The Thirty-Second AAAI Conference
on Artificial Intelligence (AAAI-18)

Truthful and Near-Optimal Mechanisms for Welfare Maximization in Multi-Winner Elections

Umang Bhaskar
Tata Institute of Fundamental Research,
Mumbai, India 400 005

Varsha Dani
University of New Mexico,
Albuquerque, NM 87131, USA

Abheek Ghosh
Indian Institute of Technology,
Guwahati, India 781039

Abstract

Mechanisms for aggregating the preferences of agents in elections need to balance many different considerations, including efficiency, information elicited from agents, and manipulability. We consider the utilitarian social welfare of mechanisms for preference aggregation, measured by the distortion. We show that for a particular input format called threshold approval voting, where each agent is presented with an independently chosen threshold, there is a mechanism with nearly optimal distortion when the number of voters is large. Threshold mechanisms are potentially manipulable, but place a low informational burden on voters.

We then consider truthful mechanisms. For the widely-studied class of ordinal mechanisms which elicit the rankings of candidates from each agent, we show that truthfulness essentially imposes no additional loss of welfare. We give truthful mechanisms with distortion $O(\sqrt{m \log m})$ for k-winner elections, and distortion $O(\sqrt{m} \log m)$ when candidates have arbitrary costs, in elections with m candidates. These nearly match known lower bounds for ordinal mechanisms that ignore the strategic behavior. We further tighten these lower bounds and show that for truthful mechanisms our first upper bound is tight. Lastly, when agents decide between two candidates, we give tight bounds on the distortion for truthful mechanisms.

1 Introduction

How should a group of agents, presented with a set of candidates, collectively decide which candidates to select? This problem of deciding how to aggregate the preferences of multiple rational agents is the fundamental challenge in social choice theory and has implications for diverse fields, including government formation, recommendation systems, and resource allocation. In many of these applications multiple candidates can be selected. Each candidate carries a cost, and a budget typically constrains the total cost of the selected set of candidates. For example, in the selection of k members for a committee, each candidate has cost $1/k$ while the budget available is 1.

A particularly appealing recent application of preference aggregation is participatory budgeting, which enables individuals who are directly affected to decide how the budget available to their local government should be spent. Here, voters are presented with a number of candidate projects with costs, and must select a budget-feasible set of projects to fund. In the US, funds worth more than \$250 million have been allocated via participatory budgeting in more than 440 community projects. Cities as diverse as Porto Alegre in Brazil to Chicago in the US use participatory budgeting to fund projects. Researchers in computational social choice have contributed significantly to this effort, both in the design and theoretical analysis of mechanisms, as well as building systems for participatory budgeting, e.g., (Goel et al. 2016; Caragiannis et al. 2017).

The classical approach for designing mechanisms for preference aggregation is axiomatic: we identify properties that are intuitively appealing and design mechanisms that satisfy these properties. Typical properties of interest include Pareto-optimality, truthfulness, and monotonicity.[1] A natural and common assumption is that the the agents possess cardinal utilities for the candidates, not just ordinal preferences over them. This allows, e.g., the characterization of truthful randomized mechanisms (Gibbard 1977). When the utilities of different agents have a common measure and are comparable, the aggregate utility, or utilitarian welfare, is a commonly used measure to design and evaluate mechanisms. This is a standard assumption, e.g., in mechanism design, agents' preferences are expressed as money. In transportation systems, the time spent in transit is often used as a measure of utility, and the aggregate time spent as a measure of efficiency (e.g., (Mclean 2016)). Good utilitarian welfare does not substitute for other properties, but a mechanism with bad utilitarian welfare arguably has little use for most practical applications.

Our work focuses on the utilitarian welfare of mechanisms. We assume that candidates have cardinal utilities for the candidates, and in addition, for each agent, the sum of utilities for the candidates is unity. This is known as the *unit-sum* representation, and is based on the premise that all agents have equal weight. In contrast, the *unit-range* representation requires the sum of candidate utilities to be ar-

[1] Informally, a mechanism is Pareto-optimal if no other outcome increases the welfare of all the agents, truthful if every agent maximizes its utility by reporting preferences truthfully, and monotone if increasing the preference by an agent for a candidate does not decrease the probability that the candidate is selected.

bitrary, but restricts the utility of a candidate for an agent to be in the range $[0, 1]$ (e.g., (Filos-Ratsikas and Miltersen 2014)). The utility of an agent for a set of candidates is the sum of utilities of the candidates. We use the concept of distortion, defined by Procaccia and Rosenschein ((2006)) informally as the ratio of the maximum utility obtained by a budget-feasible set, to the welfare obtained by the mechanism, in the worst case over all possible inputs. The distortion of any mechanism is at least 1, and the closer it is to 1 the better. The definition naturally extends to randomized mechanisms by considering the expected welfare obtained by the mechanism. Distortion is a particularly appealing measure since it is a worst-case bound, similar to the approximation ratio studied in theoretical computer science as a measure of the efficiency of algorithms. While the approximation ratio measures the loss in efficiency due to computational complexity, the distortion measures the loss due to other constraints such as truthfulness, incomplete data obtained from the agents, or computational complexity.

In fact, if our sole objective is to minimize the distortion, and if we can elicit the cardinal utility each agent obtains from the candidates, then the objective is easily met (ignoring issues of computational complexity). We simply solve the optimization problem that selects the budget-feasible subset of candidates that maximizes utilitarian welfare, which corresponds to the well-studied knapsack problem. The problem is NP-hard even for the case of a single agent, but known algorithms called *fully polynomial-time approximation schemes* can obtain an arbitrarily good approximation in polynomial time.

