilitarian guarantee of* $\frac{b-c_{\max}}{b}$*, and this guarantee is asymptotically tight.*

One simple way to derive a utilitarian guarantee for MES_μ would be to directly combine the results of Theorem 1 and Proposition 4, obtaining a guarantee of $(2\sqrt{\frac{c_{\min}}{b}} - \frac{c_{min}+c_{\max}}{b})(\frac{b-c_{\max}}{b})$. However, in the PB setting, different rules do not necessarily spend the same proportion of the budget (or spend it efficiently). By simply multiplying our two bounds, we would be double-counting this inefficiency. By accounting for this more carefully, we can derive a utilitarian guarantee for $\overline{\text{MES}_\mu}$ that coincides with the comparative guarantee from Theorem 1.

Theorem 2. *Let μ be a DNS satisfaction function. Then,* $\overline{\text{MES}_\mu}$ *has a utilitarian guarantee of* $2\sqrt{\frac{c_{\min}}{b}} - \frac{c_{\min}+c_{\max}}{b}$.

Proof sketch. We first compare the output of $\overline{\text{MES}_\mu}$ to a truncated output of $\textsc{Greedy}_\mu$, considering the subset of projects purchased by $\textsc{Greedy}_\mu$ using the first $b - c_{\max}$ units of the budget that it spends. We lower bound the ratio of the utilitarian welfare of these two outcomes by $(2\sqrt{\frac{c_{\min}}{b}} - \frac{c_{\max}+c_{\min}}{b})\frac{b}{b-c_{\max}}$. The extra $\frac{b}{b-c_{\max}}$ factor reflects the decrease in total cost of the truncated greedy outcome compared to the original outcome of $\textsc{Greedy}_\mu$. We then compare the truncated $\textsc{Greedy}_\mu$ outcome to the output of $\textsc{MaxSat}_\mu$, and find that this outcome has the same utilitarian ratio as guaranteed by Proposition 4: $\frac{b-c_{\max}}{b}$. Combining these bounds, we obtain a guarantee of $2\sqrt{\frac{c_{\min}}{b}} - \frac{c_{\min}+c_{\max}}{b}$. $\square$

The result of Theorem 2, like all of the results in this paper, can be written in terms of just two instance parameters — the minimum and maximum committee size of the instance: $2\sqrt{\frac{c_{\min}}{b}} - \frac{c_{\min}+c_{\max}}{b} = 2\sqrt{\frac{1}{k_2}} - \frac{1}{k_2} - \frac{1}{k_1}$.

In order to give more context to our utilitarian guarantees, we consider the following simple example.

Example 2. Let μ be any DNS function. Consider a PB instance class $\mathcal{I}^* \subseteq \mathcal{I}$ with $b(I) = \$1,000,000$, $c_{\min}(I) = \$10,000$ and $c_{\max}(I) = \$30,000$. For this class, MES_μ has a utilitarian guarantee of $2\sqrt{0.01} - 0.01 - 0.03 = 0.16$. $\mathcal{I}^*$ is "similar" to a class of multiwinner elections with $k = \frac{b}{c_{\min}} = 100$. (We could think of $\mathcal{I}^*$ as allowing some candidates to be bigger, taking up any — potentially non-integral — number of seats from 1 to 3.) For this election, the utilitarian guarantee of MES is 0.18 and that of GJCR (another proportional rule) is 0.19, using the bounds of Brill and Peters (2024).

Meanwhile, consider an MES utilitarian guarantee that is a function of n, the number of voters. We know that for any DNS function μ, this guarantee is bounded from above by $\frac{1}{n-1}$ (Propositions 1 and 2). For instances with $n \geq 1000$, this provides a utilitarian guarantee of at most 0.001.

In the multiwinner setting, the MES utilitarian guarantee is almost tight, with Lackner and Skowron (2020) deriving an upper bound of $\frac{2}{\lfloor\sqrt{k}\rfloor} - \frac{1}{k}$ for proportional rules. In the PB setting, we show that the MES utilitarian guarantee from Theorem 2 is asymptotically tight — not just for MES, but for any rule satisfying EJR1 —, at least for cost satisfaction.

