Preference Elicitation For Participatory Budgeting*

Gerdus Benadè†1, Ariel D. Procaccia‡1, Swaprava Nath§2, and Nisarg Shah¶3

[1]Carnegie Mellon University, USA
[2]IIT Kanpur, India
[3]University of Toronto, Canada

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 a rethinking of some of the basics of computational social choice, including the very way in which individuals express their preferences. We analytically compare four preference elicitation methods — knapsack votes, rankings by value or value for money, and threshold approval votes — through the lens of *implicit utilitarian voting*, and find that threshold approval votes are qualitatively superior. This conclusion is supported by experiments using data from real participatory budgeting elections.

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 which 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.

*This is an extended version of the paper 'Preference Elicitation for Participatory Budgeting' which appeared in the proceedings of the 31st Conference of the Association for the Advancement of Artificial Intelligence (AAAI 2017) [Benadè *et al.*, 2017]. This paper contains several new results, specifically, we study the distortion of deterministic aggregation rules in addition to randomized ones, and propose algorithms for the computation of these rules which have not appeared in print elsewhere.

†jbenade@andrew.cmu.edu
‡arielpro@cs.cmu.edu
§swaprava@cse.iitk.ac.in
¶nisarg@cs.toronto.edu

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* (`http://www.participatorybudgeting.org`) has helped allocate more than $170 million dollars of public money for more than 500 local projects, primarily in the US and Canada (including New York City, Chicago, Boston, and San Francisco).

Paraticipatory 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 Euro on participatory budgets in 2017 [Legendre *et al.*, 2017; Gutiérrez, 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. [2016] — who have facilitated several participatory budgeting elections as part of the Stanford Crowdsourced Democracy Team (`http://voxpopuli.stanford.edu`) — 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 individual solution to the knapsack problem, that is, the set of projects that maximizes his 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 his 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. [2016] 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. [2016], 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 which maximize *(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; Boutilier *et al.*, 2015; Anshelevich *et al.*, 2015; Anshelevich and Sekar, 2016; Anshelevich and Postl, 2016]; it 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 achieving social welfare maximization (lower distortion is better).

In §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 each alternative whose value for him is 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.

It is worth noting that 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 is used to evaluate our performance.

While our theoretical results in §3 bound the distortion, i.e., 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 §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. We observe that the running times of these distortion-minimizing rules scale gracefully to practical sizes.

In §5 we use these algorithms to compare different approaches to participatory budgeting using the *average-case ratio* of the optimal social welfare, and the social welfare achieved by our aggregation rules. Specifically, we experimentally evaluate approaches that use the input formats we study in conjunction with their respective optimal aggregation rules, which minimize the distortion on each profile, and compare them to two approaches currently employed in practice. (Note that these rules are not guaranteed to achieve the optimal performance in our experiments as we measure performance using the *average-case* ratio of the optimal to the achieved social welfare rather than the (worst-case) distortion. Nonetheless, such 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 the use of aggregation rules that minimize distortion on every input profile significantly outperforms the currently deployed approaches, and among the input formats we study, threshold approval votes remain superior, even in practice.

1.2 Related Work

Let us first describe the theoretical results of Goel et al. [2016] 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 his own knapsack vote. In a sense, the viewpoint underlying this model is the opposite of ours, as a voter's utility is derived from his 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. [2016] 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 bounds 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.*, 2016] — whereas Caragiannis et al. assume that a voter's utility for a subset of alternatives is his *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] present algorithms for finding matchings in weighted graphs given ordinal comparisons among the edges by their weight; Krysta *et al.* [2014] apply this notion to the house allocation problem; and, Chakrabarty and Swamy [2014] study this notion 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 non-trivial 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 $\mathcal{F}_c = \{S \subseteq A : c(S) \leqslant 1 \ \wedge \ c(T) > 1, \ \forall S \subsetneq T \subseteq A\}$ 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, i.e., 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 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 $\text{sw}(S, \vec{v}) = \sum_{a \in S} \text{sw}(a, \vec{v})$.

The utility function of a voter i is only accessible through his vote ρ_i, which is induced by v_i. The vector $\vec{\rho} = \{\rho_1, \ldots, \rho_n\}$ is called the *input profile*. Let $\vec{v} \; \triangleright \; \vec{\rho}$ denote that utility profile $\vec{v}$ is *consistent* with input profile $\vec{\rho}$. We study four specific formats for input votes. Below, 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, 1000 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 $v_i \triangleright \kappa_i$ if and only if $c(\kappa_i) \leqslant 1$ and $v_i(\kappa_i) \geqslant v_i(S)$ for all $S \in \mathcal{F}_c$. 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 him, and by decreasing ratio of value for him to cost, respectively. Formally, let $\mathcal{L} = \mathcal{L}(A)$ denote the set of rankings over the alternatives. For a ranking $\sigma \in \mathcal{L}$, let $\sigma(a)$ denote the position of alternative a in σ, and $a \succ_\sigma b$ denote $\sigma(a) < \sigma(b)$, i.e., that a is preferred to b under σ. Then, we say that utility function v_i is consistent with the ranking by value (resp. value for money) of voter $i \in N$, denoted σ_i, if and only if $v_i(a) \geqslant v_i(b)$ (resp. $v_i(a)/c_a \geqslant v_i(b)/c_b$) for all $a \succ_{\sigma_i} b$. To elicit such votes, the voters may be asked: "If you had to divide 1000 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* τ_i of voter $i \in N$ consists of the set of alternatives whose value for him is at least t, i.e., $v_i \triangleright \tau_i$ if and only if $\tau_i = \{a \in A : v_i(a) \geqslant t\}$. To elicit threshold approval votes with a threshold $t = 1/10$, the voters may be asked: "If you had to divide 1000 points among the projects based on how much you like them, select all the projects that would receive at least 100 points."