However, this approach is problematic for several reasons. Firstly, agents in many settings are strategic and may misreport their utilities to the mechanism if doing so would increase their actual utility. Strategic voting in elections is a significant problem, when candidates supporting a less-popular third candidate may vote for one of the other candidates, to prevent their least desirable candidate getting elected. Secondly, even assuming that agents are truthful, the elicitation of cardinal utilities is a complex task: the human agent needs to be explained the scale being used, and must convert the implicit utilities for all the candidates to explicit values on this scale. The problem of utility elicitation is itself an active area of research (e.g., see (Chajewska, Koller, and Parr 2000; Wakker and Deneffe 1996).

Thus, our objective in this paper is to design mechanisms for preference aggregation, and particularly for participatory budgeting, that maximize utilitarian welfare in the presence of these constraints — truthfulness and the complexity of information elicited from the agents. Most of our results will hold for participatory budgeting, and we note that k-winner selection is a particular case of participatory budgeting in our setting, where the objective is to maximize utilitarian welfare. For k-winner selection, the cost of each candidate is $1/k$ and the total budget available is 1. In this case, we will sometimes show stronger results than for participatory budgeting.

1.1 Our Contribution

Our first result is a randomized mechanism for participatory budgeting that obtains distortion close to 1 when the number of agents is large. This mechanism elicits from each agent the subset of candidates with utility above a given threshold. This particular format for preference elicitation (or input format) is studied previously (Benade et al. 2017) with upper and lower bounds of $O(\log^2 m)$ and $\Omega(\log m/\log\log m)$ shown on the distortion with m candidates. We show that a subtle modification to the mechanism — when the threshold for the agents is i.i.d., rather than identical — allows us to beat the previous lower bound and give a near-optimal mechanism.

A modification of the mechanism elicits merely a binary input from each agent. We present each voter with a threshold and a single candidate and ask if the agent's utility for the candidate is above the threshold. The mechanism, even with this severely limited information, obtains distortion close to 1, when the number of agents is large.

Threshold mechanisms, despite our near-optimal results, suffer from two weaknesses: they depend on an explicit expression of cardinal utilities, and they are not truthful. We therefore next consider the extensively-studied class of ordinal mechanisms, where agents order the candidates according to their preference.[2] For these mechanisms, we consider the effect of imposing truthfulness on the efficiency, measured as in previous work by the distortion.

For ordinal mechanisms, we show that requiring truthfulness imposes essentially no loss of the distortion. Boutilier et al. (2015) showed that any ordinal mechanism has distortion $\Omega(\sqrt{m})$, even when a single candidate is to be selected and ignoring strategic behavior. For the k-winner selection problem, where k candidates are to be selected, we show a truthful ordinal mechanism with distortion $O(\sqrt{m \log m})$. This is quite surprising given the strong characterizations of truthful mechanisms (Gibbard 1977). We extend this to a truthful ordinal $O(\sqrt{m}\log m)$ mechanism for the participatory budgeting problem, where candidates have costs and a budget-feasible subset is desired. Thus, when measuring worst-case efficiency, truthfulness comes nearly for free.

It is instructive to compare our upper bound of $O(\sqrt{m \log m})$ for unit-sum representation of utilities with the lower bound of $\Omega(m^{2/3})$ of Miltersen and Filos-Ratsikas (Filos-Ratsikas and Miltersen 2014), obtained for unit-range representation of utilities. The worse bound for the unit-range case is intuitively because while agents may have significantly different importance, due to different total utilities for the candidates, the distortion of any truthful mechanism is only as good as the best *anonymous* truthful mechanism, which must treat all agents equally. Thus, truthful mechanisms cannot be biased towards agents of larger total utility, which results in worse distortion in the unit-range representation.

For truthful mechanisms, we prove that our bound of $O(\sqrt{m \log m})$ is tight, even for the 1-winner selection problem. Based on previous characterizations of truthful ordinal mechanisms, we give a series of instances which together

[2] E.g., in order of decreasing implicit cardinal utility.

show that only mechanisms based on the harmonic scoring functions of (Boutilier et al. 2015) can obtain $O(\sqrt{m \log m})$ distortion. Finally, we show that even such mechanisms have lower bound $\Omega(\sqrt{m \log m})$ on the distortion.

Lastly, we consider mechanisms with two candidates. Any truthful mechanism for this case must be ordinal. We give an ordinal mechanism that is *input-optimal*: on every input, any truthful mechanism has distortion at least as large as our mechanism. In the worst case, the distortion is 1.5.

All missing proofs appear in the full version of the paper.

1.2 Related Work

While we consider mechanisms that maximize the utilitarian welfare, this is one of many objectives studied for multi-winner elections; see, e.g., (Elkind et al. 2017; Skowron, Faliszewski, and Lang 2016). Distortion as a measure of loss of welfare due to the input format is introduced by Procaccia and Rosenschein (2006), who show that no mechanism has unit distortion, even in simple instances, and for many popular ordinal mechanisms such as Borda and Veto, the distortion is unbounded. Boutilier et al. (2015) consider randomized ordinal mechanisms for single-winner selection, and show a lower bound of $\Omega(\sqrt{m})$ and an upper bound of $O(\sqrt{m} \log^* m)$ on the distortion, ignoring strategic behaviour. They give a randomized mechanism which uses a harmonic scoring function with distortion $O(\sqrt{m \log m})$. We use this scoring function in our work as well. Further, they show that the ordinal mechanism that obtains the least distortion in any given instance can be computed in polynomial time. For participatory budgeting, when candidates have costs and a budget-feasible subset of items is to be selected, Benade et al. (2017) show that for ordinal mechanisms the distortion is bounded from above by $O(\sqrt{m} \log m)$, and below by $\Omega(\sqrt{m})$. They also introduce a new input format called threshold approval voting, where a real-valued threshold is fixed, and each agent reports the candidates with utility above the threshold. For this input format, the distortion is shown to be bounded by $O(\log^2 m)$ and $\Omega(\log m / \log \log m)$. If a single candidate is to be selected, a mechanism with distortion $O(\log m)$ is given. When a single candidate is to be selected, and agent costs (rather than utilities) form a metric, constant upper and lower bounds on the distortion are known (Anshelevich, Bhardwaj, and Postl 2015; Anshelevich and Postl 2016).

