informs
http://pubsonline.informs.org/journal/mnsc

MANAGEMENT SCIENCE
*Articles in Advance*, pp. 1–15
ISSN 0025-1909 (print), ISSN 1526-5501 (online)

# Preference Elicitation for Participatory Budgeting

**Gerdus Benadè,[a] Swaprava Nath,[b] Ariel D. Procaccia,[c] Nisarg Shah[d]**

[a] Questrom School of Business, Boston University, Boston, Massachusetts 02215; [b] Department of Computer Science and Engineering, Indian Institute of Technology Kanpur, 208016 Kanpur, India; [c] School of Engineering and Applied Sciences, Harvard University, Cambridge, Massachusetts 02138; [d] Department of Computer Science, University of Toronto, Toronto, Ontario M5S 2E4, Canada

**Contact:** benade@bu.edu, https://orcid.org/0000-0002-0024-0026 (GB); swaprava@cse.iitk.ac.in, https://orcid.org/0000-0001-8309-5006 (SN); arielpro@seas.harvard.edu (ADP); nisarg@cs.toronto.edu (NS)

**Received:** September 4, 2018
**Revised:** December 5, 2019
**Accepted:** March 8, 2020
**Published Online in Articles in Advance:** ■■ ■■, 2020

https://doi.org/10.1287/mnsc.2020.3666



**Abstract.** Participatory budgeting enables the allocation of public funds by collecting and aggregating individual preferences. It has already had a sizable real-world impact, but making the most of this new paradigm requires rethinking some of the basics of computational social choice, including the very way in which individuals express their preferences. We attempt to maximize social welfare by using observed votes as proxies for voters' unknown underlying utilities, and analytically compare four preference elicitation methods: knapsack votes, rankings by value or value for money, and threshold approval votes. We find that threshold approval voting is qualitatively superior, and also performs well in experiments using data from real participatory budgeting elections.

**History:** Accepted by Yan Chen, decision analysis.
**Funding:** This work was supported by the Office of Naval Research [Grant N00014-16-1-3075], the National Science Foundation [Grants CCF-1215883, CCF-1525932, and IIS-1350598], the Alfred P. Sloan Foundation through a Sloan Research Fellowship, the Fulbright Association through a Fulbright-Nehru postdoctoral fellowship, and by the Natural Sciences and Engineering Research Council of Canada under the Discovery Grants program.
**Supplemental Material:** The e-companion is available at https://doi.org/10.1287/mnsc.2020.3666.

**Keywords:** group decisions • voting committees • utility preference theory • artificial intelligence

## 1. Introduction

A central societal question is how to consolidate diverse preferences and opinions into reasonable, collective decisions. Classical voting theory takes an axiomatic approach that identifies desirable properties that the aggregation method should satisfy, and studies the (non)existence and structure of such rules. A celebrated example of this is Arrow's impossibility result (Arrow 1951). By contrast, the field of computational social choice (Brandt et al. 2016) typically attempts to identify an appealing objective function and design aggregation rules to optimize this objective.

One of the best-studied problems in computational social choice deals with aggregating individual preferences over alternatives—expressed as rankings—into a collective choice of a subset of alternatives (Procaccia et al. 2012, Skowron et al. 2015, Caragiannis et al. 2016). Nascent social choice applications, though, have given rise to the harder, richer problem of budgeted social choice (Lu and Boutilier 2011), where alternatives have associated costs, and the selected subset is subject to a budget constraint.

Our interest in budgeted social choice stems from the striking real-world impact of the participatory budgeting paradigm (Cabannes 2004), which allows local governments to allocate public funds by eliciting and aggregating the preferences of residents over potential projects. Indeed, in just a few years, the Participatory Budgeting Project[1] has helped allocate more than $300 million of public money for more than 1,600 local projects, primarily in the United States and Canada (including New York City, Chicago, Boston, and San Francisco).

Participatory budgeting has also attracted attention globally. A 2007 study by the World Bank (Shah 2007) reports instances of participatory budgeting in locations as diverse as Guatemala, Peru, Romania, and South Africa. In Europe, the push for participatory budgeting is arguably led by Madrid and Paris: both cities have spent more than 100 million euros on participatory budgets in 2017 (Gutiérrez 2017, Legendre et al. 2017). Notably, a participatory budgeting application is also included in the Decide Madrid open-source tool for civic engagement, providing a framework to simplify the hosting and management of participatory budgeting elections around the world.

In the first formal analysis of this paradigm, Goel et al. (2019)—who have facilitated several participatory budgeting elections as part of the Stanford Crowdsourced Democracy Team[2]—propose and evaluate two participatory budgeting approaches. In the first approach, the input format—the way in which each voter's preferences are elicited—is knapsack votes: Each voter reports his or her individual solution to the knapsack problem, that is, the set of projects that maximizes the voter's overall value (assuming an additive valuation function), subject to the budget constraint. The second component of the approach is the aggregation rule; in this case, each voter is seen as approving all the projects in the individual's knapsack, and then projects are ordered by the number of approval votes and greedily selected for execution, until the budget runs out. The second approach uses value-for-money comparisons as input format. It asks voters to compare pairs of projects by the ratio between value and cost. These comparisons are aggregated using variants of classic voting rules, including the Borda count rule and the Kemeny rule.

In a sense, Goel et al. (2019) take a bottom-up approach: They define novel, intuitive input formats that encourage voters to take cost—not just value—into account, and justify them after the fact. By contrast, we wish to take a top-down approach, by specifying an overarching optimization goal, and using it to compare different methods for participatory budgeting.

### 1.1. Our Approach and Results

Following Goel et al. (2019), we assume that voters have additive utility functions and vote over a set of alternatives, each with a known cost. Our goal is to choose a subset of alternatives that maximizes (utilitarian) social welfare subject to a budget constraint.

This reduces to a knapsack problem when we have access to the utility functions; the problem is challenging precisely because we do not. Rather, we have access to votes, in a certain input format, which are consistent with the utility functions. This goal—maximizing social welfare based on votes that serve as proxies for latent utility functions—has been studied for more than a decade (Procaccia and Rosenschein 2006, Caragiannis and Procaccia 2011, Anshelevich et al. 2015, Boutilier et al. 2015, Anshelevich and Postl 2016, Anshelevich and Sekar 2016), and has recently been termed implicit utilitarian voting (Caragiannis et al. 2016).

Absent complete information about the utility functions, clearly social welfare cannot be perfectly maximized. Procaccia and Rosenschein (2006) introduced the notion of distortion to quantify how far a given aggregation rule is from achieving this goal. Roughly speaking, given a vote profile (a set of $n$ votes) and an outcome, the distortion is the worst-case ratio between the social welfare of the optimal outcome, and the social welfare of the given outcome, where the worst case is taken with respect to all utility profiles that are consistent with the given votes.

Previous work on implicit utilitarian voting assumes that each voter expresses his preferences by ranking the alternatives in order of decreasing utility. By contrast, the main insight underlying our work is that *the implicit utilitarian voting framework allows us to decouple the input format and aggregation rule, thereby enabling an analytical comparison of different input formats in terms of their potential for providing good solutions to the participatory budgeting problem.*

This decoupling is achieved by associating each input format with the distortion of the optimal (randomized) aggregation rule, that is, the rule that minimizes distortion on every vote profile. Intuitively, the distortion associated with an input format measures how useful the information contained in the votes is for social welfare maximization (a lower distortion is better).

In Section 3, we apply this approach to compare four input formats. The first is knapsack votes, which (disappointingly) has distortion linear in the number of alternatives, the same distortion that one can achieve in the complete absence of information. Next, we analyze two closely related input formats: rankings by value and rankings by value for money, which ask voters to rank the alternatives by their value and by the ratio of their value and cost, respectively. We find that for both of these input formats, the distortion grows no faster than the square root of the number of alternatives, which matches a lower bound up to logarithmic factors. Finally, we examine a novel input format, which we call *threshold approval votes*: each voter is asked to approve the alternatives the voter values above a threshold that we choose. We find tight bounds showing that the distortion of threshold approval votes is essentially logarithmic in the number of items. To summarize, our theoretical results show striking separations between different input formats, with threshold approval votes coming out well on top.

These results may also be interpreted as approximation ratios to the optimal solution of the classical knapsack problem, where we are given only partial information about voter utilities (a vote profile, in some format) and an adversary selects both the vote profile and a utility profile consistent with the votes, which are used to evaluate our performance.

Although our theoretical results in Section 3 bound the distortion, that is, the worst-case ratio of the optimal social welfare to the social welfare achieved over all instances, it may be possible to provide much stronger performance guarantees on any specific instance. In Section 4, we design algorithms to compute the distortion-minimizing subset of alternatives (when considering deterministic aggregation rules), and distribution over subsets of alternatives (when considering randomized aggregation rules) for a specific instance.

In Section 5, we use these algorithms to evaluate the average-case ratio of the optimal social welfare to the social welfare achieved by various aggregation rules. Specifically, we compare the distortion-minimizing aggregation rule for each of the four formats we study to two approaches currently used in practice. (Note that a distortion-minimizing rule finds the subset of alternatives that is best in the worst case over all utility profiles that could have induced the observed vote profile. There is no guarantee that this subset maximizes social welfare for the actual underlying utility profile that induced the vote profile. Nonetheless, we find distortion-minimizing rules perform extremely well.) We use data from two real-world participatory budgeting elections held in Boston in 2015 and 2016. The experiments indicate that aggregation rules that minimize distortion on every input profile significantly outperform the currently deployed approaches, and among the input formats we study, threshold approval votes remain superior, even in practice. We also observe that the running times of these distortion-minimizing rules scale gracefully to practical sizes.

Our results emphasize that the choice of input format should be an important consideration when setting up a participatory budgeting framework. We observe that greedy aggregation rules, which are the current norm, can lead to inefficient outcomes. We also find compelling evidence against using knapsack votes in their poor theoretical guarantees and weaker empirical performance. Beyond these general observations, we expect the idiosyncrasies of every individual city to guide its implementation of participatory budgeting. For example, one set of city officials revealed that they used knapsack votes specifically so that voters experience the constraints that come with having a limited budget.

### 1.2. Related Work

Let us first describe the theoretical results of Goel et al. (2019) in slightly greater detail. Most relevant to our work is a theorem that asserts that knapsack voting (i.e., knapsack votes as the input format, coupled with greedy approval-based aggregation) actually maximizes social welfare. However, the result strongly relies on their overlap utility model, where the utility of a voter for a subset of alternatives is (roughly speaking) the size of the intersection between this subset and the voter's own knapsack vote. In a sense, the viewpoint underlying this model is the opposite of ours, as a voter's utility is derived from the individual's vote, instead of the other way around. One criticism of this model is that even if certain alternatives do not fit into a voter's individual knapsack solution due to the budget constraint, the voter could (and usually will) have some utility for them. Goel et al. (2019) also provide strategyproofness results for knapsack voting, which similarly rely on the overlap utility model. Finally, they interpret their methods as maximum likelihood estimators (Young 1988, Conitzer and Sandholm 2005) under certain noise models.