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

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

The (overall) distortion of a rule f is given by $\text{dist}(f) = \max_{\vec{\rho}} \text{dist}(f, \vec{\rho})$. 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,

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

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 $f^*_{\texttt{det}}$ is given by

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

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 §3.1, we present theoretical results for the distortion associated with different input formats when no constraints are imposed on the aggregation rule, i.e., when randomized aggregation rules are allowed. Subsequently, in §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 3.1. *The distortion associated with any input format is at most m.*

Proof. 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. [2016].

Theorem 3.2. *For $n \geqslant m$, the distortion associated with knapsack votes is $\Omega(m)$.*

Proof. 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 $\lfloor n/m \rfloor \leqslant n_a \leqslant \lceil n/m \rceil$ 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 $\Pr[f(\vec{\kappa}) = \{a^*\}] \leqslant 1/m$. 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 $\vec{\kappa}$, i.e., $\vec{v} \triangleright \vec{\kappa}$. Moreover, it holds that $\text{sw}(a^*, \vec{v}) \geqslant n/2$, whereas $\text{sw}(a, \vec{v}) \leqslant n_a \leqslant n/m + 1$ for $a \in A \setminus \{a^*\}$. It follows that

$$\text{dist}(f) \geqslant \text{dist}(f, \vec{\kappa}) \geqslant \frac{n/2}{\frac{1}{m} \cdot n + \frac{m-1}{m} \cdot (\frac{n}{m} + 1)} \geqslant \frac{m}{6},$$

as desired. □

In light of Observation 3.1, 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. [2016] 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 also study rankings by value, which are prevalent in the existing literature on implicit utilitarian voting [Procaccia and Rosenschein, 2006; Caragiannis and Procaccia, 2011; Boutilier *et al.*, 2015; Anshelevich *et al.*, 2015; Anshelevich and Sekar, 2016; Anshelevich and Postl, 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.*, 2016].

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 $n \geqslant \sqrt{m}$. 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 3.3. *The distortion associated with rankings by value and rankings by value for money is* $\mathcal{O}(\sqrt{m}\ \log m)$.

Proof. 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 σ and an alternative $a \in A$, let $\sigma(a)$ denote the position of a in σ. 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 $\sigma|_S$ (resp. $\vec{\sigma}|_S$) denote the ranking (resp. preference profile) obtained by restricting σ (resp. $\vec{\sigma}$) 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: $S_0 = \{a \in A : c_a \leqslant u_0\}$ and $S_i = \{a \in A : \ell_i < c_a \leqslant u_i\}$ 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 $|S_i| \leqslant \sqrt{m} \cdot (1/u_i)$, 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 $\text{sw}(S^* \cap S^+, \vec{v})$.

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 $\text{sw}(S^* \cap S^-, \vec{v})$. Let us first bound $\text{sw}(S^* \cap S^-, \vec{v})$. Recall that $S_0^- = \emptyset$. Hence,

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

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

$$\sum_{b \in S_i} \text{sc}(b, \vec{\sigma}|_{S_i}) = n \cdot H_{|S_i|} \leqslant n \cdot H_m,$$

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

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

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

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

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

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

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 $v_j(a) \leqslant 2/\gamma_j(a)$ for $j \in [n]$, and taking a further sum over $j \in [n]$, we get

$$\text{sw}(a, \vec{v}) \leqslant 2 \cdot \text{sc}(a, \vec{\sigma}|_{S_i}). \tag{4}$$

Combining this with Equation (1), we get

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

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

$$\text{sw}(S^* \cap S^-, \vec{v}) = \sum_{i=1}^{\log m} \text{sw}(S^* \cap S_i^-, \vec{v}) \leqslant \frac{(m/2^{i-1}) \cdot 2 \cdot n \cdot (1 + \log m)}{\sqrt{m} \cdot m/2^i} = 4 \cdot n \cdot (1 + \log m)/\sqrt{m}. \tag{5}$$

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 $\text{sw}(S^* \cap S^-, \vec{v})$ 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 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 $v_i(a) \geqslant t$.

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 are approved, providing the rule little information to distinguish between the alternatives, and yielding high distortion.

Theorem 3.4. *The distortion associated with deterministic threshold approval votes is* $\Omega(\sqrt{m})$.

Proof. 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. (While 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 $t \leqslant 1/\sqrt{m}$. Let $A = \{a_1, \ldots, a_m\}$.

Suppose $t \leqslant 1/\sqrt{m}$. 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 $\Pr[f(\vec{\tau}) = \{a^*\}] \leqslant 1/(\sqrt{m}/2 + 1)$. 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 $t \leqslant 1/\sqrt{m}$. Further, $\text{sw}(a^*, \vec{v}) = n/2$ whereas $\text{sw}(a, \vec{v}) \leqslant n/\sqrt{m}$ for all $a \in A \setminus \{a^*\}$. Hence,

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

Because the optimal social welfare is $\Theta(n)$, we have that $\text{dist}(f) = \Omega(\sqrt{m})$, 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 $\Pr[f(\vec{\tau}) = a^*] \leqslant 1/m$. 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 $\text{sw}(a^*, \vec{v}) = n/\sqrt{m}$. In contrast, we have

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

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

For specific ranges of the threshold, it is possible to derive stronger lower bounds. However, the $\Omega(\sqrt{m})$ lower bound of Theorem 3.4 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:

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

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 the (overall) distortion of rule f as

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

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

Proof. 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,

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

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 $n_0^a \leqslant n$, we get

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

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 $n_{j,k}^* = \sum_{a \in S^* \cap S_k} n_j^a$. Using Equation (6), we have

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

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 ℓ_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 $|S_k| \leqslant 1/u_k$). Let $B_{j,k}$ denote the set of selected alternatives for the pair (j, k). Because we have $j > 0$ and $k > 0$,

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

where, in the first transition, we bound the welfare from below by only considering utilities that are at least ℓ_j, and the second transition holds because $u_j = 2\ell_j$, $|S^* \cap S_k| \leqslant 2|B_{j,k}|$, 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

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

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 $\text{sw}(S_0, \vec{v}) \geqslant \text{sw}(S^* \cap S_0, \vec{v})$.