The papers cited above ignore strategic behavior. Closely related to our work is the paper by Miltersen and Filos-Ratsikas (2014), which obtains an upper bound of $O(m^{3/4})$ and a lower bound of $\Omega(m^{2/3})$ on the distortion of truthful ordinal mechanisms with unit-range representation. As discussed, the apparent discrepancy between these and our results is because of the difference in representation of agent utilities, indicating the importance of these. The authors note that the distortion of any truthful ordinal mechanism can also be obtained by a mechanism that is anonymous and neutral. Informally, a mechanism is anonymous if it is not biased towards any agent, and neutral if it is not biased towards any candidate. More generally, truthful ordinal mechanisms are characterized by Gibbard (1973; 1977) and Satterthwaite (1975). This is further developed by Barbera (Barbera 1978). We use and describe these characterizations in Section 4.

When agent have costs that form a metric, truthful mechanisms are studied by Feldman et al. (2016). If the agents are on a line, a mechanism with a distortion of 2 is given, and this is tight.

2 Notation and Preliminaries

In the participatory budgeting problem, agents from a population N of size n are faced with the problem of selecting candidates from a set $\mathcal{C}$ of size m. We will use i, j, k for agents, and x, y, z for candidates. Each candidate x has non-negative cost c_x. There is a fixed budget $B = 1$, and the set of candidates selected must have total cost at most 1; we say that such a set of candidates is feasible. In the k-winner selection problem for $k \in \mathbb{Z}_+$, each candidate has cost $1/k$.

Each agent has a utility function $u_i : \mathcal{C} \to \mathbb{R}_+$ so that the sum of utilities $\sum_{x \in \mathcal{C}} u_i(x)$ for each agent is exactly 1. This is the standard unit-sum normalization assumption to ensure that all agents have equal influence. For technical reasons we will assume that if $x \neq y$, then $u_i(x) \neq u_i(y)$ for any agent i. Let $\vec{u} := (u_i)_{i \in N}$ be a vector of utility functions, or a utility profile, for the agents. For a candidate x, we define the utilitarian social welfare (or simply utility) $\mathrm{uw}(x)$ to be $\sum_i u_i(x)$, the sum of utilities of agents for that candidate. The utility of a set of candidates is the sum of the utility of the candidates in the set, i.e., $\mathrm{uw}(S) = \sum_{x \in S} \mathrm{uw}(x)$. Agent i's utility for a subset $S \subseteq \mathcal{C}$ is $u_i(S) := \sum_{x \in S} u_i(x)$.

We use $\lambda = ((u_i)_{i \in N}, (c_x)_{x \in \mathcal{C}})$ for an instance of participatory budgeting, and Λ as the set of all possible instances. $S^*(\lambda)$ is the feasible subset with maximum welfare, and $\mathbf{OPT}(\lambda)$ is the welfare of this set. If the instance λ is clear, we use S^* and $\mathbf{OPT}$ to simplify notation.

Input formats and distortion. Our objective is to design mechanisms, possibly randomized, to select feasible subsets of maximum welfare. Keeping in mind that human agents find it burdensome to report their utility functions accurately, we consider mechanisms with differing input formats.

- *Independent threshold approval votes:* Each agent i is given a real-valued threshold T_i and returns a subset $S_i \subseteq \mathcal{C}$.
- *Binary threshold approval votes:* Each agent i is given a real-valued threshold T_i and a candidate x_i, and returns a bit b_i.
- *Ordinal votes:* Each agent i returns a linear order $\prec_i$ of $\mathcal{C}$. Let $\vec{\prec} := (\prec_i)_{i \in N}$. We say that $\vec{\prec}$ is consistent with utility profile $\vec{u} = (u_i)_{i \in N}$, written $\vec{\prec} \cong \vec{u}$, if for each agent i, $u_i(x) > u_i(y)$ implies $x \prec_i y$. A mechanism with this input format is called an ordinal mechanism. We write $(\vec{\prec}/_i \prec'_i)$ to denote the vector where the ith component of $\vec{\prec}$ is replaced by $\prec'_i$.

We note that there is thus a difference between an instance of participatory budgeting and an input to a mechanism. While the former includes the utility function for each agent, the latter may not, depending on the input format. We always assume that the costs for candidates $(c_x)_{c\in\mathcal{C}}$ are implicit inputs to the mechanism.

Given an input format, a mechanism μ is defined as a map, possibly randomized, from possible inputs to distributions over feasible subsets of $\mathcal{C}$. We consider as our primary measure of the efficiency of a mechanism its distortion (Benade et al. 2017; Procaccia and Rosenschein 2006). To define distortion formally, consider the ordinal votes. Then for a mechanism μ and input $I = \left(\vec{\precsim}, (c_x)_{x\in\mathcal{C}}\right)$, the distortion is defined as the worst case ratio over all possible utility functions consistent with $\vec{\precsim}$ of the maximum utility of a feasible subset, to the expected utility obtained by the mechanism.

$$\operatorname{dist}(\mu, I) := \sup_{\vec{u}\cong\vec{\precsim}} \frac{\max\{\operatorname{uw}(S) : \sum_{x\in S} c_x \leq 1\}}{\mathbb{E}_{S\sim\mu(\vec{\precsim})}\operatorname{uw}(S)}\,.$$

The distortion of a mechanism is defined as the maximum distortion over all possible inputs.

$$\operatorname{dist}(\mu) := \sup_{I} \operatorname{dist}(\mu, I)\,.$$

For mechanisms where the inputs are obtained deterministically, even when the mechanism is randomized, taking the supremum over all utility functions consistent with the mechanism input is appropriate. However, in Section 3 we will consider mechanisms for the threshold approval input formats when the thresholds are chosen randomly, and hence, the inputs themselves are random variables. For these mechanisms, we show that for any instance, in expectation over the input to the mechanism, the welfare obtained is large. Here, taking the supremum over all utility functions consistent with the mechanism input would in effect remove the randomization and we would be doing a worst-case analysis — not just over inputs, but also over the random bits, which defeats the purpose of randomization.