As our work applies the implicit utilitarian voting approach (Boutilier et al. 2015, Caragiannis et al. 2016) to a problem in the budgeted social choice framework (Lu and Boutilier 2011), it is naturally related to both lines of work. Lu and Boutilier (2011) introduce the budgeted social choice framework, in which the goal is to collectively select a set of alternatives subject to a budget constraint. Their framework generalizes the participatory budgeting problem studied herein as it allows the cost of an alternative to also depend on the number of voters who derive utility from the alternative. However, their results are incomparable to ours because they assume that every voter's utility for an alternative is determined solely by the rank of the alternative in the voter's preference order—specifically, that the utilities of all voters follow a common underlying positional scoring rule—which is a common assumption in the literature on resource allocation (Bouveret and Lang 2011, Baumeister et al. 2017). This makes the elicitation problem trivial because eliciting ordinal preferences (i.e., rankings by value) is assumed to accurately reveal the underlying cardinal utilities. By contrast, we do not impose such a restriction on the utilities, and compare the rankings-by-value input format with three other input formats.

Previous work on implicit utilitarian voting focuses exclusively on the rankings-by-value input format. Boutilier et al. (2015) study the problem of selecting a single winning alternative, and provide an upper and lower bound on the distortion achieved by the optimal aggregation rule. Their setting is a special case of the participatory budgeting problem where the cost of each alternative equals the entire budget. Consequently, their lower bound applies to our more general setting, and our upper bound for the rankings-by-value input format generalizes theirs (up to a logarithmic factor). Caragiannis et al. (2016) extend the results of Boutilier et al. (2015) to the case where a subset of alternatives of a given size $k$ is to be selected (only for the rankings-by-value input format); this is again a special case of the participatory budgeting problem where the cost of each alternative is $B/k$. However, our results are incomparable to theirs because we assume additive utility functions—following previous work on participatory budgeting (Goel et al. 2019)—whereas Caragiannis et al. (2016) assume that a voter's utility for a subset of alternatives is the voter's maximum utility for any alternative in the subset.

The core idea behind implicit utilitarian voting—approximating utilitarian social welfare given ordinal information—has also been studied in mechanism design. Filos-Ratsikas et al. (2014) and Anshelevich and Sekar (2016) present algorithms for finding matchings in weighted graphs given ordinal comparisons among the edges by their weight, and Abramowitz and Anshelevich (2018) generalize this to a class of utility maximization problems on graphs, including the maximum weight spanning tree, maximum weight $b$-matching, and maximum traveling salesperson problems. Krysta et al. (2014) study distortion in the house allocation problem and Chakrabarty and Swamy (2014) consider it in a general mechanism design setting, but with the restriction borrowed from Lu and Boutilier (2011) that the utilities of all agents are determined by a common positional scoring rule.

A line of research on resource allocation focuses on maximizing other forms of welfare such as the egalitarian welfare or the Nash welfare (see, e.g., Moulin 2003). Maximizing the Nash welfare has the benefit that it is invariant to scaling an agent's utility function, and thus does not require normalizing the utilities. In addition, it is known to satisfy nontrivial fairness guarantees in domains that are similar to or generalize participatory budgeting (Conitzer et al. 2017, Fain et al. 2018). It remains to be seen whether maximizing the Nash welfare subject to votes that only partially reveal the underlying utilities can preserve such guarantees.

## 2. The Model

Let $[k] = \{1, \ldots, k\}$ denote the set of $k$ smallest positive integers. Let $N = [n]$ be the set of voters, and $A$ be the set of $m$ alternatives. The cost of alternative $a$ is denoted $c_a$, and the budget $B$ is normalized to 1. For $S \subseteq A$, let $c(S) = \sum_{a \in S} c_a$. Define $\mathscr{F}_c = \{S \subseteq A : c(S) \leq 1 \wedge c(T) > 1, \forall S \subsetneq T \subseteq A\}$ as the inclusion-maximal budget-feasible subsets of $A$.

We assume that each voter has a utility function $v_i : A \to \mathbb{R}_+ \cup \{0\}$, where $v_i(a)$ is the utility that voter $i$ has for alternative $a$, and that these utilities are additive, that is, the utility of voter $i$ for a set $S \subseteq A$ is defined as $v_i(S) = \sum_{a \in S} v_i(a)$. Finally, to ensure fairness among voters, we make the standard assumption (Caragiannis and Procaccia 2011, Boutilier et al. 2015) that $v_i(A) = 1$ for all voters $i \in N$. We note that alternative normalizations exist. For example, we may assume $0 \leq v_i(a) \leq 1$ for all $a \in A$ as Filos-Ratsikas and Miltersen (2014) do; however, we do not see a compelling reason to prefer this assumption. We call the vector $\vec{v} = \{v_1, \ldots, v_n\}$ of voter utility functions the utility profile.

Given the utility profile, the (utilitarian) social welfare of an alternative $a \in A$ is defined as $\text{sw}(a, \vec{v}) = \sum_{i \in N} v_i(a)$; for a set $S \subseteq A$, let $\text{sw}(S, \vec{v}) = \sum_{a \in S} \text{sw}(a, \vec{v})$. If voter utilities were known, our objective would be to return $S \in \arg\max_{X \in \mathscr{F}_c} \text{sw}(X)$, that is, the budget-feasible set of alternatives that maximizes social welfare. Observe that our unit-sum normalization implies that we are effectively maximizing normalized social welfare (Aziz 2019).

Unfortunately, the utility function of a voter $i$ is only accessible through the individual's vote $\rho_i$, which is induced by $v_i$. The vector $\vec{\rho} = \{\rho_1, \ldots, \rho_n\}$ is called the input profile. Let $\vec{v} \rhd \vec{\rho}$ denote that utility profile $\vec{v}$ is consistent with input profile $\vec{\rho}$. We study four specific formats for input votes. In the following, we describe each input format along with a sample question that may be asked to the voters to elicit votes in that format. The voters can be induced to think of their utilities for the different alternatives (i.e., projects) in a normalized fashion by asking them to (mentally) divide a constant sum of points—say, 1,000 points—among the alternatives based on how much they like each alternative.

- The *knapsack vote* $\kappa_i \subseteq A$ of voter $i \in N$ represents a feasible subset of alternatives with the highest value for the voter. We have $v_i \rhd \kappa_i$ if and only if $c(\kappa_i) \leq 1$ and $v_i(\kappa_i) \geq v_i(S)$ for all $S \in \mathscr{F}_c$. If the total budget is \$100,000, the voters may be asked: "Select the best set of projects according to you subject to a total budget of \$100,000."
- The *rankings-by-value* and the *rankings-by-value-for-money* input formats ask voter $i \in N$ to rank the alternatives by decreasing value for the voter, and by decreasing ratio of value for the voter to cost, respectively. Formally, let $\mathscr{L} = \mathscr{L}(A)$ denote the set of rankings over the alternatives. For a ranking $\sigma \in \mathscr{L}$, let $\sigma(a)$ denote the position of alternative $a$ in $\sigma$, and $a \succ_\sigma b$ denote $\sigma(a) < \sigma(b)$, that is, that $a$ is preferred to $b$ under $\sigma$. Then, we say that utility function $v_i$ is consistent with the ranking by value (respectively, value for money) of voter $i \in N$, denoted $\sigma_i$, if and only if $v_i(a) \geq v_i(b)$ (respectively, $v_i(a)/c_a \geq v_i(b)/c_b$) for all $a \succ_{\sigma_i} b$. To elicit such votes, the voters may be asked: "If you had to divide 1,000 points among the projects based on how much you like them, rank the projects in the decreasing order of the number of points they would receive (divided by the cost)."
- For a threshold $t$, the *threshold approval vote* $\tau_i$ of voter $i \in N$ consists of the set of alternatives whose value for the voter is at least $t$, that is, $v_i \rhd \tau_i$ if and only if $\tau_i = \{a \in A : v_i(a) \geq t\}$. To elicit threshold approval votes with a threshold $t = 1/10$, the voters may be asked: "If you had to divide 1,000 points among the projects based on how much you like them, select all the projects that would receive at least 100 points." Due to our normalization, we are effectively asking voters whether their utility for an alternative is at least some fraction of their total utility, making this input format scale-invariant like the others.

In our setting, a (randomized) aggregation rule $f$ for an input format maps each input profile $\vec{\rho}$ in that format to a distribution over $\mathcal{F}_c$. The rule is deterministic if it returns a particular set in $\mathcal{F}_c$ with probability 1.

In the implicit utilitarianism framework, the ultimate goal is to maximize the (utilitarian) social welfare. Procaccia and Rosenschein (2006) use the notion of distortion to quantify how far an aggregation rule $f$ is from achieving this goal. The distortion of $f$ on a vote profile $\vec{\rho}$ is given by

$$\mathrm{dist}(f,\vec{\rho}) = \sup_{\vec{v}:\vec{v}\triangleright\vec{\rho}} \frac{\max_{T\in\mathcal{F}_c}\mathrm{sw}(T,\vec{v})}{\mathbb{E}[\mathrm{sw}(f(\vec{\rho}),\vec{v})]}.$$

The (overall) distortion of a rule $f$ is given by $\mathrm{dist}(f) = \max_{\vec{\rho}}\ \mathrm{dist}(f,\vec{\rho})$. The optimal (randomized) aggregation rule $f^*$, which we term the distortion-minimizing aggregation rule, selects the distribution minimizing distortion on each input profile individually, that is,

$$f^*(\vec{\rho}) = \arg\min_{p\in\Delta(\mathcal{F}_c)} \sup_{\vec{v}:\vec{v}\triangleright\vec{\rho}} \frac{\max_{T\in\mathcal{F}_c}\mathrm{sw}(T,\vec{v})}{\mathbb{E}[\mathrm{sw}(p,\vec{v})]},$$

where $\Delta(\mathcal{F}_c)$ is the set of distributions over $\mathcal{F}_c$. Needless to say, $f^*$ achieves the best possible overall distortion. Similarly, the deterministic distortion-minimizing aggregation rule $f^*_{\mathrm{det}}$ is given by

$$f^*_{\mathrm{det}}(\vec{\rho}) = \arg\min_{S\in\mathcal{F}_c} \sup_{\vec{v}:\vec{v}\triangleright\vec{\rho}} \frac{\max_{T\in\mathcal{F}_c}\mathrm{sw}(T,\vec{v})}{\mathrm{sw}(S,\vec{v})}.$$

Finally, we say that the distortion associated with an input format (i.e., elicitation method) is the overall distortion of the (randomized) distortion-minimizing aggregation rule for that format; this, in a sense, quantifies the effectiveness of the input format in achieving social welfare maximization. In a setting where deterministic rules must be used, we say that the distortion associated with deterministic aggregation of votes in an input format is the overall distortion of the deterministic distortion-minimizing aggregation rule for that format. Observe that we always mention deterministic aggregation explicitly, and the distortion associated with an input format allows randomized aggregation by default.