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 $\text{sw}(S^*, \vec{v})/(2+16\log^2 m)$. 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 3.6. *The distortion associated with randomized threshold approval votes is* $\Omega(\log m/\log\log m)$.

Proof. 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 $\lceil \log m/\log(2\log m)\rceil$ sub-intervals: 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 ℓ_j denote the upper and lower end points of I_j and observe that $u_j \leqslant 2\log m \cdot \ell_j$ 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 $\Pr[t \in I_k] \leqslant \log(2\log m)/\log m$. 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

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

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 $\Pr[f \text{ returns } a^* | t \notin I_k] \leqslant 1/\log m$. 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 $\text{sw}(a^*, \vec{v}) = n \cdot u_k/2$, whereas $\text{sw}(a, \vec{v}) \leqslant n \cdot \ell_k$ 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

$$\Pr[t \in I_k] \cdot \frac{n \cdot u_k}{2} + \Pr[t \notin I_k] \left(\frac{1}{\log m} \cdot \frac{n \cdot u_k}{2} + \frac{\log m - 1}{\log m} \cdot n \cdot \ell_k \right) \leqslant \frac{\log(2\log m) + 2}{\log m} \cdot \frac{n \cdot u_k}{2},$$

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

Our proof of Theorem 3.6 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 $a \in A$, i.e., exactly one alternative needs to be selected. This is exactly the setting studied by Boutilier et al. [2015]. On the other hand, Theorem 3.5 establishes a slightly weaker upper bound of $\mathcal{O}(\log^2 m)$ for the general participatory budgeting problem. For the restricted setting of Boutlier 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 3.5, 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.

Theorem 3.7. *If $c_a = 1$ for all $a \in A$, the distortion associated with randomized threshold approval votes is $\mathcal{O}(\log m)$.*

Proof. This proof is along the lines of the more general proof of Theorem 3.5, 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. In our special case, with the alternatives having an equal cost, there is no longer a need to partition them based on their cost, which leads to an improvement in the bound by a factor of $\log m$.

Formally, for $j \in [\log m]$, let $\ell_j = 2^{j-1}/m$ and $u_j = 2 \cdot \ell_j$. Consider the rule which chooses $j \in [\log m]$ uniformly at random, elicits approval votes with threshold $t = \ell_j$, and returns an alternative with the greatest number of approval votes. We show that the distortion of this rule is $\mathcal{O}(\log m)$.

Let $\vec{v}$ denote the underlying utility profile, and $a^* = \arg\max_{a \in A} \operatorname{sw}(a, \vec{v})$ be the welfare-maximizing alternative. If there exists $j \in [\log m]$ such that our rule returns a^* when it sets the threshold $t = \ell_j$ (which happens with probability $1/\log m$), we immediately obtain $\mathcal{O}(\log m)$ distortion. Let us assume that our rule never returns a^*. For $a \in A$ and $j \in [\log m]$, let n_j^a denote the number of approval votes a receives when the threshold $t = \ell_j$, and let $a_j \in A$ be the alternative returned by our rule when $t = \ell_j$. Because our rule returns an alternative with the greatest number of approval votes, we have

$$\forall j \in [\log m], \; \sum_{k=j}^{\log m} n_k^{a_j} \geqslant \sum_{k=j}^{\log m} n_k^{a^*} \geqslant n_j^{a^*}. \tag{9}$$

Now, the expected social welfare achieved by our rule is at least

$$\sum_{j=1}^{\log m} \Pr[t = \ell_j] \cdot \operatorname{sw}(a_j, \vec{v}) \geqslant \frac{1}{\log m} \sum_{j=1}^{\log m} \ell_j \left(\sum_{k=j}^{\log m} n_k^{a_j} \right) \geqslant \frac{1}{2 \log m} \sum_{j=1}^{\log m} u_j \cdot n_j^{a^*} \geqslant \frac{1}{2 \log m} \cdot \operatorname{sw}(a^*, \vec{v}),$$

where the first transition follows from Equation (9), and the second transition holds because $\ell_j = u_j/2$. Hence, the distortion of our rule is $\mathcal{O}(\log m)$, as desired. $\square$

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. While our construction requires the number of voters to be extremely large compared to 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 3.8. *For sufficiently large n, the distortion associated with deterministic aggregation of knapsack votes is $\Omega(2^m/\sqrt{m})$.*

Proof. 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 $n \geqslant \binom{m}{\lfloor m/2 \rfloor}$ 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 $\lfloor n/\binom{m}{\lfloor m/2 \rfloor} \rfloor \leqslant n_i \leqslant \lceil n/\binom{m}{\lfloor m/2 \rfloor} \rceil$, 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 $\text{sw}(a^*, \vec{v}) \geqslant n \cdot \frac{\lfloor m/2 \rfloor}{m} \geqslant n/3$ (for $m \geqslant 2$).

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 $\text{sw}(f(\vec{\kappa}), \vec{v}) \leqslant n_{k^*} \leqslant n/\binom{m}{\lfloor m/2 \rfloor} + 1$. By assumption, $n \geqslant \binom{m}{\lfloor m/2 \rfloor}$, so we have

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

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

Proof. 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 $v_i(S^*) \geqslant 1/m$ because there must exist $a \in A$ such that $v_i(a) \geqslant 1/m$, and $v_i(S^*) \geqslant v_i(a)$.

Hence, we have $\text{sw}(S^*, \vec{v}) \geqslant (n/2^m) \cdot 1/m$, 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 3.10. *The distortion associated with deterministic aggregation of rankings by value for money is unbounded.*

Proof. 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 $\text{dist}(f) \geqslant \text{dist}(f, \vec{v}) = \infty$.

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

Hence, in either case, $\text{dist}(f) \geqslant (1 - \epsilon)/\epsilon$. Because ϵ 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 3.11 (Caragiannis *et al.* [2016])**.** *For $n \geqslant m - 1$, 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 3.12. *The distortion associated with deterministic aggregation of rankings by value is $\mathcal{O}(m^2)$.*

Proof. 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 3.13. *For a fixed threshold $t > 1/(m - 1)$, the distortion associated with deterministic aggregation of deterministic threshold approval votes is unbounded.*

Proof. 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 3.14. *For the fixed threshold $t = 1/m$, the distortion associated with deterministic aggregation of deterministic threshold approval votes is $\mathcal{O}(m^2)$.*

Proof. 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 $\vec{\tau}$. 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

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

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

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

Using a similar line of argument, we also have

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

Hence, the distortion of f is at most

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

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

□

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.

