

Preference Elicitation For Participatory Budgeting

Gerdus Benade

Carnegie Mellon University
jbenade@andrew.cmu.edu

Swaprava Nath

Carnegie Mellon University
swapravn@cs.cmu.edu

Ariel D. Procaccia

Carnegie Mellon University
arielpro@cs.cmu.edu

Nisarg Shah

Harvard University
nisarg@g.harvard.edu

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

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

Our interest in budgeted social choice stems from the striking real-world impact of the *participatory budgeting* paradigm (Cabannes 2004), which allows local governments to allocate public funds by eliciting and aggregating the preferences of residents over potential projects. Indeed, in just a few years, the *Participatory Budgeting Project*¹ has helped allocate more than \$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).

In pioneering work, Goel et al. (2016) — who have facilitated a number of participatory budgeting elections as part of the Stanford Crowdsourced Democracy Team² — propose and evaluate two participatory budgeting approaches.

Copyright © 2017, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹<http://www.participatorybudgeting.org>

²<http://voxpopuli.stanford.edu>

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 the 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

Let us define the participatory budgeting problem a bit more formally, following Goel et al. (2016). A set N of n voters are voting over a set A of m alternatives (projects), where each alternative a has cost c_a . The utility voter i has for alternative a is denoted $v_i(a)$. Moreover, utility functions are additive, that is, the utility of a voter for a set of alternatives $A' \subseteq A$ is $\sum_{a \in A'} v_i(a)$. Our goal is to choose a set $W \subseteq A$ of winning alternatives that maximizes the (*utilitarian*) *social welfare*, subject to the total cost not exceeding the budget B :

$$\arg \max_{W \subseteq A : \sum_{a \in W} c_a \leq B} \sum_{i \in N} \sum_{a \in W} v_i(a). \quad (1)$$

We make essentially³ no assumptions about the utility functions. Nevertheless, solving (1) would be easy if we

³Other than a standard normalization assumption that we discuss later.

had 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, that 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, Bhardwaj, and Postl 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 thus associated with an input format measures how useful the information contained in the votes is for achieving social welfare maximization.

In §3, we apply this approach to compare four input formats. The first is knapsack votes, which (disappointingly) is associated with trivial distortion of $\Theta(m)$.⁴ 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 both admit an upper bound of $\mathcal{O}(\sqrt{m} \cdot \log m)$ on distortion, which almost matches a lower bound of $\Omega(\sqrt{m})$. 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 that its associated distortion is $\mathcal{O}(\log^2 m)$, and establish a lower bound of $\Omega(\log m / \log \log m)$. To summarize, our theoretical results show striking separations between different input formats, with threshold approval votes coming out well on top.

While our theoretical results in §3 bound the distortion, i.e., the *worst-case ratio* of the optimal social welfare to

⁴As we later show, an upper bound of $\mathcal{O}(m)$ can be achieved trivially irrespective of the input format, by selecting a single alternative uniformly at random. Knapsack votes, unfortunately, do not help improve it.

the social welfare achieved, in §4 we compare different approaches to participatory budgeting using the *average-case ratio* of the two. 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,⁵ and compare them to two approaches currently employed in practice. 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. We also observe that the running times of the distortion minimizing rules scale gracefully for most practical scenarios.

1.2 Related Work

Due to space constraints, we only discuss a few closely related papers. 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.

Naturally, our work is also closely related to previous work on implicit utilitarian voting. Crucially, as noted above, this line of work 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 bound of $\mathcal{O}(\sqrt{m} \log^* m)$ and a lower bound of $\Omega(\sqrt{m})$ 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 alterna-

⁵Note that such 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.

tives 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.

2 The Model

Let $[k] \triangleq \{1, \dots, k\}$. 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 $\mathcal{F}_c = \{S \subseteq A : c(S) \leq 1 \wedge c(T) > 1, \forall S \subsetneq T \subseteq A\}$ as the inclusion-maximal budget-feasible subsets of A .