## 3. Theoretical Results

In Section 3.1, we present theoretical results for the distortion associated with different input formats when no constraints are imposed on the aggregation rule, that is, when randomized aggregation rules are allowed. Subsequently, in Section 3.2, we study the distortion associated with deterministic aggregation under these input formats.

### 3.1. Randomized Aggregation Rules

We begin by making a simple observation that holds for (randomized) aggregation of votes in any input format.

**Observation 1.** The distortion associated with any input format is at most $m$.

**Proof of Observation 1.** Consider the rule that selects a single alternative uniformly at random; this is clearly budget-feasible. Due to the normalization of utility functions, the expected welfare achieved by this rule is $(1/m)\cdot\sum_{i\in N}\sum_{a\in A} v_i(a) = n/m$. On the other hand, the maximum welfare that any subset of alternatives can achieve is at most $n$. Hence, the distortion of this rule, which does not require any input, is at most $m$. □

**3.1.1. Knapsack Votes.** We now present our analysis for knapsack votes—an input format advocated by Goel et al. (2019).

**Theorem 2.** *For $n \geq m$, the distortion associated with knapsack votes is $\Omega(m)$.*

**Proof of Theorem 2.** Consider the case where every alternative has cost 1 (i.e., equal to the budget). Consider the input profile $\vec{\kappa}$, in which voters are partitioned into $m$ subsets $\{N_a\}_{a\in A}$ of roughly equal size; specifically, let $n_a = |N_a|$ and enforce $\lfloor n/m\rfloor \leq n_a \leq \lceil n/m\rceil$ for all $a \in A$. For every $a \in A$ and $i \in N_a$, let $\kappa_i = \{a\}$.

Consider a randomized aggregation rule $f$. There must exist an alternative $a^* \in A$ such that $\Pr[f(\vec{\kappa}) = \{a^*\}] \leq 1/m$. Now, construct a utility profile $\vec{v}$ such that i) for all $i \in N_{a^*}$, we have $v_i(a^*) = 1$, and $v_i(a) = 0$ for $a \in A\setminus\{a^*\}$; and ii) for all $a \in A\setminus\{a^*\}$ and $i \in N_a$, we have $v_i(a) = v_i(a^*) = 1/2$, and $v_i(b) = 0$ for $b \in A\setminus\{a, a^*\}$.

Note that $\vec{v}$ is consistent with the input profile $\vec{\kappa}$, that is, $\vec{v}\triangleright\vec{\kappa}$. Moreover, it holds that $\mathrm{sw}(a^*,\vec{v}) \geq n/2$, whereas $\mathrm{sw}(a,\vec{v}) \leq n_a \leq n/m+1$ for $a \in A\setminus\{a^*\}$. It follows that

$$\mathrm{dist}(f) \geq \mathrm{dist}(f,\vec{\kappa}) \geq \frac{n/2}{\frac{1}{m}\cdot n + \frac{m-1}{m}\cdot\left(\frac{n}{m}+1\right)} \geq \frac{m}{6},$$

as desired. □

In light of Observation 1, this result indicates that the distortion associated with knapsack votes is asymptotically indistinguishable from the distortion one can achieve with absolutely no information about voter preferences, suggesting that knapsack votes may not be an appropriate input format if the goal is to maximize social welfare. Our aim now is to find input formats that achieve better results when viewed through the implicit utilitarianism lens.

**3.1.2. Rankings by Value and by Value for Money.** Goel et al. (2019) also advocate the use of comparisons between alternatives based on value for money, which, like knapsack votes, encourage voters to consider the trade-off between value and cost. We study rankings by value for money as an input format; observe that such rankings convey more information than specific pairwise comparisons.

In addition, we study rankings by value, which are prevalent in the existing literature on implicit utilitarian voting (Procaccia and Rosenschein 2006, Caragiannis and Procaccia 2011, Anshelevich et al. 2015, Boutilier et al. 2015, Anshelevich and Postl 2016, Anshelevich and Sekar 2016). Rankings by value convey more information than $k$-approval votes, in which each voter submits the set of top $k$ alternatives by their value—this is the input format of choice for most real-world participatory budgeting elections (Goel et al. 2019).

Boutilier et al. (2015) prove a lower bound of $\Omega(\sqrt{m})$ on distortion in the special case of our setting where all alternatives have cost 1, the input format is rankings by value, and $n \geq \sqrt{m}$. This result carries over to our more general setting, not only with rankings by value, but also with rankings by value for money, as both input formats coincide in case of equal costs. Our goal is to establish an almost matching upper bound.

We start from a mechanism of Boutilier et al. (2015) that has distortion $\mathcal{O}(\sqrt{m \log m})$ in their setting. It carefully balances between high-value and low-value alternatives (where value is approximately inferred from the positions of the alternatives in the input rankings). In our more general participatory budgeting problem, it is crucial to also take into account the costs, and find the perfect balance between selecting many low-cost alternatives and fewer high-cost ones. We modify the mechanism of Boutilier et al. precisely to achieve this goal. Specifically, we partition the alternatives into $\mathcal{O}(\log m)$ buckets based on their costs, and differentiate between alternatives within a bucket based on their (inferred) value. Our mechanism for rankings by value for money requires more careful treatment as values are obfuscated in value-for-money comparisons.

At first glance, our setting seems much more difficult, distortion-wise, than the simple setting of Boutilier et al. (2015). But ultimately, we obtain only a slightly weaker upper bound on the distortion associated with both rankings by value and by value for money. In other words, to our surprise, incorporating costs and a budget constraint comes at almost no cost (no pun intended) to social welfare maximization.

**Theorem 3.** *The distortion associated with rankings by value and rankings by value for money is* $\mathcal{O}(\sqrt{m}\ \log m)$.

**Proof of Theorem 3.** We first present the proof for rankings by value for money as it is trickier, and later describe how an almost identical proof works for rankings by value.

Let us begin by introducing additional notation. For a ranking $\sigma$ and an alternative $a \in A$, let $\sigma(a)$ denote the position of $a$ in $\sigma$. For a preference profile $\vec{\sigma}$ with $n$ votes, let the harmonic score of $a$ in $\vec{\sigma}$ be defined as $\mathrm{sc}(a, \vec{\sigma}) = \sum_{j=1}^{n} 1/\sigma_j(a)$. Finally, given a set of alternatives $S \subseteq A$, let $\sigma|_S$ (respectively, $\vec{\sigma}|_S$) denote the ranking (respectively, preference profile) obtained by restricting $\sigma$ (respectively, $\vec{\sigma}$) to the alternatives in $S$.

For ease of exposition assume $m$ is a power of 2. Let $\vec{\sigma}$ denote the input profile consisting of voter preferences in the form of rankings by value for money. Let $\vec{v}$ denote the underlying utility profile consistent with $\vec{\sigma}$. Let $S^* = \arg\max_{S \in \mathcal{F}_c} \mathrm{sw}(S, \vec{v})$ be the budget-feasible set of alternatives maximizing the social welfare.

Define $\ell_0 = 0$ and $u_0 = 1/m$. For $i \in [\log m]$, define $\ell_i = 2^{i-1}/m$ and $u_i = 2^i/m$. Let us partition the alternatives into $\log m + 1$ buckets based on their costs: $S_0 = \{a \in A : c_a \leq u_0\}$ and $S_i = \{a \in A : \ell_i < c_a \leq u_i\}$ for $i \in [\log m]$. Note that for $i \in \{0\} \cup [\log m]$, selecting at most $1/u_i$ alternatives from $S_i$ is guaranteed to be budget-feasible.

Next, let us further partition the buckets into two parts: for $i \in \{0\} \cup [\log m]$, let $S_i^+$ consist of the $\sqrt{m} \cdot (1/u_i)$ alternatives from $S_i$ with the largest harmonic scores in the reduced profile $\vec{\sigma}|_{S_i}$, and $S_i^- = S_i \setminus S_i^+$. If $|S_i| \leq \sqrt{m} \cdot (1/u_i)$, we let $S_i^+ = S_i$ and $S_i^- = \emptyset$. Note that $S_0^+ = S_0$. Let $S^+ = \cup_{i=0}^{\log m} S_i^+$ and $S^- = A \setminus S^+$.

We are now ready to define our randomized aggregation rule, which randomizes over two separate mechanisms:

- Mechanism *A:* Select a bucket $S_i$ uniformly at random, and select a $(1/u_i)$-size subset of $S_i^+$ uniformly at random.
- Mechanism *B:* Select a single alternative uniformly at random.

Our aggregation rule executes each mechanism with an equal probability 1/2. We now show that this rule achieves distortion that is $\mathcal{O}(\sqrt{m} \log m)$.

First, note that mechanism $A$ selects each bucket $S_i$ with probability $1/(\log m + 1)$, and when $S_i$ is selected, it selects each alternative in $S_i^+$ with probability at least $1/\sqrt{m}$. (This is because the mechanism selects $1/u_i$ alternatives at random from $S_i^+$, which has at most $\sqrt{m} \cdot (1/u_i)$ alternatives.) Hence, the mechanism selects each alternative in $S^+$ (and therefore, each alternative in $S^* \cap S^+$) with probability at least $1/(\sqrt{m}(\log m + 1))$. In other words, the expected social welfare achieved under mechanism $A$ is $\mathcal{O}(\sqrt{m} \log m)$ approximation of $\mathrm{sw}(S^* \cap S^+, \vec{v})$.