Proposition 5. *Let $\mu = \mu^c$ be the cost satisfaction function and let* PROP_μ *be a rule satisfying EJR1*$_\mu$*. For each $k_1, k_2 \in \mathbb{N}$ with $k_2 \geq k_1$, there exists a PB instance I with* $c_{\min}(I) = \frac{b(I)}{k_2}$ *and* $c_{\max}(I) = \frac{b(I)}{k_1}$ *such that* PROP_μ *has a utilitarian ratio of at most* $\frac{2}{\lfloor\sqrt{b/c_{\min}}\rfloor} - \frac{c_{\min}+xc_{\max}}{b}$*, where* $x = \lfloor\sqrt{k_2}\rfloor\frac{c_{\min}}{c_{\max}} - \lfloor\lfloor\sqrt{k_2}\rfloor\frac{c_{\min}}{c_{\max}}\rfloor$.

When $c_{\min} = c_{\max}$, we get $x = 0$ and the utilitarian ratio above reduces to the multiwinner bound.

If we assume that $\sqrt{k_2} \in \mathbb{N}$, we get an upper bound of $2\sqrt{\frac{c_{\min}}{b}} - \frac{c_{\min}+xc_{\max}}{b}$. We note that $x < 1$ and that x approaches 1 for appropriate choices of k_1 and k_2. For instance, consider the sequences $k_1(a) = a - 1$ and $k_2(a) = a^2$ for $a \in \mathbb{N}$, resulting in $x = \frac{a-1}{a} \to 1$. For $x \to 1$, the utilitarian ratio in the proof above approaches the MES utilitarian guarantee of $2\sqrt{\frac{c_{\min}}{b}} - \frac{c_{\min}+c_{\max}}{b}$ from Theorem 2.

6 Guarantees with Incorrect Satisfaction

Finally, we consider the impact on welfare from using an "incorrect" satisfaction function for the $\textsc{Greedy}$ rule.

Theorem 3. *Let μ be the actual DNS satisfaction function, and let μ' be another DNS satisfaction function. Then, the utilitarian guarantee (as measured in terms of μ) of* $\textsc{Greedy}_{\mu'}$ *is* $\frac{b-c_{\max}}{b} \times \frac{c_{\min}}{c_{\max}}$*. This bound is asymptotically tight.*

This guarantee is made up of two terms. The first, $\frac{b-c_{\max}}{b}$, captures the fact that $\textsc{Greedy}$ is only guaranteed to pick the most efficient set of projects for the first $b - c_{\max}$ dollars it spends. The second, $\frac{c_{\min}}{c_{\max}}$, can be thought of as the distortion of the instance, dictating how far DNS satisfaction functions can diverge. We can show that this bound is tight by constructing an example in which the actual satisfaction function is μ^c and our $\textsc{Greedy}$ rule is parametrised with $\mu^\#$.

7 Conclusion and Future Work

We studied the utilitarian guarantee of the Method of Equal Shares. In particular, we obtained a utilitarian guarantee of $2\sqrt{\frac{c_{\min}}{b}} - \frac{c_{\max}+c_{\min}}{b}$ (with $c_{\min}$ and $c_{\max}$ being the minimum and maximum cost of a project and b the budget limit in a given instance) when the satisfaction of voters is measured by a DNS satisfaction function. We further showed that this bound is tight for rules satisfying EJR1, and thus for MES.

There are multiple ways one could move forward from here. Firstly, in a very recent work, Papasotiropoulos et al. (2025) introduced "MES with bounded overspending" in an attempt to fix some shortcomings of MES. Does this rule behave better than MES with regard to its utilitarian guarantee or are there perhaps other rules which allow to trade-off between proportionality and welfare?