For mechanisms where the input is itself randomized, we therefore propose and use the following simpler definition:[3]

$$\operatorname{dist}(\mu) := \sup_{\lambda\in\Lambda} \frac{\mathbf{OPT}(\lambda)}{\mathbb{E}_{S\sim\mu(\lambda)}[\operatorname{uw}(S)]}$$

Here, Λ is the set of all possible instances. The definition corresponds to the approximation ratio studied for randomized algorithms, however for the approximation ratio, typically the constraint is computational complexity, whereas for us there are many constraints, including the input format and truthfulness.

Truthfulness. Truthful ordinal mechanisms are our focus in Sections 4 and 5, and we define truthfulness with respect to these mechanisms. A mechanism is truthful if in any instance, each agent obtains maximum utility in expectation by reporting the linear order $\prec_i$ consistent with its utility function. Formally, mechanism μ is truthful if for any vector $\vec{\precsim}' = (\prec'_i)_{i\in N}$ of linear orders over $\mathcal{C}$, and every agent j, if $\prec_j$ is the linear order consistent with agent j's utility function u_j, then

$$\mathbb{E}_{S\sim\mu(\vec{\precsim}'/_j\prec_j)}[u_j(S)] \geq \mathbb{E}_{S\sim\mu(\vec{\precsim}')}[u_j(S)]$$

Note that we do not insist that the components of $\vec{\precsim}'$ are consistent with the utility functions of the agents other than j. Hence we require truthfulness to be a dominant strategy.

For a nonnegative integer n, $H_n = \sum_{i=1}^{n} 1/i$ is the nth Harmonic number, and $\log(n+1) \leq H_n \leq 1 + \log n$.

3 Independent Threshold Mechanisms

We start with a mechanism that achieves nearly optimal distortion (close to 1) when the number of agents is large. Our mechanism uses the idea of randomized thresholds as in (Benade et al. 2017) but presents each agent with an independent random threshold. We note that Benade et al. show a lower bound of $\Omega(\log m/\log\log m)$ on the distortion for mechanisms that present the same threshold to all agents. Our mechanism places the same informational load on each agent but obtains a significantly lower distortion. The mechanism chooses a threshold for each agent uniformly and independently from the interval $[0, 1]$, and asks the agent for all candidates with utility above the threshold. For each candidate, the count of all thresholds that it exceeds serves as an unbiased estimator of its total utility over the agents, which is then used to compute the optimal budget-feasible set.

Mechanism 1 INDEPENDENT-THRESHOLDS

1: **for** each agent i **do**
2: $T_i \sim \mathcal{U}[0, 1]$ ▷ Randomized threshold for agent i
3: $S_i \leftarrow \{x \in \mathcal{C} : u_i(x) \geq T_i\}$ ▷ agent i returns all candidates with value above the threshold
4: **for** each candidate $x \in \mathcal{C}$ **do**
5: **if** $x \in S_i$ **then** $V_{i,x} \leftarrow 1$ **else** $V_{i,x} \leftarrow 0$
6: $\bar{V}_x \leftarrow \frac{1}{n}\sum_i V_{i,x}$ for each candidate x
7: **return** $S \leftarrow \arg\max_{T\subseteq\mathcal{C}:\sum_{x\in T} c_x\leq 1}\{\sum_{x\in T}\bar{V}_x\}$

Theorem 1. *Let* $\delta = m^2\sqrt{\frac{2\log(2mn)}{n}}$. *With probability at least* $1 - 1/n$ *the set* S *output by* INDEPENDENT-THRESHOLDS *has* $\operatorname{uw}(S) \geq \mathbf{OPT}(1-\delta)$

The proof uses the fact that $\bar{V}_x$ is an unbiased estimator the the utility $\operatorname{uw}(x)/n$, and then uses Chernoff bounds, appropriately.

Next we observe that we can obtain near-optimal distortion even by mechanisms with a significantly smaller reporting load on the agents. Indeed, the following mechanism elicits a binary input from each voter. Each voter is presented with a random threshold and a candidate, and asked if their utility for the candidate is above the threshold. Similar to INDEPENDENT-THRESHOLDS, these binary inputs are then used to obtain an estimate of the welfare of each candidate, which in turn is used to determine the budget-feasible set of candidates with maximum utilitarian welfare.

[3]Benade et al. (2017) also analyze a version of independent threshold approval votes when each agent gets the same threshold. Their mechanism input is thus also randomized. And in communication with the authors, it appears the distortion bounds they obtained hold under the simpler definition of distortion.