We assume that each voter has a *utility function* $v_i : A \rightarrow \mathbb{R}_+ \cup \{0\}$, 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 $\vec{v} = \{v_1, \dots, v_n\}$ 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 $\vec{\rho} = \{\rho_1, \dots, \rho_n\}$ 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:

- 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 $c(\kappa_i) \leq 1$ and $v_i(\kappa_i) \geq v_i(S)$ for all $S \in \mathcal{F}_c$.
- 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 $v_i(a) \geq v_i(b)$ (resp. $v_i(a)/c_a \geq v_i(b)/c_b$) for all $a \succ_{\sigma_i} b$.

- 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 $\tau_i = \{a \in A : v_i(a) \geq t\}$.

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

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

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

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

where $\Delta(\mathcal{F}_c)$ is the set of distributions over \mathcal{F}_c . Needless to say, f^* achieves the best possible overall distortion.

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.⁶

3 Theoretical Results

Before we present our analysis of the different input formats from the perspective of implicit utilitarianism, let us make a simple observation that holds across all input formats.

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 . \square

3.1 Knapsack Votes

We now present our analysis for knapsack votes — an input format advocated by Goel et al. (2016).

Theorem 3.2. *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). For ease of exposition, assume that m divides n . Consider the input profile $\vec{\kappa}$, in which voters

⁶In a setting where deterministic rules must be used, one could similarly associate each input format with its best deterministic rule. We study this less motivated and technically less interesting setting in the full version of the paper available at <http://procaccia.info/papers>.

are partitioned into m subsets $\{N_a\}_{a \in A}$ of equal size, and for every $a \in A$ and $i \in N_a$, we have $\kappa_i = \{a\}$.

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

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

$$\text{dist}(f) \geq \text{dist}(f, \vec{\kappa}) \geq \frac{n/2}{\frac{1}{m} \cdot n + \frac{m-1}{m} \cdot \frac{n}{m}} \geq \frac{m}{4},$$

as desired. \square

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.

3.2 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, Bhardwaj, and Postl 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).

As noted in §1.2, 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, and the input format is rankings by value. 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 (ironically) 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. The proof of this result appears in the full version.

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

3.3 Threshold Approval Votes

Approval voting — where voters can choose to approve any subset of alternatives, and the 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 $v_i(a) \geq t$.

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. The proof of this result appears in the full version of the paper.

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

We thus turn our attention to *randomized threshold approval votes*, in which the threshold is selected in a randomized fashion. Technically, this is a distribution over input formats, one for each value of the threshold.⁷ The (overall) distortion of a rule f , which now determines both the distribution over thresholds D , and the aggregation of input profile $\vec{\rho}(\vec{v}, t)$ induced by utility profile \vec{v} and threshold t , is

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

Interested readers can refer to the full version of the paper for a discussion on how this definition relates to the definition of distortion of a rule for a fixed input format. We find

⁷While we study deterministic and randomized threshold selection, we always allow randomized aggregation rules. The full version of the paper presents an analysis of the case where the aggregation rule has to be deterministic.

that the 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 $j = 1, \dots, 2 \log m$.

Let \vec{v} denote a utility profile that is consistent with the input profile. For $a \in A$ and $j \in \{0, \dots, 2 \log m\}$, 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,

$$\begin{aligned} \text{sw}(a, \vec{v}) &= \sum_{i \in N} v_i(a) \leq \sum_{j=0}^{2 \log m} \sum_{i \in N} \mathbb{I}\{v_i(a) \in I_j\} \cdot u_j \\ &= \sum_{j=0}^{2 \log m} n_j^a \cdot u_j, \end{aligned}$$

A similar argument also yields a lower bound, and after substituting $\ell_0 = 0$, $u_0 = 1/m^2$, and $n_0^a \leq n$, we get

$$\sum_{j=1}^{2 \log m} n_j^a \cdot \ell_j \leq \text{sw}(a, \vec{v}) \leq \frac{n}{m^2} + \sum_{j=1}^{2 \log m} n_j^a \cdot u_j. \quad (2)$$