While these profile-wise optimal aggregation rules dominate all other aggregation rules, they may be computationally difficult to implement, because they optimize a non-linear 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 two-stage algorithm in the spirit of the cutting-set approach of Mutapcic and Boyd [2009]. More details can be found in Appendix A.

We evaluate the practicality of this approach by comparing the running times of computing deterministic voting rules, averaged over 10 trials, on data from participatory budgeting elections held in Boston in 2016. Voters were asked to choose from 10 alternatives; $4{,}430$ votes were cast.

Our discussions with officials from several cities have revealed a hesitance to use randomized voting rules, so we are particularly interested in the performance of the determinstic rules. (Our computational results in the following section also show that deterministic worst case rules typically perform better in terms of average welfare ratio than randomized rules.)

Figure 1 summarizes the average time to compute the deterministic worst-case optimal set of alternatives on a log-log scale. The experiments were run on an 8-core Intel(R) Xeon(R) CPU with 2.27GHz processor speed and 50GB memory. We observe that the running time scales gracefully with the number of agents. When sampling 500 voters, rules such as Det Th Ap and Det Val take less than 5 minutes, indicating the practicability of these methods for the participatory budgeting elections at the scale of those in Boston, MA [Goel *et al.*, 2016]. We also note that, due to the once-off nature of participatory budgeting elections, it is conceivable to use an aggregation algorithm which takes several days or even weeks to compute the optimal set of alternatives.

5 Empirical Results

Our theoretical results in §3 characterize how well we can optimize distortion on an observed input profile. Recall that distortion is the *worst-case* ratio of the optimal social welfare to the social

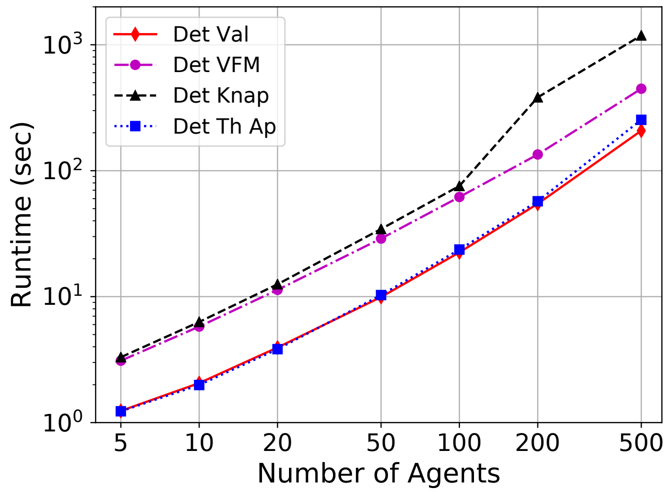

Figure 1: Average running time of the deterministic voting rules on the Boston 2016 dataset.

welfare achieved, where the worst case is taken over all utility profiles consistent with the observed input profile. In practice we care about this ratio according to the *actual* underlying utility profile. Thus, a distortion-minimizing aggregation rule is not guaranteed to be optimal in practice. This is why an empirical study is called for.

In this section, we compare the performance of different approaches to participatory budgeting, where the performance is measured by the *average-case* ratio of the optimal and achieved social welfare, and the average is taken over utility profiles drawn to be consistent with input profiles from two real-world participatory budgeting elections.

Datasets: We use data from participatory budgeting elections held in 2015 and 2016 in Boston, Massachusetts. Both elections offered voters 10 alternatives. The 2015 dataset contains 2600 4-approval votes (voters were asked to approve their four most preferred alternatives) and the 2016 dataset contains 4430 knapsack votes.

For each dataset, we conduct three independent trials. In each trial, we create r sub-profiles, each consisting of n voters drawn at random from the population. For each sub-profile, we draw k random utility profiles $\vec{v}$ consistent with the sub-profile, 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.

Approaches: We use the utility profile $\vec{v}$ drawn to create an input profile in four input formats we study. For each format, we use the deterministic as well as randomized distortion-minimizing aggregation rule. The non-trivial algorithms we devise for these rules are presented in §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. While randomizing over the threshold is required to

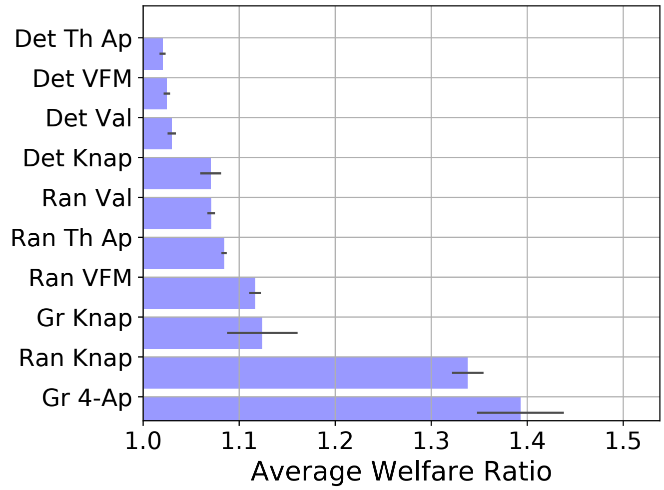

Figure 2: Average welfare ratio of different approaches to participatory budgeting based on data from Boston 2015 and 2016 elections (lower is better).

minimize the distortion (the *worst-case* ratio of the optimal and achieved social welfare), as is our goal in the theoretical results of §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 hold-out 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.*, 2016]: 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.

As the performance measure for the ten approaches, we use the average ratio of the optimal and the achieved social welfare according to the actual utility profile used to induce the input profiles — termed *average welfare ratio* — where the average is taken across the entire experiment.