Secondly, our utilitarian guarantees are worst-case guarantees, and thus are unlikely to be representative of real-world scenarios. For instance, the experiments of Bredereck et al. (2019), Elkind et al. (2024), and Revel et al. (2025) indicate that proportional voting rules behave significantly better than what the worst-case guarantee would suggest. It would be interesting to understand to what extent empirical performance depends on the parameters of the instance, such as those used in our utilitarian guarantee.

Acknowledgments

This work was supported by a *Structural Democracy Fellowship* through the Brooks School of Public Policy at Cornell University, an *EPSRC Doctoral Training Partnership* award, and by the Singapore Ministry of Education under grant number MOE-T2EP20221-0001.

References

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.; and Lee, B. E. 2021. Proportionally Representative Participatory Budgeting with Ordinal Preferences. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, 5110–5118. AAAI Press.

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 Multiagent Systems (AAMAS)*, 23–31. IFAAMAS.

Aziz, H.; Lu, X.; Suzuki, M.; Vollen, J.; and Walsh, T. 2024. Fair Lotteries for Participatory Budgeting. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*, 9469–9476.

Bredereck, R.; Faliszewski, P.; Niedermeier, R.; and Kaczmarczyk, A. 2019. An Experimental View on Committees Providing Justified Representation. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, 109–115.

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. AAAI Press.

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 (ACM-EC)*, 301. ACM Press. Full version arXiv:2302.01989 [cs.GT].

Brill, M.; and Peters, J. 2024. Completing Priceable Committees: Utilitarian and Representation Guarantees for Proportional Multiwinner Voting. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*, 9528–9536. AAAI Press.

Dong, C.; Bullinger, M.; Wąs, T.; Birnbaum, L.; and Elkind, E. 2025. Selecting Interlacing Committees. In *Proceedings of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 630 – 638.

Elkind, E.; Faliszewski, P.; Igarashi, A.; Manurangsi, P.; Schmidt-Kraepelin, U.; and Suksompong, W. 2024. The Price of Justified Representation. *ACM Transactions on Economics and Computation*, 12(3): 11:1–11:27.

Fairstein, R.; Vilenchik, D.; Meir, R.; and Gal, K. 2022. Welfare vs. Representation in Participatory Budgeting. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 409–417.

Faliszewski, P.; Fils, J.; Peters, D.; Pierczyński, G.; Skowron, P.; Stolicki, D.; Szufa, S.; and Talmon, N. 2023. Participatory Budgeting: Data, Tools, and Analysis. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*, 2667–2674.

Kellerer, H.; Pferschy, U.; and Pisinger, D. 2004. *Knapsack Problems*. Springer.

Lackner, M.; and Skowron, P. 2020. Utilitarian Welfare and Representation Guarantees of Approval-Based Multiwinner Rules. *Artificial Intelligence*, 288: 103366.

Los, M.; Christoff, Z.; and Grossi, D. 2022. Proportional Budget Allocations: A Systematization. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*, 398–404.

Maly, J. 2025. The core in Participatory Budgeting can be empty. *Economics Letters*, 254: 112472.

Papasotiropoulos, G.; Pishbin, S. Z.; Skibski, O.; Skowron, P.; and Wąs, T. 2025. Method of Equal Shares with Bounded Overspending. In *Proceedings of the 26th ACM Conference on Economics and Computation (ACM-EC)*, 841–868.

Peters, D.; Pierczyński, G.; and Skowron, P. 2021. Proportional Participatory Budgeting with Additive Utilities. In *Proceedings of the 34th Conference on Neural Information Processing Systems (NeurIPS)*, 12726–12737.

Revel, M.; Milli, S.; Lu, T.; Watson-Daniels, J.; and Nickel, M. 2025. Representative Ranking for Deliberation in the Public Sphere. In *Proceedings of the 42nd International Conference on Machine Learning (ICML)*, 51583–51613.

Rey, S.; Schmidt, F.; and Maly, J. 2025. The (Computational) Social Choice Take on Indivisible Participatory Budgeting. Technical report, arXiv:2303.00621 [cs.GT].

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.