Mechanism 2 I.I.D.-THRESHOLDS-AND-CANDIDATES

Let $A_1, A_2, \ldots, A_m$ be a uniformly random partition of N into m subsets of size $\lceil n/m \rceil$ or $\lfloor n/m \rfloor$.
for each candidate $x \in [m]$ **do**
  for each agent $i \in A_x$ **do**
    $T_i \sim \mathcal{U}[0,1]$ ▷ Random threshold for agent i
    $b_i \leftarrow 1$ if $u_i(x) \geq T_i$, else $b_i \leftarrow 0$ ▷ agent i returns if he values candidate x above threshold T_i
    if $b_i = 1$ **then**
      $V_{i,x} \leftarrow 1$
    else
      $V_{i,x} \leftarrow 0$
  $\bar{V}_x \leftarrow \frac{1}{|A_x|} \sum_{i \in A_x} V_{i,x}$
$S \leftarrow \arg\max_{T \subseteq \mathcal{C}: \sum_{x \in T} c_x \leq 1} \{\sum_{x \in T} \bar{V}_x\}$
return S

Theorem 2. *Let* $\delta = m^{5/2}\sqrt{\frac{72\log(4mn)}{n}}$. *With probability at least* $(1 - 1/n)$, *the set* S *output by* I.I.D.-THRESHOLDS-AND-CANDIDATES *satisfies* $\mathrm{uw}(S) \geq \mathbf{OPT}(1-\delta)$.

4 Truthful Ordinal Mechanisms

We now consider ordinal mechanisms. We first show that the lower bound of $\Omega(\sqrt{m})$ on the distortion of ordinal mechanisms (Boutilier et al. 2015) can in fact nearly be achieved by *truthful* ordinal mechanisms, and give truthful mechanisms with distortion $O(\sqrt{m \log m})$ for the k-selection problem and $O(\sqrt{m} \log m)$ for participatory budgeting. We then show a lower bound of $\Omega(\sqrt{m \log m})$ for truthful mechanisms for the k-selection problem, proving our mechanism in this case has optimal distortion.

Our results in this section rely on the characterization of truthful ordinal mechanisms by Gibbard (1977) and Barbera (1978), and the welfare of truthful ordinal mechanisms by Filos-Ratsikas and Miltersen (2014). We first present these characterizations.

We will need the following definitions. A decision scheme is *anonymous* if it does not depend on the identities of agents. A decision scheme is *neutral* if it does not depend on the identities of candidates. Formally, let ν be a permutation of the set of voters. Thus if $\nu(\overrightarrow{\prec}) = \overrightarrow{\prec}'$, then $\prec_i = \prec'_{\nu(i)}$, for $i \in N$. Mechanism μ is anonymous if for any permutation ν of the set of voters, and for any input $\overrightarrow{\prec}$, $\mu(\nu(\overrightarrow{\prec})) = \mu(\overrightarrow{\prec})$.

Similarly, to define neutrality, let κ be a permutation of the set of candidates. If $\kappa(\overrightarrow{\prec}) = \overrightarrow{\prec}'$, then for each voter i and candidates $x, x' \in \mathcal{C}$, $x \preceq_i x'$ iff $\kappa(x) \preceq'_i \kappa(x')$. A mechanism μ is neutral if for any permutation κ on the set of candidates, and for any input $\overrightarrow{\prec}$, $\mu(\kappa(\overrightarrow{\prec})) = \kappa(\mu(\overrightarrow{\prec}))$.

A *point voting mechanism* assigns a weight w_i to each position $i \in [m]$ in a ranking. Such a mechanism picks a voter uniformly at random, and chooses the candidate ranked r with probability w_r. Finally, a supporting size mechanism is given by $n+1$ real numbers $w_0, w_1, \ldots, w_n$, such that $b_r + b_{n-r} = 1$. It picks two candidates a, b uniformly at random, and if k voters prefer a to b, selects a with probability w_k (and b with probability w_{n-k}). Using the characterization by Gibbard, Barbera (1978) then shows the following result.

Theorem 3. *A decision scheme is anonymous, neutral, and truthful if and only if it is a distribution over point-voting and supporting-size decision schemes.*

Further, Filos-Ratsikas and Miltersen (2014) show that the distortion obtained by any truthful mechanism, can also be obtained by an anonymous and neutral truthful mechanism.

Theorem 4. *For any ordinal mechanism* μ, *there is an anonymous and neutral truthful mechanism* μ' *with distortion at most that of* μ.

We note that the above theorem is stated for the unit-range case, where for each agent, the utility for each candidate lies in $[0,1]$. In particular, the sum of candidate utilities may not be 1. However, it can be easily seen that the proof in the earlier paper extends to the unit-sum case as well.

4.1 Near-optimal truthful ordinal mechanisms

Recall that in an ordinal mechanism, each agent i returns a linear order $\prec_i$ of $\mathcal{C}$. Let $\overrightarrow{\prec} := (\prec_i)_{i \in N}$. We say that $\overrightarrow{\prec}$ is consistent with a utility profile $\vec{u} = (u_i)_{i \in N}$, written $\overrightarrow{\prec} \cong \vec{u}$, if for each agent i, $u_i(x) > u(y)$ implies $x \prec_i y$. Given a linear order $\prec_i$ for agent i and a candidate x, $\mathrm{rk}_i(x)$ is the number of candidates that i prefers to x (including x), i.e., $|\{y : y \preceq_i x\}|$. Define $\mathrm{Sc}(x) := \sum_{i \in N} 1/\mathrm{rk}_i(x)$. As before, S^* is the feasible set of candidates with maximum welfare, and $\mathbf{OPT} = \mathrm{uw}(S^*)$.

We first give a randomized mechanism with distortion $O(\sqrt{m \log m})$ for the k-selection problem, where k candidates are to be chosen. Our mechanism runs the Harmonic Scoring mechanism (Boutilier et al. 2015) as the sampling subroutine, and outputs either the resulting *single candidate* with probability $1/2$, or a randomly chosen subset of size k.[4] We will use this mechanism as a subroutine later on with subsets of $\mathcal{C}$ as possible candidates, and hence explicitly give the set of candidates as an input.