Finally, to complete the proof, we show that the expected welfare achieved under mechanism $B$ is an $\mathcal{O}(\sqrt{m}\log m)$ approximation of $\mathrm{sw}(S^* \cap S^-, \vec{v})$. Let us first bound $\mathrm{sw}(S^* \cap S^-, \vec{v})$. Recall that $S_0^- = \emptyset$. Hence,

$$\mathrm{sw}\left(S^* \cap S^-, \vec{v}\right) = \sum_{i=1}^{\log m} \mathrm{sw}\left(S^* \cap S_i^-, \vec{v}\right).$$

Fix $i \in [\log m]$ and $a \in S_i^-$. One can easily check that

$$\sum_{b \in S_i} \mathrm{sc}\left(b, \vec{\sigma}|_{S_i}\right) = n \cdot H_{|S_i|} \le n \cdot H_m,$$

where $H_k$ is the $k$th harmonic number. Because $S_i^+$ consists of the $\sqrt{m}/u_i$ alternatives in $S_i$ with the largest harmonic scores, we have

$$\mathrm{sc}\left(a, \vec{\sigma}|_{S_i}\right) \le \frac{n \cdot H_m}{\sqrt{m} \cdot (1/u_i)} = \frac{n \cdot (1 + \log m)}{\sqrt{m} \cdot m/2^i}. \tag{1}$$

Next, we connect this bound on the harmonic score of $a$ to a bound on its social welfare. For simplicity, let us denote $\vec{\gamma} \triangleq \vec{\sigma}|_{S_i}$. Due to our definition of the partitions, we have

$$c_a \le 2 \cdot c_b, \forall b \in S_i. \tag{2}$$

Further, fix a voter $j \in [n]$. For each alternative $b$ such that $b \succ_{\gamma_j} a$, we also have $v_j(b)/c_b \ge v_j(a)/c_a$. Substituting Equation (2), we get

$$v_j(a) \le 2 v_j(b), \forall j \in [n], b \in S_i \text{ s.t. } b \succ_{\gamma_j} a. \tag{3}$$

Taking a sum over all $b \in S_i$ with $b \succ_{\gamma_j} a$, and using the fact that the values of each voter $j$ sum to 1, we get $v_j(a) \le 2/\gamma_j(a)$ for $j \in [n]$, and taking a further sum over $j \in [n]$, we get

$$\mathrm{sw}\left(a, \vec{v}\right) \le 2 \cdot \mathrm{sc}\left(a, \vec{\sigma}|_{S_i}\right). \tag{4}$$

Combining this with Equation (1), we get

$$\mathrm{sw}\left(a, \vec{v}\right) \le \frac{2 \cdot n \cdot (1 + \log m)}{\sqrt{m} \cdot m/2^i}, \forall a \in S_i^-.$$

Note that $S^*$ can contain at most $2/u_i = m/2^{i-1}$ alternatives from $S_i$ while respecting the budget constraint. Hence,

$$\begin{aligned}\mathrm{sw}\left(S^* \cap S^-, \vec{v}\right) &= \sum_{i=1}^{\log m} \mathrm{sw}\left(S^* \cap S_i^-, \vec{v}\right) \\ &\le \frac{(m/2^{i-1}) \cdot 2 \cdot n \cdot (1 + \log m)}{\sqrt{m} \cdot m/2^i} \\ &= 4 \cdot n \cdot (1 + \log m)/\sqrt{m}.\end{aligned} \tag{5}$$

Because the utilities sum to 1 for each voter, the expected social welfare achieved under mechanism $B$ is $(1/m) \cdot \sum_{i \in N} \sum_{a \in A} v_i(a) = n/m$, which is an $\mathcal{O}(\sqrt{m}\log m)$ approximation of $\mathrm{sw}(S^* \cap S^-, \vec{v})$ due to Equation (5).

This completes the proof of $\mathcal{O}(\sqrt{m}\log m)$ distortion associated with rankings by value for money. The proof for rankings by value is almost identical. In fact, one can make two simplifications.

First, the factor of 2 from Equation (3), and therefore from Equation (4) disappears because the rankings already dictate comparison by value. This leads to an improvement in Equation (5) by a factor of 2.

Second, Equation (3) not only holds for $b \in S_i$ such that $b \succ_{\gamma_j} a$, but also holds more generally for $b \in A$ such that $b \succ_{\sigma_j} a$. Hence, there is no longer a need to compute the harmonic scores on the restricted profile $\vec{\sigma}|_{S_i}$; one can simply work with the original input profile $\vec{\sigma}$. □

### 3.1.3. Threshold Approval Votes.

Approval voting—where voters can choose to approve any subset of alternatives, and a most widely approved alternative wins—is well studied in social choice theory (Brams and Fishburn 2007). In our utilitarian setting, we reinterpret this input format as threshold approval votes, where the principal sets a threshold $t$, and each voter $i \in N$ approves every alternative $a$ for which $v_i(a) \ge t$.

We first investigate deterministic threshold approval votes, in which the threshold is selected deterministically, but find that it does not help us (significantly) improve over the distortion we can already obtain using rankings by value or by value for money. Specifically, for a fixed threshold, we are always able to construct cases in which alternatives have significantly different welfares, but either no alternative is approved or an extremely large set of alternatives is approved, providing the rule little information to distinguish between the alternatives, and yielding high distortion.

**Theorem 4.** *The distortion associated with deterministic threshold approval votes is* $\Omega(\sqrt{m})$.

**Proof of Theorem 4.** Imagine the case where $c_a = 1$ for all alternatives $a \in A$. Recall that the budget is 1. Let $f$ denote a randomized aggregation rule. (Although we study deterministic and randomized threshold selection, we still allow randomized aggregation rules. Section 3.2 studies the case where the aggregation rule has to be deterministic.) It must return a single alternative, possibly chosen in a randomized fashion. We construct our adversarial input profile based on whether $t \le 1/\sqrt{m}$. Let $A = \{a_1, \ldots, a_m\}$.

Suppose $t \le 1/\sqrt{m}$. Fix a set of alternatives $S \subseteq A$ such that $|S| = \sqrt{m}/2 + 1$ (assume for ease of exposition $\sqrt{m}$ is an even integer). Construct the input profile $\vec{\tau}$ such that $\tau_i = S$ for all $i \in N$. Now, there must exist $a^* \in S$ such that $\Pr[f(\vec{\tau}) = \{a^*\}] \le 1/(\sqrt{m}/2 + 1)$. Construct the underlying utility profile $\vec{v}$ such that for each voter

$i \in N$, $v_i(a^*) = 1/2$, $v_i(a) = 1/\sqrt{m}$ for $a \in S \setminus \{a^*\}$, and $v_i(a) = 0$ for $a \in A \setminus S$. Note that this is consistent with the input profile given that $t \le 1/\sqrt{m}$. Further, $\mathrm{sw}(a^*, \vec{v}) = n/2$ whereas $\mathrm{sw}(a, \vec{v}) \le n/\sqrt{m}$ for all $a \in A \setminus \{a^*\}$. Hence,

$$\mathbb{E}\big[\mathrm{sw}\big(f(\vec{\tau}), \vec{v}\big)\big] \le \frac{1}{\sqrt{m}/2+1} \cdot \frac{n}{2} + \frac{\sqrt{m}/2}{\sqrt{m}/2+1} \cdot \frac{n}{\sqrt{m}} = \mathcal{O}\left(\frac{n}{\sqrt{m}}\right).$$

Because the optimal social welfare is $\Theta(n)$, we have that $\mathrm{dist}(f) = \Omega(\sqrt{m})$, as required.

Now suppose that $t > 1/\sqrt{m}$. Construct an input profile $\vec{\tau}$ in which $\tau_i = \emptyset$ for every voter $i \in N$. In this case, there exists an alternative $a^* \in A$ such that $\Pr[f(\vec{\tau}) = a^*] \le 1/m$. Let us construct the underlying utility profile $\vec{v}$ as follows. For every voter $i \in N$, let $v_i(a^*) = 1/\sqrt{m}$, and $v_i(a) = (1 - 1/\sqrt{m})/m$ for all $a \in A \setminus \{a^*\}$. Note that this is consistent with the input profile given that $t > 1/\sqrt{m}$. Clearly, the optimal social welfare is achieved by $\mathrm{sw}(a^*, \vec{v}) = n/\sqrt{m}$. In contrast, we have

$$\mathbb{E}\big[\mathrm{sw}\big(f(\vec{\tau}), \vec{v}\big)\big] \le \frac{1}{m} \cdot \frac{n}{\sqrt{m}} + \left(1 - \frac{1}{\sqrt{m}}\right) \cdot \frac{1 - 1/\sqrt{m}}{m} = \mathcal{O}\left(\frac{n}{m}\right).$$

Hence, we again have $\mathrm{dist}(f) = \Omega(\sqrt{m})$, as desired. □

For specific ranges of the threshold, it is possible to derive stronger lower bounds. However, the $\Omega(\sqrt{m})$ lower bound of Theorem 4 is sufficient to establish a clear asymptotic separation between the power of deterministic and randomized threshold approval votes.

Under randomized threshold approval votes, we can select the threshold in a randomized fashion. Technically, this is a distribution over input formats, one for each value of the threshold. Before we define the (overall) distortion of a rule that randomizes over input formats, let us recall the definition of the overall distortion of a rule for a fixed input format:

$$\mathrm{dist}\big(f\big) = \max_{\vec{\rho}} \sup_{\vec{v}: \vec{v} \triangleright \vec{\rho}} \frac{\max_{T \in \mathcal{F}_c} \mathrm{sw}\big(T, \vec{v}\big)}{\mathbb{E}\big[\mathrm{sw}\big(f\big(\vec{\rho}\big), \vec{v}\big)\big]} = \sup_{\vec{v}} \frac{\max_{T \in \mathcal{F}_c} \mathrm{sw}\big(T, \vec{v}\big)}{\mathbb{E}\big[\mathrm{sw}\big(f\big(\vec{\rho}\big), \vec{v}\big)\big]}.$$

Here, $\vec{\rho}(\vec{v})$ denotes the input profile induced by utility profile $\vec{v}$. In the case of randomized threshold approval votes, rule $f$ specifies a distribution $D$ over the threshold $t$, as well as the aggregation of input profile $\vec{\rho}(\vec{v}, t)$ induced by utility profile $\vec{v}$ and a given choice of threshold $t$. We define the (overall) distortion of rule $f$ as

$$\mathrm{dist}\big(f\big) = \sup_{\vec{v}} \; \mathbb{E}_{t \sim D} \; \frac{\max_{T \in \mathcal{F}_c} \mathrm{sw}\big(T, \vec{v}\big)}{\mathbb{E}\big[\mathrm{sw}\big(f\big(\vec{\rho}(\vec{v}, t)\big), \vec{v}\big)\big]}.$$

Interestingly, observe that due to the expectation over threshold $t$, which affects the induced input profile $\vec{\rho}(\vec{v}, t)$, we can no longer decompose the maximum over $\vec{v}$ into a maximum over $\vec{\rho}$ followed by a maximum over $\vec{v}$ such that $\vec{v} \triangleright \vec{\rho}$, in contrast to the case of a fixed input format.

This flexibility of randomizing the threshold value allows us to dramatically reduce the distortion.

**Theorem 5.** *The distortion associated with randomized threshold approval votes is $\mathcal{O}(\log^2 m)$.*

**Proof of Theorem 5.** For ease of exposition, assume $m$ is a power of 2. Let $I_0 = [0, 1/m^2]$, and $I_j = (2^{j-1}/m^2, 2^j/m^2]$, $\ell_j = 2^{j-1}/m^2$, and $u_j = 2^j/m^2$ for $j = 1, \ldots, 2\log m$.