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 $S^* = \arg \max_{S \in \mathcal{F}_c} \text{sw}(S, \vec{v})$ be the feasible set of alternatives maximizing the social welfare. For $j, k \in \{0, \dots, 2 \log m\}$, let $n_{j,k}^* = \sum_{a \in S^* \cap S_k} n_j^a$. Using Equation (2), we have

$$\begin{aligned} \sum_{j=1}^{2 \log m} n_{j,k}^* \cdot \ell_j &\leq \text{sw}(S^* \cap S_k, \vec{v}) \\ &\leq |S^* \cap S_k| \cdot \frac{n}{m^2} + \sum_{j=1}^{2 \log m} n_{j,k}^* \cdot u_j. \quad (3) \end{aligned}$$

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

$$\begin{aligned} \text{sw}(B_{j,k}, \vec{v}) &\geq \sum_{a \in B_{j,k}} \left(\sum_{p=j}^{2 \log m} n_p^a \right) \cdot \ell_j \\ &\geq \frac{1}{4} \cdot \left(\sum_{p=j}^{2 \log m} n_{p,k}^* \right) \cdot u_j \geq \frac{1}{4} \cdot n_{j,k}^* \cdot u_j, \quad (4) \end{aligned}$$

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 $u_j = 2\ell_j$, $|S^* \cap S_k| \leq 2|B_{j,k}|$, 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

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

where the first transition follows from Equation (4), and the second transition follows from Equation (3).

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. Further, the social welfare achieved by this mechanism is $\text{sw}(S_0, \vec{v}) \geq \text{sw}(S^* \cap S_0, \vec{v})$.

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 $\mathcal{O}(\log^2 m)$. \square

We also show that at least logarithmic distortion is inevitable even when using randomized threshold approval votes; the proof appears in the full version of the paper.

Theorem 3.6. *The distortion associated with randomized threshold approval votes is $\Omega(\log m / \log \log m)$.*

4 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 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 4 most preferred alternatives) and the 2016 dataset contains 4430 knapsack votes.

For each dataset, we conduct 3 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 the full version of the paper. 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”).

It is important to note that, 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 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. Thus, in our experiments, we learn the optimal threshold value based on a holdout set that is not subsequently used. This learning approach is practical as it only uses observed input votes rather than underlying actual utilities. In other words, we acknowledge that this choice 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-App”), 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 1 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.




Figure 1: Average welfare ratio (lower is better) of different approaches to participatory budgeting based on data from Boston 2015 and 2016 elections.

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 in order to cast a vote.

Figure 2 shows the running times of our deterministic voting rules, averaged over 10 trials, on a log-log scale. We consider only the deterministic voting rules as they outperform their randomized counterparts in terms of the average welfare ratio. We observe that the running time scales gracefully with the number of agents. The experiments used the Boston 2016 dataset with 10 alternatives, and were run on an 8-core Intel(R) Xeon(R) CPU with 2.27GHz processor speed and 50GB main memory. Even with 500 voters, rules such as Det Th Ap and Det Val take less than 5 minutes, indicating the practicability of these methods even for the largest real-world participatory budgeting elections that we are aware of, which have no more than 5,000 voters.

5 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 is-




Figure 2: Average running time of the deterministic voting rules on the Boston 2016 dataset.

sue 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.) A participatory budgeting system based on threshold approval votes might ask voters to “mark each project on which you would be happy to see the city spend \$10,000”. While this seems reasonable enough (and probably easier than casting knapsack votes), human subject experiments are needed to rigorously determine whether threshold approval votes, 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.

Acknowledgements

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

References

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

Anshelevich, E., and Sekar, S. 2016. Blind, greedy, and random: Algorithms for matching and clustering using only

ordinal information. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*, 390–396.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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