Results: Figure 2 shows the average 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. While randomized rules do achieve better *distortion*, there always exists a deterministic rule minimizing the *average welfare ratio* objective; although, it is not necessarily the deterministic distortion-minimizing aggregation rule.

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 distortion 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.

6 Discussion

Our results indicate that threshold approval votes should receive serious consideration as the input format of choice for participatory budgeting. But there is one important issue we have not studied: the cognitive load imposed on voters by different input formats. (If it were not for this issue, we would just elicit the full utility functions — the whole point is to reduce cognitive load.) While determining threshold approval votes, as discussed in §2, seems reasonable enough (and probably easier than casting knapsack votes), human subject experiments are needed to rigorously determine whether threshold approval, and other input formats, require an acceptable cognitive effort.

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

This work was partially supported by the NSF under grants CCF-1215883, CCF-1525932, and IIS-1350598, by the ONR grant N00014-16-1-3075, by a Sloan Research Fellowship, and by a Fulbright-Nehru Post-doctoral Fellowship and by the NSERC under the Discovery Grants program. We thank Ashish Goel and Anilesh Krishnaswamy of the Stanford Crowdsourced Democracy Team for generously sharing their data with us, and Tuomas Sandholm for providing the hardware to run the experiments.

References

E. Anshelevich and J. Postl. Randomized social choice functions under metric preferences. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 46–52, 2016.

E. Anshelevich and S. Sekar. Blind, greedy, and random: Algorithms for matching and clustering using only ordinal information. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*, pages 390–396, 2016.

E. Anshelevich, O. Bhardwaj, and J. Postl. Approximating optimal social choice under metric preferences. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, pages 777–783, 2015.

K. Arrow. *Social Choice and Individual Values*. Wiley, 1951.

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

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

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

S. Bouveret and J. Lang. A general elicitation-free protocol for allocating indivisible goods. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 73–78, 2011.

S. J. Brams and P. C. Fishburn. *Approval Voting*. Springer, 2nd edition, 2007.

F. Brandt, V. Conitzer, U. Endress, J. Lang, and A. D. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.

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

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

I. Caragiannis, S. Nath, A. D. Procaccia, and N. Shah. Subset selection via implicit utilitarian voting. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 151–157, 2016.

D. Chakrabarty and C. Swamy. Welfare maximization and truthfulness in mechanism design with ordinal preferences. In *Proceedings of the 5th Conference on Innovations in Theoretical Computer Science (ITCS)*, pages 105–120, 2014.

A. Charnes and W. W. Cooper. Programming with linear fractional functionals. *Naval Research Logistics Quarterly*, 9(3–4):181–186, 1962.

V. Conitzer and T. Sandholm. Common voting rules as maximum likelihood estimators. In *Proceedings of the 21st Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 145–152, 2005.

V. Conitzer, R. Freeman, and N. Shah. Fair public decision making. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, pages 629–646, 2017.

G. B. Dantzig. Discrete-variable extremum problems. *Operations Research*, 5(2):266–288, 1957.

Brandon Fain, Kamesh Munagala, and Nisarg Shah. Fair allocation of indivisible public goods. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, 2018. Forthcoming.

A. Filos-Ratsikas, S. K. S. Frederiksen, and J. Zhang. Social welfare in one-sided matchings: Random priority and beyond. In *Proceedings of the 7th International Symposium on Algorithmic Game Theory (SAGT)*, pages 1–12, 2014.

A. Goel, A. K. Krishnaswamy, S. Sakshuwong, and T. Aitamurto. Knapsack voting for participatory budgeting. Manuscript, 2016.

B. Gutiérrez. Madrid as a democracy lab, 2017. Accessed: 2018-05-02.

Valen E Johnson. Revised standards for statistical evidence. *Proceedings of the National Academy of Sciences*, 110(48):19313–19317, 2013.

P. Krysta, D. Manlove, B. Rastegari, and J. Zhang. Size versus truthfulness in the house allocation problem. In *Proceedings of the 15th ACM Conference on Economics and Computation (EC)*, pages 453–470, 2014.

J. Legendre, H Madénian, and P. L. Scully. Participatory budgeting in Paris, France, 2017. Accessed: 2018-05-02.

T. Lu and C. Boutilier. Budgeted social choice: From consensus to personalized decision making. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 280–286, 2011.

H. Moulin. *Fair Division and Collective Welfare*. MIT Press, 2003.

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

A. D. Procaccia and J. S. Rosenschein. The distortion of cardinal preferences in voting. In *Proceedings of the 10th International Workshop on Cooperative Information Agents (CIA)*, pages 317–331, 2006.

A. D. Procaccia, S. J. Reddi, and N. Shah. A maximum likelihood approach for selecting sets of alternatives. In *Proceedings of the 28th Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 695–704, 2012.

A. Shah. *Participatory Budgeting*. The World Bank, 2007.

P. Skowron, P. Faliszewski, and J. Lang. Finding a collective set of items: From proportional multirepresentation to group recommendation. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, pages 2131–2137, 2015.

H. P. Young. Condorcet's theory of voting. *The American Political Science Review*, 82(4):1231–1244, 1988.

Appendix

A Algorithmic Details Omitted From §4

Throughout this section, we assume that it is practically feasible to explicitly enumerate the collection of inclusion-maximal feasible sets of alternatives $\mathcal{F}_c$. This assumption is justified given that real-world participatory budgeting problems typically involve up to 20 alternatives [Goel *et al.*, 2016].