Let $\vec{v}$ denote a utility profile that is consistent with the input profile. For $a \in A$ and $j \in \{0, \ldots, 2\log m\}$, define $n_j^a = |\{i \in N : v_i(a) \in I_j\}|$ to be the number of voters whose utility for $a$ falls in the interval $I_j$. We now bound the social welfare of $a$ in terms of the numbers $n_j^a$. Specifically,

$$\mathrm{sw}\big(a, \vec{v}\big) = \sum_{i \in N} v_i(a) \le \sum_{j=0}^{2\log m} \sum_{i \in N} \mathbb{I}\big\{v_i(a) \in I_j\big\} \cdot u_j = \sum_{j=0}^{2\log m} n_j^a \cdot u_j,$$

where $\mathbb{I}$ indicates the indicator variable. A similar argument also yields a lower bound, and after substituting $\ell_0 = 0$, $u_0 = 1/m^2$, and $n_0^a \le n$, we get

$$\sum_{j=1}^{2\log m} n_j^a \cdot \ell_j \le \mathrm{sw}\big(a, \vec{v}\big) \le \frac{n}{m^2} + \sum_{j=1}^{2\log m} n_j^a \cdot u_j. \tag{6}$$

Next, divide the alternatives into $1 + 2\log m$ buckets based on their costs, with bucket $S_j = \{a \in A : c_a \in I_j\}$. Note that selecting at most $1/u_j$ alternatives from $S_j$ is guaranteed to satisfy the budget constraint.

Let $S^* = \arg\max_{S \in \mathcal{F}_c} \mathrm{sw}(S, \vec{v})$ be the feasible set of alternatives maximizing the social welfare. For $j, k \in \{0, \ldots, 2\log m\}$, let $n^*_{j,k} = \sum_{a \in S^* \cap S_k} n_j^a$ denote the number of voters whose utility for an alternative falls in interval $I_j$, summed over the alternatives that appear in both $S^*$ and $S_k$. Using Equation (6), we have

$$\sum_{j=1}^{2\log m} n^*_{j,k} \cdot \ell_j \le \mathrm{sw}\big(S^* \cap S_k, \vec{v}\big) \le |S^* \cap S_k| \cdot \frac{n}{m^2} + \sum_{j=1}^{2\log m} n^*_{j,k} \cdot u_j. \tag{7}$$

We now construct three different mechanisms; our final mechanism will randomize between them:

*Mechanism A:* Pick a pair $(j, k)$ uniformly at random from the set $T = \{(j, k) : j, k \in [2\log m]\}$. Then, set the threshold to $\ell_j$, and using the resulting input profile, greedily select the $1/u_k$ alternatives from $S_k$ with the largest number of approval votes (or select $S_k$ if $|S_k| \le 1/u_k$). Let $B_{j,k}$ denote the set of selected alternatives for the pair $(j,k)$. Because we have $j > 0$ and $k > 0$,

$$\mathrm{sw}(B_{j,k}, \vec{v}) \ge \sum_{a \in B_{j,k}} \left( \sum_{p=j}^{2\log m} n_p^a \right) \cdot \ell_j \ge \frac{1}{4} \cdot \left( \sum_{p=j}^{2\log m} n_{p,k}^* \right) \cdot u_j \ge \frac{1}{4} \cdot n_{j,k}^* \cdot u_j, \qquad (8)$$

where, in the first transition, we bound the welfare from below by only considering utilities that are at least $\ell_j$, and the second transition holds because $u_j = 2\ell_j$, $|S^* \cap S_k| \le 2|B_{j,k}|$, and $B_{j,k}$ consists of greedily selected alternatives with the highest number of approval votes. Thus, the expected social welfare achieved by mechanism $A$ is

$$\begin{aligned} &\frac{1}{(2\log m)^2} \sum_{j=1}^{2\log m} \sum_{k=1}^{2\log m} \mathrm{sw}(B_{j,k}, \vec{v}) \\ &\ge \frac{1}{4 \cdot (2\log m)^2} \sum_{j=1}^{2\log m} \sum_{k=1}^{2\log m} n_{j,k}^* \cdot u_j \\ &\ge \frac{1}{16\ \log^2 m} \left( \mathrm{sw}(S^* \setminus S_0, \vec{v}) - |S^* \setminus S_0| \cdot \frac{n}{m^2} \right) \\ &\ge \frac{1}{16\ \log^2 m} \left( \mathrm{sw}(S^* \setminus S_0, \vec{v}) - \frac{n}{m} \right), \end{aligned}$$

where the first transition follows from Equation (8), and the second transition follows from Equation (7).

*Mechanism B:* Select all the alternatives in $S_0$. Because each alternative in $S_0$ has cost at most $1/m^2$, this is clearly budget-feasible. The social welfare achieved by this mechanism is $\mathrm{sw}(S_0, \vec{v}) \ge \mathrm{sw}(S^* \cap S_0, \vec{v})$.

*Mechanism C:* Select a single alternative uniformly at random from $A$. This is also budget-feasible, and due to normalization of values, its expected social welfare is $n/m$.

Our final mechanism executes mechanism $A$ with probability $16\log^2 m/(2 + 16\log^2 m)$, and mechanisms $B$ and $C$ with probability $1/(2 + 16\log^2 m)$ each. It is easy to see that its expected social welfare is at least $\mathrm{sw}(S^*, \vec{v})/(2+16\log^2 m)$. Hence, its distortion is $\mathcal{O}(\log^2 m)$. □

We also show that at least logarithmic distortion is inevitable even when using randomized threshold approval votes.

**Theorem 6.** *The distortion associated with randomized threshold approval votes is* $\Omega(\log m/\log\log m)$.

**Proof of Theorem 6.** Imagine the case where $c_a = 1$ for all $a \in A$. Recall that the budget is 1. Let $f$ denote a rule that elicits randomized threshold approval votes and aggregates them to return a distribution over $A$ (as only a single project can be executed at a time). Note that $f$ is not simply the aggregation rule, but the elicitation method and the aggregation rule combined.

Divide the interval $(1/m, 1]$ into $\lceil \log m/\log(2\log m) \rceil$ subintervals: For $j \in [\lceil \log m/\log(2\log m) \rceil]$, let

$$I_j = \left( \frac{(2\log m)^{j-1}}{m}, \min\left\{ \frac{(2\log m)^j}{m}, 1 \right\} \right],$$

note that the minimum in the upper bound only affects the last interval. Let $u_j$ and $\ell_j$ denote the upper and lower end points of $I_j$ and observe that $u_j \le 2\log m \cdot \ell_j$ for all $j \in [\lceil \log m/\log(2\log m) \rceil]$.

Let $t$ denote the threshold picked by $f$ (in a randomized fashion). There must exist $k \in [\lceil \log m/\log(2\log m) \rceil]$ such that $\Pr[t \in I_k] \le \log(2\log m)/\log m$. Fix a subset $S \subseteq A$ of size $\log m$, and let $V = u_k/2 + (\log m - 1) \cdot \ell_k$. Construct a (partial) utility profile $\vec{v}$ such that for each voter $i \in N$, $v_i(a) \in I_k$ for $a \in S$, $\sum_{a \in S} v_i(a) = V$, and $v_i(a) = (1-V)/(m - \log m)$ for $a \in A \setminus S$. First, this is feasible because

$$V = \frac{u_k}{2} + (\log m - 1) \cdot \ell_k \le \frac{1}{2} + \frac{\log m - 1}{2\log m} \le 1.$$

Second, this partial description completely dictates the induced input profile when $t \notin I_k$. Because $f$ can only distinguish between alternatives in $S$ when $t \in I_k$, there must exist $a^* \in S$ such that $\Pr[f \text{ returns } a^* | t \notin I_k] \le 1/\log m$. Suppose the underlying utility profile $\vec{v}$ satisfies, for each voter $i \in N$, $v_i(a^*) = u_k/2$ and $v_i(a) = \ell_k$ for $a \in S \setminus \{a^*\}$. Observe that this is consistent with the partial description provided before.

In this case, the optimal social welfare is given by $\mathrm{sw}(a^*, \vec{v}) = n \cdot u_k/2$, whereas $\mathrm{sw}(a, \vec{v}) \le n \cdot \ell_k$ for all $a \in A \setminus \{a^*\}$. The latter holds because $\ell_k > (1-V)/(m - \log m)$. The expected social welfare achieved by $f$ under $\vec{v}$ is at most

$$\begin{aligned} &\Pr[t \in I_k] \cdot \frac{n \cdot u_k}{2} \\ &+ \Pr[t \notin I_k] \times \left( \frac{1}{\log m} \cdot \frac{n \cdot u_k}{2} + \frac{\log m - 1}{\log m} \cdot n \cdot \ell_k \right) \\ &\le \frac{\log(2\log m) + 2}{\log m} \cdot \frac{n \cdot u_k}{2}, \end{aligned}$$

where the final transition holds because $u_k \le 2\log m \cdot \ell_k$. Thus, the distortion achieved by $f$ is $\Omega(\log m/\log\log m)$, as desired. □

Our proof of Theorem 6 establishes a lower bound of $\Omega(\log m/\log\log m)$ on the distortion associated with randomized threshold approval votes by only using the special case of the participatory budgeting problem in which $c_a = 1$ for each $a \in A$, that is, exactly one alternative needs to be selected. This is exactly the setting studied by Boutilier et al. (2015). On the other hand, Theorem 5 establishes a slightly weaker upper bound of $\mathcal{O}(\log^2 m)$ for the general participatory budgeting problem. For the restricted setting of Boutilier et al. (2015), one can improve the general $\mathcal{O}(\log^2 m)$ upper bound to $\mathcal{O}(\log m)$, thus leaving a very narrow gap from the $\Omega(\log m/\log\log m)$ lower bound. This proof is similar to the proof of Theorem 5, whose $\mathcal{O}(\log^2 m)$ bound is the result of a randomization over $\mathcal{O}(\log m)$ partitions of the alternatives based on their cost and $\mathcal{O}(\log m)$ possible values of the threshold. When costs are identical, there is no need to partition based on cost, reducing the partitions by a logarithmic factor. We defer the proof to the appendix.

**Theorem 7.** *If $c_a = 1$ for all $a \in A$, the distortion associated with randomized threshold approval votes is $\mathcal{O}(\log m)$.*

### 3.2. Deterministic Aggregation Rules

We next study the distortion that can be achieved under different input formats if we are forced to use a deterministic aggregation rule. Recall that the distortion associated with deterministic aggregation of votes under an input format is the least distortion a deterministic aggregation rule for that format can achieve. Specifically, we study the distortion associated with deterministic aggregation of knapsack votes, rankings by value and value for money, and deterministic threshold approval votes. We omit randomized threshold approval votes as the inherent randomization involved in elicitation makes the use of deterministic aggregation rules hard to justify.