Mechanism 3 k-WINNER-SELECTION

Require: Set A of m candidates, k
Let i be a randomly chosen agent. Sample z at random from A with probability proportional to $1/\mathrm{rk}_i(z)$. Let $Z \leftarrow \{z\}$.
Y is a set of size k, sampled uniformly from A.
S is chosen from $\{Y, Z\}$ with equal probability
return S

For any agent i, $\sum_{x \in A} \frac{1}{\mathrm{rk}_i(x)} = H_m$. Hence, for any candidate x,

$$\mathbb{P}[x \in Z] = \sum_{i \in N} \frac{1}{n} \frac{1}{H_m \mathrm{rk}_i(x)} = \frac{\mathrm{Sc}(x)}{n H_m}\,. \tag{1}$$

[4] If it is important that a subset of candidates of size k be returned, in the first case, we can always add $k-1$ randomly chosen candidates. This affects neither the truthfulness nor the upper bound on the distortion.

The truthfulness of k-WINNER-SELECTION is immediate from Theorem 3 since Y is chosen irrespective of agent preferences, and Z is chosen by a point-voting mechanism. We now prove the bound on the social welfare.

Theorem 5. *The expected social welfare of S is at least $\mathbf{OPT}/4\sqrt{m \log m}$.*

Proof sketch. Let S^* be the optimal set of candidates, S_1^* the set of candidates in S^* with score at least $n\sqrt{\log m/m}$, and $S_2^* = S^* \setminus S_1^*$. We will show that in expectation, the social welfare of Y and Z are at least $1/2\sqrt{m \log m}$ times the social welfare of S_2^* and S_1^* respectively. The expected welfare of $Y + Z$ is then at least $\mathbf{OPT}/2\sqrt{m \log m}$. Since $\mathbb{E}[\mathrm{uw}(S)] = (\mathbb{E}[\mathrm{uw}(Y) + \mathrm{uw}(Z)])/2$, the proof follows.

For any candidate x, it can be shown that $\mathrm{Sc}(x) \geq \mathrm{uw}(x)$. This gives the following upper bound on the utility of S_2^* :

$$\mathrm{uw}(S_2^*) \leq \sum_{x \in S_2^*} \mathrm{Sc}(x) \leq |S_2^*|\, n\sqrt{\frac{\log m}{m}} \leq kn\sqrt{\frac{\log m}{m}}\,.$$

The expected social welfare of Y is exactly kn/m, and substituting gives us that $\sqrt{m \log m}\mathbb{E}[\mathrm{uw}(Y)] \geq \mathrm{uw}(S_2^*)$, as required. For the bound on Z, we use the fact that candidates are chosen in proportion to their score, which for each candidate in S_1^* is at least $n\sqrt{\log m/m}$, to show that

$$\mathbb{E}[\mathrm{uw}(Z)] \geq \sum_{x \in S_1^*} \mathrm{uw}(x)\frac{\mathrm{Sc}(x)}{nH_m} \geq \frac{\mathrm{uw}(S_1^*)}{2\sqrt{m \log m}}\,.$$

The first inequality is from (1). The bound on $\mathbb{E}[\mathrm{uw}(S)]$ follows. $\square$

We now adapt k-WINNER-SELECTION to general costs. Our mechanism uses the Ranking-by-Value Mechanism from Benade et al. (2017), except that we use the above k-winner selection mechanism in place of Mechanism A by Benade et al. to recover truthfulness. The mechanism first divides the interval $[0, 1]$ of costs into $1 + \log m$ buckets, with bucket s (roughly) consisting of all candidates with costs in the interval $[2^{s-1}/m, 2^s/m]$. If m_s is the number of candidates in bucket s, it chooses bucket s with probability proportional to $\sqrt{m_s \log m_s}$ (Benade et al. choose a bucket with uniform probability). It then uses k-WINNER-SELECTION to select $m/2^s$ candidates from bucket s.

Mechanism 4 TRUTHFUL-RANKING-BY-VALUE

For $s \in [\log m]$, define $l_s = 2^{s-1}/m$, $u_s = 2^s/m$
Let $T_0 := \{x : c_x \leq \frac{1}{m}\}$, and $T_s = \{x : l_s < c_x \leq u_s\}$ for $s \in [\log m]$. Let $m_s = |T_s|$, $s \in [\log m] \cup \{0\} \triangleright 1/u_s$ candidates can be chosen from T_s within the budget.
Pick $r \in [\log m] \cup \{0\}$, where $\mathbb{P}[r = s] \propto \sqrt{m_s \log m_s}$
Run k-WINNER-SELECTION with inputs T_r and $k = \frac{1}{u_r}$. Let U be the set of candidates returned.
return U

For truthfulness of Mechanism 4, note that the sets T_s as well as r are decided independent of the reported linear orders. Given T_r, we can restrict attention to the valuations given by an agent j to candidates in T_r. The proof of truthfulness follows from the truthfulness of k-WINNER-SELECTION.

Theorem 6. TRUTHFUL-RANKING-BY-VALUE *has distortion $O\left(\sqrt{m} \log m\right)$.*

Proof. For $s \in [\log m] \cup \{0\}$, let T_s^* be the set of candidates in T_s that are budget-feasible and maximize the social welfare. Then since $c_x \geq l_s$ for $x \in T_s$ and the budget is 1, $|T_s^*| \leq 1/l_s = 2/u_s$. Let $T_s' \subseteq T_s$ be the set of candidates of size $1/u_s$ with maximum social welfare. Then

$$\mathrm{uw}(T_s') \geq \frac{1}{2}\mathrm{uw}(T_s^*) \geq \frac{1}{2}\mathrm{uw}(S^* \cap T_s)\,.$$

By Theorem 5, if $r = s$, then for the set U returned by the k-selection mechanism,

$$\mathbb{E}[\mathrm{uw}(U)] \geq \frac{1}{4}\frac{\mathrm{uw}(S^* \cap T_s)}{\sqrt{m_s \log m_s}}\,.$$

The expected social welfare of the set U returned is thus

$$\sum_{s=0}^{\log m} \mathbb{P}[r = s]\frac{1}{4}\frac{\mathrm{uw}(S^* \cap T_s)}{\sqrt{m_s \log m_s}} = \frac{\mathrm{uw}(S^*)}{4\sum_{s=0}^{\log m}\sqrt{m_s \log m_s}}\,.$$

Since $\sum_s m_s = m$ and $\sqrt{x \log x}$ is concave, the expression on the right is maximized when the m_s's are all equal, in which case the expected social welfare is $\mathbf{OPT}/\left(\sqrt{m} \log m\right)$. $\square$

4.2 A tight lower bound for truthful mechanisms

We now show that the bound on distortion obtained for the k-winner selection problem is tight.