A.1 Deterministic Rules

Let $V(\vec{\rho}) = \{\vec{v} : \vec{v} \triangleright \vec{\rho}\}$ denote the set of utility profiles consistent with input profile $\vec{\rho}$. Hence, we are interested in computing

$$\arg\min_{S \in \mathcal{F}_c} \max_{\vec{v} \in V(\vec{\rho})} \frac{\max_{T \in \mathcal{F}_c} \mathrm{sw}(T, \vec{v})}{\mathrm{sw}(S, \vec{v})} = \arg\min_{S \in \mathcal{F}_c} \max_{T \in \mathcal{F}_c} \left\{ \max_{\vec{v} \in V(\vec{\rho})} \frac{\mathrm{sw}(T, \vec{v})}{\mathrm{sw}(S, \vec{v})} \right\}.$$

An algorithm is self-evident: compute $d(\vec{\rho}, S, T) = \max_{\vec{v} \in V(\vec{\rho})} \mathrm{sw}(T, \vec{v}) / \mathrm{sw}(S, \vec{v})$ for every pair $S, T \in \mathcal{F}_c$, and return $\arg\min_{S \in \mathcal{F}_c} \max_{T \in \mathcal{F}_c} d(\vec{\rho}, S, T)$.

For the input methods we study in this paper, we can describe the set of consistent utility profiles $V(\vec{\rho})$ using linear constraints. Observe that $V(\vec{\rho}) = V(\rho_1) \times \cdots \times V(\rho_n)$, where $V(\rho_i) = \{v \geqslant 0 : v \triangleright \rho_i\}$ is the set of m-dimensional utility functions consistent with voter i's input ρ_i. It is therefore sufficient to describe each $V(\rho_i)$ using linear constraints.

For a ranking by value σ_i, we use:

$$V(\sigma_i) = \left\{ v_i \in \mathbb{R}_+^m : \sum_{a \in A} v_i(a) = 1 \wedge \left(v_i(\sigma_i^{-1}(k)) \geqslant v_i(\sigma_i^{-1}(k+1)),\ \forall k \in [m-1] \right) \right\}.$$

For a ranking by value for money σ_i, we use:

$$V(\sigma_i) = \left\{ v_i \in \mathbb{R}_+^m : \sum_{a \in A} v_i(a) = 1 \wedge \left(\frac{v_i(\sigma_i^{-1}(k))}{c_{\sigma_i^{-1}(k)}} \geqslant \frac{v_i(\sigma_i^{-1}(k+1))}{c_{\sigma_i^{-1}(k+1)}}, \forall k \in [m-1] \right) \right\}.$$

For a knapsack vote κ_i, we use:

$$V(\kappa_i) = \left\{ v_i \in \mathbb{R}_+^m : \sum_{a \in A} v_i(a) = 1 \wedge \left(\sum_{a \in \kappa_i} v_i(a) \geqslant \sum_{a \in S} v_i(a),\ \forall S \in \mathcal{F}_c \right) \right\}.$$

For a threshold approval vote τ_i elicited using threshold t, we use:

$$V(\tau_i) = \left\{ v_i \in \mathbb{R}_+^m : \sum_{a \in A} v_i(a) = 1 \wedge \left(v_i(a) \geqslant t,\ \forall a \in \tau_i \right) \wedge \left(v_i(a) \leqslant t,\ \forall a \in A \setminus \tau_i \right) \right\}.$$

Note that the polytope for knapsack votes has exponentially many constraints, while the other polytopes have polynomially many constraints. When necessary, heuristics may be devised to approximate $V(\kappa_i)$, however, in our experiments with real data we only encountered instances with $m \leqslant 20$, where it was possible to enumerate the constraints in $V(\kappa_i)$.

This polytope $V(\rho_i)$ is the only part of our generic algorithm that is dependent on the input format. Generically, let $A(\vec{\rho})\,\vec{v} \leqslant b(\vec{\rho})$ be the set of linear constraints describing $V(\vec{\rho})$.

Our next goal is to use this characterization of $V(\vec{\rho})$ to compute $d(\vec{\rho}, S, T)$ for specific $S, T \in \mathcal{F}_c$. Recall that $\text{sw}(S, \vec{v}) = \sum_{a \in A} x^S(a) \sum_{i \in [n]} v_i(a)$, where x^S is the characteristic vector for the set of alternatives S, and that $d(\vec{\rho}, S, T) = \max\{\frac{\text{sw}(T,\vec{v})}{\text{sw}(S,\vec{v})} : A(\vec{\rho})\,\vec{v} \leqslant b(\vec{\rho})\}$. This is a standard *linear-fractional program*, which can be converted to a linear program $LP(\vec{\rho}, S, T)$ using the famous Charnes-Cooper transformation [Charnes and Cooper, 1962]. We omit these standard transformations in the interest of space.

The complete algorithm for resolving the deterministic optimal aggregation rule on an input profile $\vec{\rho}$ is given as Algorithm 1.

Algorithm 1: Computing the worst-case optimal deterministic rule

---

Data: Input profile $\vec{\rho}$
Result: A set $S \in \mathcal{F}_c$ yielding the least distortion
dist[S] = 0, $\forall S \in \mathcal{F}_c$
for $S \in \mathcal{F}_c$ **do**
  for $T \in \mathcal{F}_c$, $T \neq S$ **do**
    dist[S] = $\max($dist[S], $LP(\vec{\rho}, S, T))$
  end
end
return $\arg\min_{S \in \mathcal{F}_c}$ dist[S]

A.2 Randomized Rules

Using a similar line of argument as before, observe that the optimal randomized aggregation rule returns the following distribution p over feasible sets of alternatives:

$$\arg\min_{p \in \Delta(\mathcal{F}_c)} \max_{T \in \mathcal{F}_c} \max_{\vec{v} \in V(\vec{\rho})} \frac{\text{sw}(T, \vec{v})}{\sum_{S \in \mathcal{F}_c} p(S) \cdot \text{sw}(S, \vec{v})}.$$

We introduce an additional continuous variable z representing the optimal distortion, and reformulate the problem as

$$\begin{aligned} \min_{p,z} \quad & z \\ s.t. \quad & \text{sw}(T, \vec{v}) - z \cdot \sum_{S \in \mathcal{F}_c} p(S) \cdot \text{sw}(S, \vec{v}) \leqslant 0, \ \forall T \in \mathcal{F}_c, \vec{v} \in V(\vec{\rho}) \\ & p \in \Delta(\mathcal{F}_c). \end{aligned} \tag{10}$$

At this point, it is possible to handle the constraints in (10) by formulating the problem in terms of the vertices of the polytope $V(\vec{\rho})$. Instead, we couple binary search with a constraint-generation approach.