We find that rankings by value achieve $\Theta(m^2)$ distortion, which is significantly better than the distortion of knapsack votes (exponential in $m$) and that of rankings by value for money (unbounded). This separation between rankings by value and value for money in this setting stands in stark contrast to the setting with randomized aggregation rules, where both input formats admit similar distortion. One important fact, however, does not change with the use of deterministic aggregation rules: using threshold approval votes still performs at least as well as using any of the other input formats considered here. Specifically, we show that setting the threshold to be $t = 1/m$ results in $\mathcal{O}(m^2)$ distortion. The choice of the threshold is crucial as, for example, setting a slightly higher threshold $t > 1/(m-1)$ results in unbounded distortion.

#### 3.2.1. Knapsack Votes.

Our first result is an exponential lower bound on the distortion associated with knapsack votes when the aggregation rule is deterministic. Although our construction requires the number of voters to be extremely large compared with the number of alternatives, we remark that this is precisely the case in real participatory budgeting elections, in which a large number of citizens vote over much fewer projects.

**Theorem 8.** *For sufficiently large $n$, the distortion associated with deterministic aggregation of knapsack votes is $\Omega(2^m/\sqrt{m})$.*

**Proof of Theorem 8.** Imagine a case where every alternative has cost $2/m$ (recall that the budget is 1). It follows that no more than $\lfloor m/2 \rfloor$ alternatives may be selected while respecting the budget constraints. Let $S_1, \ldots, S_{\binom{m}{\lfloor m/2 \rfloor}}$ denote the $\binom{m}{\lfloor m/2 \rfloor}$ subsets of $A$ of size $\lfloor m/2 \rfloor$.

Assume $n \geq \binom{m}{\lfloor m/2 \rfloor}$ and partition the voters into $\binom{m}{\lfloor m/2 \rfloor}$ sets $N_1, \ldots, N_{\binom{m}{\lfloor m/2 \rfloor}}$, each consisting of roughly $n/\binom{m}{\lfloor m/2 \rfloor}$ voters; specifically, ensure that $\lfloor n/\binom{m}{\lfloor m/2 \rfloor} \rfloor \leq n_i \leq \lceil n/\binom{m}{\lfloor m/2 \rfloor} \rceil$, where $n_i = |N_i|$, for all $i \in [\binom{m}{\lfloor m/2 \rfloor}]$. Construct an input profile of knapsack votes $\vec{\kappa}$, where $\kappa_i = S_k$ for all $k \in [\binom{m}{\lfloor m/2 \rfloor}]$ and $i \in N_k$.

Let $f$ denote a deterministic aggregation rule. We can safely assume that $|f(\vec{\kappa})| = \lfloor m/2 \rfloor$ as otherwise we can add alternatives to $f(\vec{\kappa})$, which can only improve the distortion. Let $f(\vec{\kappa}) = S_{k^*}$.

Construct a utility profile $\vec{v}$ consistent with the input profile $\vec{\kappa}$ as follows. Fix $b \in S_{k^*}$, and for all $i \in N_{k^*}$, let $v_i(b) = 1$ and $v_i(a) = 0$ for all $a \in A \setminus \{b\}$. Note that these valuations are consistent with the votes of voters in $N_{k^*}$.

Next, fix $a^* \in A \setminus S_{k^*}$. Our goal is to make $a^*$ an attractive alternative that $f(\vec{\kappa})$ missed. Note that $a^*$ appears in about half of the $\lfloor m/2 \rfloor$-sized subsets of $A$. For all $k \in [\binom{m}{\lfloor m/2 \rfloor}]$ such that $a^* \in S_k$, and all voters $i \in N_k$, let $v_i(a^*) = 1$ and $v_i(a) = 0$ for all $a \in A \setminus \{a^*\}$. This ensures $\mathrm{sw}(a^*, \vec{v}) \geq n \cdot \lfloor m/2 \rfloor / m \geq n/3$ (for $m \geq 2$).

For $k \in [\binom{m}{\lfloor m/2 \rfloor}] \setminus \{k^*\}$ such that $a^* \notin S_k$, and all voters $i \in N_k$, let $v_i(a') = 1$ for some $a' \in S_k \setminus S_{k^*}$, and $v_i(a) = 0$ for all $a \in A \setminus \{a'\}$.

Observe that all voters who do not belong to $N_{k^*}$ assign zero utility to all the alternatives in $S_{k^*}$, yielding $\mathrm{sw}(f(\vec{\kappa}), \vec{v}) \leq n_{k^*} \leq n/\binom{m}{\lfloor m/2 \rfloor} + 1$. By assumption, $n \geq \binom{m}{\lfloor m/2 \rfloor}$, so we have

$$\mathrm{dist}(f, \vec{v}) \geq \frac{n/3}{n/\binom{m}{\lfloor m/2 \rfloor} + 1} = \frac{1}{6} \cdot \binom{m}{\lfloor m/2 \rfloor} = \Omega\left(\frac{2^m}{\sqrt{m}}\right),$$

as required. □

We next show that an almost matching upper bound can be achieved by the natural plurality knapsack rule that selects the subset of alternatives submitted by the largest number of voters.

**Theorem 9.** *The distortion associated with deterministic aggregation of knapsack votes is* $\mathcal{O}(m \cdot 2^m)$.

**Proof of Theorem 9.** Let $\vec{v}$ denote the underlying utility profile, and let $S^* \subseteq A$ be the set of alternatives reported by the largest number of voters. Due to the pigeonhole principle, it must be reported by at least $n/2^m$ voters. Further, each voter $i$ who reports $S^*$ must have $v_i(S^*) \geq 1/m$ because there must exist $a \in A$ such that $v_i(a) \geq 1/m$, and $v_i(S^*) \geq v_i(a)$.

Hence, we have $\mathrm{sw}(S^*, \vec{v}) \geq (n/2^m) \cdot 1/m$, whereas the maximum welfare any set of alternatives can achieve is at most $n$. Hence, the distortion of the proposed rule is at most $m \cdot 2^m$. □

#### 3.2.2. Rankings by Value and by Value for Money.

While rankings by value and by value for money have similar distortion in case of randomized aggregation rules, deterministic aggregation rules lead to a clear separation between the distortion of the two input formats.

We first show that deterministic aggregation of rankings by value for money cannot offer bounded distortion. Our counterexample exploits the uncertainty in values induced when alternatives have vastly different costs.

**Theorem 10.** *The distortion associated with deterministic aggregation of rankings by value for money is unbounded.*

**Proof of Theorem 10.** Fix $a, b \in A$. Let $c_a = \epsilon > 0$, and $c_k = 1$ for all $k \in A \setminus \{a\}$. Recall that the budget is 1. Hence, every deterministic aggregation rule must select a single alternative.

Construct an input profile $\vec{\sigma}$ in which each input ranking has alternatives $a$ and $b$ in positions 1 and 2, respectively. Let $f$ be a deterministic aggregation rule.

If $f(\vec{\sigma}) \in A \setminus \{a\}$, the utility profile $\vec{v}$ in which every voter has utility 1 for $a$, and 0 for every alternative in $A \setminus \{a\}$ ensures $\mathrm{dist}(f) \geq \mathrm{dist}(f, \vec{v}) = \infty$.

If $f(\vec{\sigma}) = a$, the utility profile $\vec{v}$ in which every voter has utility $\epsilon$ for $a$, $1 - \epsilon$ for $b$, and 0 for every alternative in $A \setminus \{a, b\}$ ensures that $\mathrm{dist}(f) \geq \mathrm{dist}(f, \vec{v}) = (1-\epsilon)/\epsilon$.

Hence, in either case, $\mathrm{dist}(f) \geq (1-\epsilon)/\epsilon$. Because $\epsilon$ can be arbitrarily small, the distortion is unbounded. □

We now turn our attention to rankings by value. Caragiannis et al. (2016) study deterministic aggregation of rankings by value in the special case of our setting where the cost of each alternative equals the entire budget, and establish a lower bound of $\Omega(m^2)$ on the distortion, which carries over to our more general setting.

**Theorem 11** (Caragiannis et al. 2016). *For* $n \geq m - 1$, *the distortion associated with deterministic aggregation of rankings by value is* $\Omega(m^2)$.

Caragiannis et al. (2016) also show that selecting the plurality winner—the alternative that is ranked first by the largest number of voters—results in distortion at most $m^2$. We show that this holds true even in our more general setting, giving us an asymptotically tight bound on the distortion.

**Theorem 12.** *The distortion associated with deterministic aggregation of rankings by value is* $\mathcal{O}(m^2)$.

**Proof of Theorem 12.** Due to the pigeonhole principle, the plurality winner, say $a \in A$, must be ranked first by at least $n/m$ voters, each of which must have utility at least $1/m$ for $a$. Hence, the social welfare of $a$ is at least $n/m^2$, while the maximum social welfare that any set of alternatives can achieve is at most $n$, yielding a distortion of at most $m^2$. □

#### 3.2.3. Threshold Approval Votes.

We now turn our attention to threshold approval votes. As mentioned earlier, our use of deterministic aggregation rules makes randomized threshold selection less motivated; we thus focus on deterministic threshold approval votes.

First, we show that for some choices of the threshold, the distortion can be unbounded.

**Theorem 13.** *For a fixed threshold* $t > 1/(m-1)$, *the distortion associated with deterministic aggregation of deterministic threshold approval votes is unbounded.*

**Proof of Theorem 13.** Suppose $c_a = 1$ for each $a \in A$. Recall that the budget is 1. Let $f$ denote a deterministic aggregation rule for threshold approval votes. Suppose the rule receives an input profile $\vec{\tau}$ in which no voter approves any alternative. Without loss of generality, let $f(\vec{\tau}) = a^*$.

We construct an underlying utility profile such that for each voter $i \in N$, $v_i(a) = 1/(m-1)$ for $a \in A \setminus \{a^*\}$, and $v_i(a^*) = 0$. Note that this is consistent with $\vec{\tau}$. Now, the optimal social welfare is $n \cdot 1/(m-1)$, whereas the welfare achieved by $f$ is zero, yielding an unbounded distortion. □

We next show that slightly reducing the threshold to $1/m$ reduces the distortion to $\mathcal{O}(m^2)$, which is at least as good as the distortion associated with any other input format.

**Theorem 14.** *For the fixed threshold* $t = 1/m$, *the distortion associated with deterministic aggregation of deterministic threshold approval votes is* $\mathcal{O}(m^2)$.

**Proof of Theorem 14.** Let $\vec{\tau}$ denote an input profile, and $\vec{v}$ the underlying utility profile. Let $S^* \in \mathcal{F}_c$ denote the feasible set of alternatives with the highest number of total approvals. The set $S \in \mathcal{F}_c$ is returned by the following algorithm: label the alternatives in order of the number of approvals received to cost, where $a_1$ has the greatest ratio. Return whichever of $\{a_1, \ldots, a_{k-1}\}$ and $\{a_k\}$ has more approvals, with $k$ chosen so that

$\{a_1, \ldots, a_{k-1}\} \in \mathcal{F}_c$ and $\{a_1, \ldots, a_k\} \notin \mathcal{F}_c$. Let $P^*$ and $P$ denote the total number of approvals received by alternatives in $S^*$ and $S$, respectively.