Theorem 7. *Let μ be a truthful ordinal mechanism for the 1-selection problem. Then the distortion of μ is at least $\sqrt{m \log m}/8$.*

Proof sketch. Let $\gamma = \sqrt{m \log m}$ and μ be a mechanism with distortion strictly better than $\gamma/8$. From Theorem 4, there exists a mechanism μ' that is anonymous, ordinal, truthful, and with distortion better than $\gamma/8$. Further, from Theorem 3, mechanism μ' is a distribution over point-voting and supporting-size decision schemes. Our proof now proceeds by contradiction as follows.

We first show that we can ignore support-size mechanisms in the support of μ. A support-size mechanism picks a pair of candidates uniformly at random to compare, and hence the probability that any fixed candidate is chosen is at most $2/m$. Fix a candidate x_0, which will be our optimal candidate. We then focus on point-voting mechanisms. Define p_r to be the average probability that the point-voting mechanism selects the candidate ranked r. We construct a series of instances, one for each $r \leq 2m/\gamma$, to show that to get distortion better than $\gamma/8$, the probabilities p_r must follow a roughly harmonic progression, i.e., $p_r \geq 8/(r \log m)$. But summing p_r for $r \leq 2m/\gamma$ is at least 2, which is a contradiction. We recall that in a point-voting decision scheme, an agent is chosen uniformly at random.

The instance for a fixed $r \leq 2m/\gamma$ is constructed as follows. We first select a subset N_r of agents so that $|N_r| = (nr\gamma)/2m$, and choose utility functions and preference orders for the instance with the following properties:

1. All agents in N_r have rank r for candidate x_0, while all other agents place x_0 at rank m.
2. All agents in N_r have utility $1/r$ for the first r candidates, and 0 for the others. All other agents have utility $1/m$ for all candidates. To maintain our assumption that no agent has the same utility for two candidates, we can slightly perturb the utilities functions.
3. Restricted to agents in N_r, and for any rank $s \in [m]$, each candidate other than x_0 appears with approximately the same frequency at rank s.

It can be calculated that candidate x_0 has welfare at least $|N_r|/r$, while any other candidate has welfare at most $2n/m$. Also, the probability that μ' chooses x_0 is at most $3/m + p_r|N_r|/n$. Then the *inverse* of the distortion for the mechanism is at most

$$\frac{4}{\gamma} + \frac{3}{m} + p_r \frac{r\gamma}{2m}$$

which is at least $8/\gamma$, by assumption on the distortion. But solving gives us that $p_r \geq 8/(r \log m)$, giving us the required contradiction. □

5 A Mechanism for Two Candidates

In this section, we use a and b to denote the two candidates. We assume that exactly one candidate is to be selected, and as is in the rest of the paper, that the utilities of an agent for the two candidates are not equal. We start by showing that any truthful mechanism for selecting one of two candidates, whether randomized or deterministic, must be ordinal. Similar results were earlier obtained in the setting where the utilities of agents for the candidates are drawn from independent distributions (Azrieli and Kim 2014; Schmitz and Tröger 2012).

Theorem 8. *Let μ be a truthful mechanism for two candidates. If utility profiles $\vec{u}$, $\vec{u}'$ are consistent with the same linear order $\vec{\prec}$, then $\mu(\vec{u}) = \mu(\vec{u}')$.*

We thus focus on ordinal mechanisms. Candidates a and b can be represented by the end-points of the interval $[0, 1]$ respectively, and each voter i as the point $u_i(b)$ in the interval. A voter's utility for a candidate is thus 1 minus its distance from the candidate.

After querying all the agents and seeing that αn of them prefer a while the rest prefer b, our mechanism TEMPERED-MAJORITY chooses candidate a with probability $p(\alpha) = \frac{2\alpha - \alpha^2}{1+2\alpha-2\alpha^2}$.

Theorem 9. TEMPERED-MAJORITY *is truthful. Further, on any instance in which α fraction of the agents prefer one candidate a and $(1 - \alpha)$ prefer the other,* TEMPERED-MAJORITY *achieves a distortion of $1+2\alpha-2\alpha^2$. The worst-case distortion is thus 3/2.*

The truthfulness follows simply because $p(\alpha)$ is an increasing function of α. The proof for the distortion explicitly optimizes $p(\alpha)$ for every fraction of voters α. Since any truthful mechanism must also be ordinal (Theorem 8) in fact we prove that no truthful mechanism has distortion better than $1 + 2\alpha - 2\alpha^2$ on an input where α fraction of candidates prefer one candidate. Thus, not only does TEMPERED-MAJORITY have optimal expected distortion over truthful mechanisms in the worst-case, but in fact is optimal (in terms of expected distortion) for every ordinal input.