Our algorithm performs a binary search on z, the optimal distortion. For a fixed value of z, say $\tilde{z}$, an iterative two-stage procedure determines whether there exists a distribution p whose distortion on the input profile $\vec{\rho}$ is at most $\tilde{z}$. If such a distribution p exists, then $\tilde{z}$ serves as an upper bound on the smallest distortion; otherwise, it serves as a lower bound. After adjusting the bounds on the optimal distortion, the value of $\tilde{z}$ is updated as in traditional binary search.

We now describe the iterative two-stage procedure that ascertains the existence of a distribution p with distortion at most $\tilde{z}$. This procedure alternately finds a distribution satisfying a limited subset of the constraints in (10), then attempts to add omitted constraints from (10) which are violated by the current distribution. At iteration t, a set of constraints defined by $\mathcal{C}_{t-1}$ have been added and we check the feasibility of

$$\left.\begin{array}{l} \mathrm{sw}(T,\vec{v}) - \tilde{z}\cdot\sum_{S\in\mathcal{F}_c} p^t(S)\cdot \mathrm{sw}(S,\vec{v}) \leqslant 0,\ \forall(\vec{v},T)\in\mathcal{C}_{t-1} \\ p^t \in \Delta(\mathcal{F}_c) \end{array}\right\} \mathrm{CF}(\tilde{z},\mathcal{C}_{t-1}).$$

If no feasible distribution p^t exists, $\tilde{z}$ is the new lower bound on the optimal distortion, and we proceed to the next step in our binary search over z. If a feasible p^t exists, we check whether it violates any constraint from (10) by solving the following linear program (which serves as an oracle) for every $T \in \mathcal{F}_c$:

$$\left.\begin{array}{ll} \max & \mathrm{sw}(T,\vec{v}) - z\cdot\displaystyle\sum_{S\in\mathcal{F}_c} p^t(S)\cdot \mathrm{sw}(S,\vec{v}) \\ & s.t.\ \ \vec{v}\in V(\vec{\rho}) \end{array}\right\} \mathrm{LP}(T,z,p^t,\vec{\rho}).$$

If the objective value of LP$(T, z, p^t, \vec{\rho})$ exceeds 0, a violated constraint is found and $(\vec{v}^*, T)$ is added to $\mathcal{C}_{t-1}$ to form $\mathcal{C}_t$, where $\vec{v}^*$ is the optimal solution to LP$(T, z, p^t, \vec{\rho})$. The algorithm then returns to solving CF$(\tilde{z}, \mathcal{C}_t)$. If no violated constraints are found, the current distribution p^t indeed has distortion at most $\tilde{z}$, and establishes an upper bound on the optimal distortion.

This complete procedure is summarized in Algorithm 2. A finite number of violated constraints can be added for each $\tilde{z}$, so we may conclude that Algorithm 2 will terminate.

B Empirical Results Omitted From §5

We provide a more detailed representation of the results summarized in §5, and investigate the usefulness of learning the optimal threshold for threshold approval voting from holdout data.

B.1 Additional Empirical Results

Figure 2 in §5 presented the average welfare ratio of ten different approaches to participatory budgeting, where the average was taken over repeated independent trials in each of two datasets. The results for the Boston 2015 and 2016 datasets are presented separately in Figure 3.

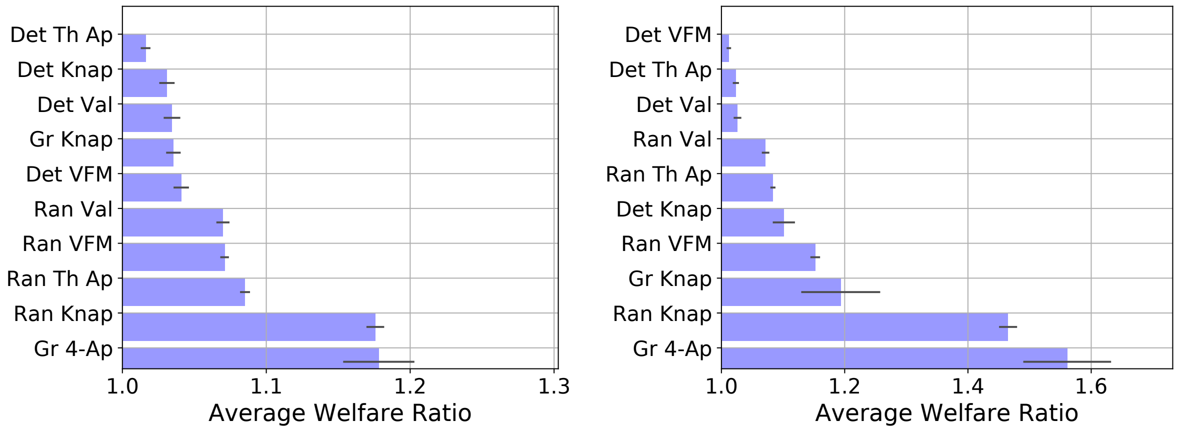

Figure 3: Average welfare ratios on the Boston datasets from 2015 (containing 4-approval votes, left) and 2016 (containing knapsack votes, right).

Many of the trends highlighted in §5 are reflected across both datasets. First, approaches based on deterministic distortion-minimizing aggregation rules, excluding the one using knapsack votes, still outperform their randomized counterparts. Further, among these approaches, the one using threshold approval votes is the most consistent, achieving the lowest average welfare ratio for the Boston 2015 dataset and the second lowest for the Boston 2016 dataset. Second, the approaches currently used in real-world elections ("Gr Knap" and "Gr 4-Ap") perform worse than most other approaches, and have high variance in their performance.

The main difference between the datasets is that, somewhat surprisingly, Greedy Knapsack performs significantly better on knapsack votes induced from random utility profiles drawn to be consistent with real 4-approval votes, than it does on real knapsack votes. This may in part be due to the fact that the families of utility profiles consistent with 4-approval (Boston 2015) and knapsack (Boston 2016) votes, are very different.