Consider a knapsack problem where the value of an alternative is the number of approvals it receives under $\tau$. Then, $P^*$ is the optimal knapsack solution, whereas $P$ is the solution quality achieved by the greedy algorithm. Using the fact that this algorithm achieves a 2-approximation of the (unbounded) knapsack problem (Dantzig 1957), we have

$$P \geq (1/2) \cdot P^*.$$

We can now establish an upper bound on the distortion of our rule. Let $T$ be the feasible set of alternatives maximizing the social welfare. Then, $T$ achieves at most $P^*$ total approvals under $\vec{\tau}$. Each approval of an alternative in $T$ by a voter can contribute at most 1 to the welfare of $T$, and each non-approval of an alternative in $T$ by a voter can contribute at most $1/m$ to the welfare of $T$. Hence, we have

$$\mathrm{sw}(T, \vec{v}) \leq P^* \cdot 1 + (n \cdot m - P^*) \cdot (1/m).$$

Using a similar line of argument, we also have

$$\mathrm{sw}(S, \vec{v}) \geq P \cdot (1/m).$$

Hence, the distortion of $f$ is at most

$$\frac{P^* + (n \cdot m - P^*)/m}{P/m} \leq 2 \cdot \frac{1 + (n \cdot m/P^* - 1)/m}{1/m} = 2 \cdot \left(m + \frac{n \cdot m}{n/m} - 1\right) = \mathbb{O}(m^2),$$

where the first transition follows from $P \geq P^*/2$. For the second transition, note that with the threshold being $1/m$, each voter must approve at least one alternative. Hence, there must exist an alternative with at least $n/m$ approvals, implying that $P^* \geq n/m$. □

## 4. Computing Worst-Case Optimal Aggregation Rules

Our theoretical results focus on the best worst-case (over all input profiles) distortion we can achieve using different input formats. However, specific input profiles may admit distortion much better than this worst-case bound. In practice, we are more interested in the deterministic or randomized aggregation rule that, on each input profile, returns the feasible set of alternatives or a distribution thereover which minimizes distortion, thus achieving the optimal distortion on each input profile individually. The optimal deterministic aggregation rule is given by

$$f^*(\vec{\rho}) = \arg\min_{S \in \mathcal{F}_c} \max_{\vec{v} \triangleright \vec{\rho}} \frac{\max_{T \in \mathcal{F}_c} \mathrm{sw}(T, \vec{v})}{\mathrm{sw}(S, \vec{v})}, \ \forall \vec{\rho},$$

and the optimal randomized aggregation rule is given by

$$\bar{f}^*(\vec{\rho}) = \arg\min_{p \in \Delta(\mathcal{F}_c)} \max_{\vec{v} \triangleright \vec{\rho}} \frac{\max_{T \in \mathcal{F}_c} \mathrm{sw}(T, \vec{v})}{\mathbb{E}[\mathrm{sw}(p, \vec{v})]}, \ \forall \vec{\rho},$$

where $\Delta(X)$ denotes the set of distributions over the elements of $X$.

Although these profile-wise optimal aggregation rules dominate all other aggregation rules, they may be computationally difficult to implement, because they optimize a nonlinear objective function (a ratio) over a complicated space.

We design practical generic algorithms for computing the deterministic and randomized profile-wise optimal aggregation rules for the input formats we study. The deterministic rule reformulates this problem as a linear-fractional program, which is recast to a linear program with the Charnes-Cooper transformation (Charnes and Cooper 1962). The procedure for the randomized aggregation rule is a two-stage algorithm in the spirit of the cutting-set approach of Mutapcic and Boyd (2009). More details can be found in Online Appendix EC.2.

## 5. Empirical Results

Our theoretical results in Section 3 bound the distortion of an observed input profile. Recall that distortion is the worst-case ratio of the optimal social welfare to the social welfare achieved, where the worst case is taken over all utility profiles consistent with the observed input profile. In practice, we care about this welfare ratio according to the actual underlying utility profile. Thus, a distortion-minimizing aggregation rule is not guaranteed to be optimal for a specific utility profile. Furthermore, it may be the case that for some input format almost all utility profiles lead to welfare ratios close to the distortion, whereas for another only a very small fraction of degenerate utility profiles have a welfare ratio close to the distortion and all the remaining utility profiles have a welfare ratio close to one, that is, lead to almost optimal outcomes. Our empirical study attempts to determine what welfare ratio we should expect from a "usual" utility profile—in other words, how closely we are able to approximate the optimal social welfare in practice.

### 5.1. Evaluation Metrics

We are interested in efficiency and scalability. For efficiency, we measure the empirical average (across instances) of the welfare ratio—the ratio of the optimal social welfare to that achieved by the aggregation rule (which only has access to the input profile, not the utility profile). Note that this ratio is different from distortion, which is a worst-case guarantee on the social welfare ratio across all utility profiles that could have induced an input profile. For scalability, we report the time it takes to compute the outcome (in seconds).

### 5.2. Data Sets

We use data from participatory budgeting elections held in 2015 and 2016 in Boston, Massachusetts. Both elections offered voters 10 alternatives. The 2015 data set contains 2,600 four-approval votes (voters were asked to approve their four most preferred alternatives) and the 2016 data set contains 4,430 knapsack votes.

For each data set, we conduct three independent trials. In each trial, we create $r$ subprofiles, each consisting of $n$ voters drawn at random from the population. For each subprofile, we draw $k$ random utility profiles $\vec{v}$ consistent with the subprofile, and use these to analyze the performance of different approaches. We use the real costs of the projects throughout. The choices of parameters $(r, n, k)$ for the three trials are $(5, 10, 10)$, $(8, 7, 10)$, and $(10, 5, 10)$. We choose this experimental design to yield sufficiently many samples to verify statistical significance of the results while completing in a reasonable amount of time. The bottleneck in these experiments is the computation of the randomized distortion-minimizing voting rules, which is very slow. Fortunately, their deterministic counterparts generally lead to outcomes with higher social welfare and scale to realistic instance sizes (see Figures 1 and 2).

### 5.3. Approaches

We use the utility profile $\vec{v}$ drawn to create an input profile in the four input formats we study. For each format, we use the deterministic as well as randomized distortion-minimizing aggregation rule.

**Figure 1.** (Color online) Average Welfare Ratio of Different Approaches to Participatory Budgeting

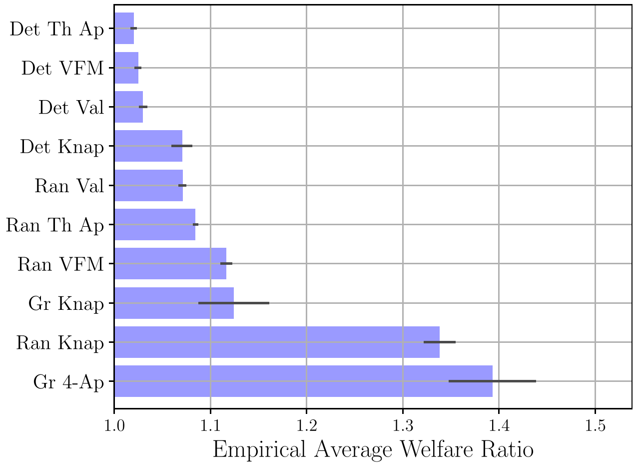


*Note.* A lower welfare ratio is better.

**Figure 2.** (Color online) Average Running Time of the Deterministic Voting Rules on the Boston 2016 Data Set

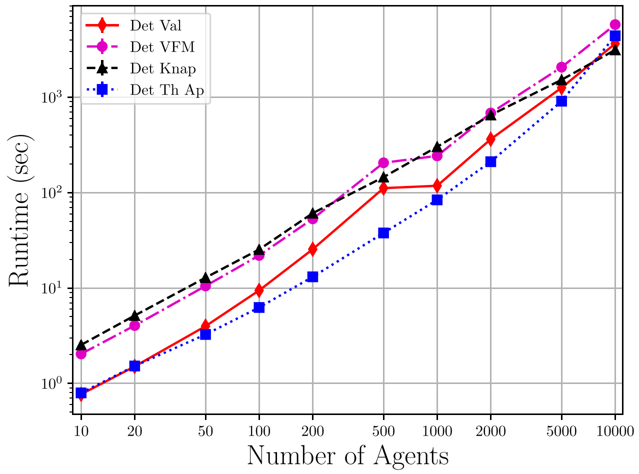


The nontrivial algorithms we devise for these rules are presented in Section 4. These eight approaches are referred to using the type of aggregation rule used (Det or Ran), and the type of input format (Knap, Val, VFM, or Th Ap).

Unlike the other input formats, threshold approval votes are technically a family of input formats, one for each value of the threshold. Although randomizing over the threshold is required to minimize the distortion (the worst-case ratio of the optimal and achieved social welfare), as is our goal in the theoretical results of Section 3, minimizing the expected ratio of the two can be achieved by a deterministic threshold. In our experiments, we learn the optimal threshold value based on a holdout set that is not subsequently used. This learning approach is practical as it only uses observed input votes rather than underlying actual utilities. This choice likely gives threshold approval votes an edge, but arguably it is an advantage this input format would also enjoy in practice.

In addition to our eight approaches, we also test two approaches used in real-world elections (Goel et al. 2019): greedy 4-approval (Gr 4-Ap) and greedy knapsack (Gr Knap). The former elicits 4-approval votes, and greedily selects the most widely approved alternatives until the budget is depleted. The latter is almost identical, except for interpreting a knapsack vote as an approval for each alternative in the knapsack.

### 5.4. Results

Figure 1 shows the empirical average of the welfare ratio of the different approaches with 95% confidence intervals, sorted from best to worst. The differences in performance between all pairs of rules—except between Det Knap and Ran Val, and between Ran VFM and Gr Knap—are statistically significant (Johnson 2013) at a 95% confidence level.

A few comments are in order. First, deterministic distortion-minimizing aggregation rules generally outperform their randomized counterparts. This is not entirely unexpected. Randomized rules achieve better distortion, that is better worst-case guarantees when the utility profile is unknown. When the utility profile is drawn from a distribution, as it is here, there exists a deterministic rule minimizing the expected welfare ratio objective.

Second, approaches based on deterministic rules are able to limit the loss in social welfare due to incomplete information about voters' utility functions to only 2%–3%. Among these approaches, the one using threshold approval votes incurs the minimum loss.

Third, knapsack votes consistently lead to higher empirical welfare ratios than alternative input formats. This, together with the poor theoretical guarantees for knapsack votes, suggests that it may not be worthwhile to ask voters to solve their personal $\mathcal{NP}$-hard knapsack problems before casting vote.