Instead of utilities, we can also define a voter's cost for a candidate as its distance from the candidate, and can then define the distortion of a mechanism appropriately. Randomized ordinal mechanisms for minimizing the distortion in the case of costs are studied previously (Anshelevich and Postl 2016). In particular, for the case of two candidates, the mechanism with minimum distortion is the PROPORTIONAL-TO-SQUARES mechanism (also called the square-weighted dictator mechanism (Meir, Procaccia, and Rosenschein 2012)), which has distortion 2, and this is tight. Surprisingly, however, with distortion defined in terms of utilities rather than costs, the PROPORTIONAL-TO-SQUARES mechanism is *not* optimal. Consider the instance where $\alpha = 0.64$ fraction of the candidates prefer candidate a. The PROPORTIONAL-TO-SQUARES mechanism chooses candidate a with probability $\frac{\alpha^2}{(1-\alpha)^2+\alpha^2}$, which evaluates to 0.7596. The distortion in terms of utilities in this instance can be shown to be larger than 1.672. Conversely, the TEMPERED-MAJORITY mechanism performs poorly when applied to the setting with costs.

Conclusion. Our work studies mechanisms that maximize utilitarian welfare, while simultaneously placing low informational burden on voters. Our threshold mechanisms are interesting because the conditions for welfare maximization — a large number of voters — are satisfied in practice. The mechanisms are not ordinal, but informally use minimal cardinal information since each voter must only report a subset of the candidates, or a simple binary input, rather than the cardinal utilities themselves. The use of cardinal information is necessary for any mechanism to circumvent the lower bound of $\Omega(\sqrt{m})$ on the distortion for ordinal mechanisms (Boutilier et al. 2015). Our mechanism is also not truthful, but this may be compensated in practice by the near-optimal welfare obtained.

For truthful ordinal mechanisms, we obtain a tight bound of $\Theta(\sqrt{m \log m})$ for the distortion of k-winner selection mechanisms and a nearly tight upper bound of $O(\sqrt{m} \log m)$ for participatory budgeting. It is interesting to note that changing the normalization of agent utilities — unit-range vs. unit-sum — significantly changes the distortion achievable, from $O(m^{2/3})$ in the former case to $O(\sqrt{m \log m})$ in the latter case. Similarly, in the case of two candidates, the PROPORTIONAL-TO-SQUARES mechanism which is optimal when agents have costs performs poorly when agents have utilities. This sounds a note of caution if minimizing distortion is the goal: Changes in the input format that appear inconsequential may significantly change the distortion of mechanisms.

Acknowledgments. Part of this work was done while the second and third authors were visiting TIFR. The first author was partially funded by a Ramanujan Fellowship. We thank Swaprava Nath, Tom Hayes, Ariel Procaccia, and Nisarg Shah for discussing the problem, as well as anonymous referees for helping improve the presentation.

References

Anshelevich, E., and Postl, J. 2016. Randomized social choice functions under metric preferences. In *IJCAI '16, New York, NY, USA*, 46–59.

Anshelevich, E.; Bhardwaj, O.; and Postl, J. 2015. Approximating optimal social choice under metric preferences. In *AAAI '15, Austin, Texas, USA*, 777–783.

Azrieli, Y., and Kim, S. 2014. Pareto efficiency and weighted majority rules. *International Economic Review* 55(4):1067–1088.

Barbera, S. 1978. Nice decision schemes. *Decision Theory and Social Ethics: Issues in Social Choice Dordrecht, Holland: D. Reidel Publishing Co* 101–117.

Benade, G.; Nath, S.; Procaccia, A. D.; and Shah, N. 2017. Preference elicitation for participatory budgeting. In *AAAI '17, San Francisco, California, USA*. Forthcoming.

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

Caragiannis, I.; Nath, S.; Procaccia, A. D.; and Shah, N. 2017. Subset selection via implicit utilitarian voting. *J. Artif. Intell. Res. (JAIR)* 58:123–152.

Chajewska, U.; Koller, D.; and Parr, R. 2000. Making rational decisions using adaptive utility elicitation. In *AAAI/IAAI*, 363–369.

Elkind, E.; Faliszewski, P.; Laslier, J.; Skowron, P.; Slinko, A.; and Talmon, N. 2017. What do multiwinner voting rules do? an experiment over the two-dimensional euclidean domain. In *AAAI '17, San Francisco, California, USA*, 494–501.

Feldman, M.; Fiat, A.; and Golomb, I. 2016. On voting and facility location. In *EC '16, Maastricht, The Netherlands*, 269–286.

Filos-Ratsikas, A., and Miltersen, P. B. 2014. Truthful approximations to range voting. In *WINE '14, Beijing, China*, 175–188. Springer.

Gibbard, A. 1973. Manipulation of voting schemes: a general result. *Econometrica: Journal of the Econometric Society* 587–601.

Gibbard, A. 1977. Manipulation of schemes that mix voting with chance. *Econometrica: Journal of the Econometric Society* 665–681.

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

Mclean, R. 2016. Americans were stuck in traffic for 8 billion hours in 2015. *CNN Money*.

Meir, R.; Procaccia, A. D.; and Rosenschein, J. S. 2012. Algorithms for strategyproof classification. *Artif. Intell.* 186:123–156.

Procaccia, A. D., and Rosenschein, J. S. 2006. The distortion of cardinal preferences in voting. In *International Workshop on Cooperative Information Agents*, 317–331. Springer.

Satterthwaite, M. A. 1975. Strategy-proofness and arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory* 10(2):187–217.

Schmitz, P. W., and Tröger, T. 2012. The (sub-) optimality of the majority rule. *Games and Economic Behavior* 74(2):651–665.

Skowron, P.; Faliszewski, P.; and Lang, J. 2016. Finding a collective set of items: From proportional multirepresentation to group recommendation. *Artificial Intelligence* 241:191–216.

Wakker, P., and Deneffe, D. 1996. Eliciting von neumann-morgenstern utilities when probabilities are distorted or unknown. *Management Science* 42(8):1131–1150.