B.2 Is It Useful to Learn the Threshold?

In our experiments, when using threshold approval votes, we use the threshold that achieves the best performance on a holdout/training set, to evaluate the performance of threshold approval votes on the test set.

After sampling a random utility profile consistent with an input vote in the training set, we generate threshold approval votes for all the thresholds in $[0, 1]$ at intervals of 0.05, and compute the average distortion (per threshold) across multiple samples. (This step will not be required in practice once sufficiently many real votes are elicited).

We select the threshold value that achieves the least average distortion. Importantly, note that we use distortion — which is only a function of the input profile — rather than the average distortion to select the optimal threshold value. Hence, this method is robust, and does not use any knowledge of the distribution of utility profiles that we later use in evaluating performance.

This optimal threshold value when evaluating the performance (average distortion) of threshold approval votes, in conjunction with both the deterministic and the randomized distortion-minimizing aggregation rules.

While threshold approval votes with deterministic aggregation rule achieves excellent performance with this method of threshold selection, it is not immediately clear whether the threshold selection was useful. Indeed, learning a threshold is only useful if the optimal threshold value remains reasonably consistent across the instances. We now investigate the usefulness of threshold selection in multiple ways.

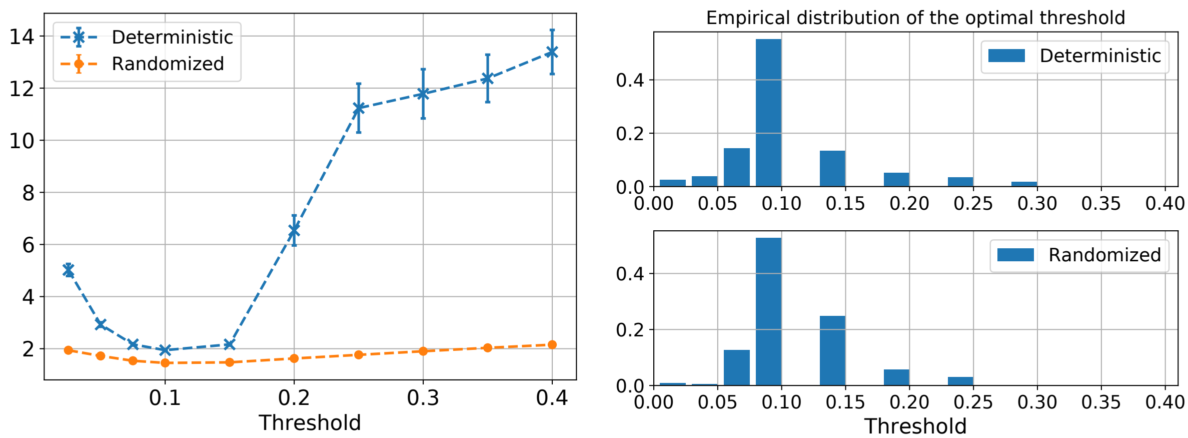

Figure 4: (Left) Average distortion achieved by different threshold values in threshold approval votes; (Right) Empirical distribution of the optimal threshold under deterministic and randomized aggregation.

First, Figure 4 shows the average distortion achieved by different values of the threshold on the training instances, when used in conjunction with the deterministic and the randomized distortion-minimizing aggregation rules. Recall that the final threshold value we select is the one that minimizes this measure. For every threshold value on the x-axis, the error bars indicate the range that contains the distortion on 95% of the training instances. We do not plot threshold values above 0.4 as the distortion is non-decreasing beyond this point.

We observe that the thresholds values that lead to the smallest average distortion are exactly those with the smallest variation across instances. Interestingly, the average distortion of different values of the threshold is wildly different under the deterministic aggregation rule, but rather similar under the randomized aggregation rule. This effect perhaps manifests itself in the improved performance of threshold approval votes with deterministic aggregation than with randomized aggregation in all of our experiments; see Figures 2 and 3.

Next, we measure the usefulness of training the threshold value in a different way. In Figure 4 (right) we plot the empirical distribution of the optimal threshold value, i.e., for each threshold value, we plot the percentage of training instances in which that value led to the smallest distortion. For both deterministic and randomized aggregation rules, the distribution of the optimal threshold value is (quite strongly) centered at 0.1. In fact, the optimal threshold value was in $[0.075, 0.15]$ in more than 80% of the training instances.

The consistency with which a single threshold value (0.1) remains the optimal value suggests that learning this value from the holdout set is very likely to be valuable.

Finally, we note that the datasets we used contain votes over 10 alternatives. That is, $m = 10$. Interestingly, this makes the empirically optimal threshold value $1/m$, which is precisely the value for which we achieve the best performance in the worst case in our theoretical results (see Theorem 3.14).

Algorithm 2: Computing the optimal randomized aggregation rule

```
Data: Input profile ρ⃗, tolerance TOL, 𝓕_c
Result: A probability distribution in Δ(𝓕_c), the optimal distortion
z⁻ = 1, z⁺ = 100, z̃ = (z⁻ + z⁺)/2
while z⁻ − z⁺ > TOL do
    𝓒_0 = ∅, t = 0
    robustFeasibleFlag ← false
    while robustFeasibleFlag is false do
        robustFeasibleFlag ← true
        t ← t + 1
        if CF(z̃, 𝓒_{t−1}) is feasible then
            p^t ← optimal solution of CF(z̃, 𝓒_{t−1})
            for T ∈ 𝓕_c do
                𝓒_t = 𝓒_{t−1}
                if optimum of LP(T, z, p^t, ρ⃗) exceeds 0 then
                    v⃗* ← optimal solution of LP(T, z, p^t, ρ⃗)
                    𝓒_t ← 𝓒_t ∪ (v⃗*, T)
                    robustFeasibleFlag ← false
                end
            end
            if robustFeasibleFlag then
                z⁺ = z̃
            end
        else
            z⁻ = z̃
        end
    end
    z̃ = (z⁺ + z⁻)/2
end
return p^t, z⁺
```