### 5.5. Scalability

We are particularly interested in deterministic distortion-minimizing rules, both because of their superior empirical performance (see Figure 1) and the fact that our discussions with officials from several cities have revealed a hesitance to use randomized voting rules. Figure 2 reports the average time to compute the deterministic worst-case optimal set of alternatives on a log-log scale, averaged over 20 trials on the Boston 2016 data set. Since this data set contains only 4,430 votes, voters were sampled with replacement to mimic larger instances. The experiments were run on an eight-core Intel Core i7-7700 CPU with 3.6 GHz processor speed and 65 gigabytes memory.

Observe that the running time scales gracefully with the number of voters. When sampling 10,000 voters with replacement, rules such as Det Th Ap and Det Val take less than three hours, indicating the practicability of these methods for participatory budgeting elections at the scale of those in Boston (Goel et al. 2019). We also note that, due to the one-off nature of participatory budgeting elections, it is conceivable to use an aggregation algorithm that takes several days or even weeks to compute the optimal set of alternatives.

## 6. Discussion

Our results indicate that approval voting should receive serious consideration as the input format of choice for participatory budgeting, and that knapsack voting may not perform well with respect to social welfare. We also observe that optimization-based aggregation leads to outcomes with significantly higher social welfare than greedy aggregation.

One advantage of traditional greedy aggregation methods over distortion-minimizing aggregations is that it is easy to publish the summarized votes in such a manner that any voter can verify the outcome. It is unclear to what extent optimization-based aggregation is hindered by this relative lack of transparency, however, the recent adoption of ranked-choice voting in New York City and several Democratic primaries (Fortin 2020) suggests that voters are increasingly comfortable with elections in which verifying the outcome is nontrivial.

It bears repeating that distortion, and the worst-case analysis in the presence of only ordinal information employed here, can be applied to many combinatorial optimization problems including maximum weight spanning tree and maximum *b*-matching.

Some issues that are beyond the scope of this paper have been addressed in subsequent works. First, central to our approach is the cognitive burden an input format places on a voter (were it not for this, we would elicit full utility functions; the whole point is to reduce cognitive load). Benadè et al. (2018) conduct human subject experiments to measure how difficult voters find different input formats, and find that although threshold approval voting is slightly harder to understand at first, it is no more difficult to use than the other common input formats. Second, we show that asking a voter about a single threshold is sufficient to achieve low distortion. However, one may ask whether there is a benefit to asking voters to reflect on two, 10, or 100 different thresholds. Mandal et al. (2019, 2020) answer this question by studying the trade-off between the number of bits an input format requires from a voter, and the distortion it achieves. Finally, Bhaskar et al. (2018) show that the voting rule we use in the proof of Theorem 3, which relies on the harmonic scoring rule, is truthful. They also observe that independently (uniformly) randomizing every voter's threshold leads to distortion approaching one.

Whatever the best approach to participatory budgeting is, now is the time to identify it, before various heuristics become hopelessly ingrained. We believe that this is a grand challenge for computational social choice, especially at a point in the field's evolution where it is gaining real-world relevance by helping people make decisions in practice.

### Acknowledgments

The authors thank Ashish Goel and Anilesh Krishnaswamy of the Stanford Crowdsourced Democracy Team for generously sharing their data.

### Endnotes

[1] See http://www.participatorybudgeting.org, accessed December 1, 2019.

[2] See http://voxpopuli.stanford.edu, accessed December 1, 2019.

## References

Abramowitz B, Anshelevich E (2018) Utilitarians without utilities: Maximizing social welfare for graph problems using only ordinal preferences. *Proc. 32nd AAAI Conf. Artificial Intelligence* (AAAI, Palo Alto, CA), 894–901.

Anshelevich E, Postl J (2016) Randomized social choice functions under metric preferences. *Proc. 25th Internat. Joint Conf. Artificial Intelligence* (IJCAI, Pasadena, CA), 46–52.

Anshelevich E, Sekar S (2016) Blind, greedy, and random: Algorithms for matching and clustering using only ordinal information. *Proc. 30th AAAI Conf. Artificial Intelligence* (AAAI, Palo Alto, CA), 390–396.

Anshelevich E, Bhardwaj O, Postl J (2015) Approximating optimal social choice under metric preferences. *Proc. 29th AAAI Conf. Artificial Intelligence* (AAAI, Palo Alto, CA), 777–783.

Arrow K (1951) *Social Choice and Individual Values* (John Wiley & Sons, Hoboken, NJ).

Aziz H (2019) Justifications of welfare guarantees under normalized utilities. *SIGecom Exchanges* 17(2):71–75.

Baumeister D, Bouveret S, Lang J, Nguyen N, Nguyen T, Rothe J, Saffidine A (2017) Positional scoring-based allocation of indivisible goods. *J. Autonomous Agents Multi-Agent Systems* 31(3): 628–655.

Benadè JG, Itzhak N, Shah N, Procaccia AD, Gal Y (2018) Efficiency and usability of participatory budgeting methods. Working paper, Boston University, Boston.

Bhaskar U, Dani V, Ghosh A (2018) Truthful and near-optimal mechanisms for welfare maximization in multi-winner elections. *Proc. 32nd AAAI Conf. Artificial Intelligence* (AAAI, Palo Alto, CA), 925–932.

Boutilier C, Caragiannis I, Haber S, Lu T, Procaccia AD, Sheffet O (2015) Optimal social choice functions: A utilitarian view. *Artificial Intelligence* 227:190–213.

Bouveret S, Lang J (2011) A general elicitation-free protocol for allocating indivisible goods. *Proc. 22nd Internat. Joint Conf. Artificial Intelligence* (IJCAI, Pasadena, CA), 73–78.

Brams SJ, Fishburn PC (2007) *Approval Voting*, 2nd ed. (Springer, New York).

Brandt F, Conitzer V, Endress U, Lang J, Procaccia AD, eds. (2016) *Handbook of Computational Social Choice* (Cambridge University Press, Cambridge, UK).

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

Caragiannis I, Procaccia AD (2011) Voting almost maximizes social welfare despite limited communication. *Artificial Intelligence* 175(9–10):1655–1671.

Caragiannis I, Nath S, Procaccia AD, Shah N (2016) Subset selection via implicit utilitarian voting. *Proc. 25th Internat. Joint Conf. Artificial Intelligence* (IJCAI, Pasadena, CA), 151–157.

Chakrabarty D, Swamy C (2014) Welfare maximization and truthfulness in mechanism design with ordinal preferences. *Proc. 5th Conf. Innovations Theoret. Comput. Sci.* (ITCS) (LIPIcs, Leibniz, Germany), 105–120.

Charnes A, Cooper WW (1962) Programming with linear fractional functionals. *Naval Res. Logist. Quart.* 9(3–4):181–186.

Conitzer V, Sandholm T (2005) Common voting rules as maximum likelihood estimators. *Proc. 21st Annual Conf. Uncertainty Artificial Intelligence* (UAI) (AUAI, Corvallis, OR), 145–152.

Conitzer V, Freeman R, Shah N (2017) Fair public decision making. *Proc. 18th ACM Conf. Econom. Comput.* (EC) (ACM, New York), 629–646.

Dantzig GB (1957) Discrete-variable extremum problems. *Oper. Res.* 5(2):266–288.

Fain B, Munagala K, Shah N (2018) Fair allocation of indivisible public goods. *Proc. 19th ACM Conf. Econom. Comput.* (EC) (ACM, New York), 575–592.

Filos-Ratsikas A, Miltersen P (2014) Truthful approximations to range voting. *Proc. 10th Conf. Web Internet Econom.* (WINE) (Springer, Cham), 175–188.

Filos-Ratsikas A, Frederiksen SKS, Zhang J (2014) Social welfare in one-sided matchings: Random priority and beyond. *Proc. 7th Internat. Symp. Algorithmic Game Theory* (SAGT), 1–12.

Fortin J (2020) Why ranked-choice voting is having a moment. *New York Times* (February 10), https://www.nytimes.com/2020/02/10/us/politics/ranked-choice-voting.html.

Goel A, Krishnaswamy AK, Sakshuwong S, Aitamurto T (2019) Knapsack voting for participatory budgeting. *ACM Trans. Econom. Comput.* 7(2):1–27.

Gutiérrez B (2017) Madrid as a democracy laboratory. *Open Democracy* (July 10), https://www.opendemocracy.net/democraciaabierta/bernardo-guti-rrez/madrid-as-democracy-lab.

Johnson VE (2013) Revised standards for statistical evidence. *Proc. Natl. Acad. Sci. USA* 110(48):19313–19317.

Krysta P, Manlove D, Rastegari B, Zhang J (2014) Size vs. truthfulness in the house allocation problem. *Proc. 15th ACM Conf. Econom. Comput.* (EC) (ACM, New York), 453–470.

Legendre J, Madénian H, Scully PL (2017) Participatory budgeting in Paris, France. Accessed February 5, 2018, http://participedia.net/en/cases/participatory-budgeting-paris-france.

Lu T, Boutilier C (2011) Budgeted social choice: From consensus to personalized decision making. *Proc. 22nd Internat. Joint Conf. Artificial Intelligence* (IJCAI, Pasadena, CA), 280–286.

Mandal D, Shah N, Woodruff DP (2020) Optimal communication-distortion tradeoff in voting. *Proc. 21st ACM Conf. Econom. Comput.* (ACM, New York). Forthcoming.

Mandal D, Procaccia AD, Shah N, Woodruff DP (2019) Efficient and thrifty voting by any means necessary. *Proc. 33rd Annual Conf. Neural Inform. Processing Systems* (NIPS) (NeurlPS, San Diego, CA), 7178–7189.

Moulin H (2003) *Fair Division and Collective Welfare* (MIT Press, Cambridge, MA).

Mutapcic A, Boyd SP (2009) Cutting-set methods for robust convex optimization with pessimizing oracles. *Optim. Methods Software* 24(3):381–406.

Procaccia AD, Rosenschein JS (2006) The distortion of cardinal preferences in voting. *Proc. 10th Internat. Workshop Cooperative Inform. Agents* (CIA) (Springer, Berlin, Heidelberg), 317–331.

Procaccia AD, Reddi SJ, Shah N (2012) A maximum likelihood approach for selecting sets of alternatives. *Proc. 28th Annual Conf. Uncertainty Artificial Intelligence* (UAI) (AUAI, Corvallis, OR), 695–704.

Shah A (2007) *Participatory Budgeting* (The World Bank, Washington, DC).

Skowron P, Faliszewski P, Lang J (2015) Finding a collective set of items: From proportional multirepresentation to group recommendation. *Proc. 29th AAAI Conf. Artificial Intelligence* (AAAI, Palo Alto, CA), 2131–2137.

Young HP (1988) Condorcet's theory of voting. *Amer. Political Sci. Rev.* 82(4):1231–1244.