

Worst-Case Voting When the Stakes are High

Anson Kahng,¹ Gregory Kehne²

¹ University of Toronto

² Harvard University

anson@cs.toronto.edu, gkehne@g.harvard.edu

Abstract

We study the additive distortion of social choice functions in the implicit utilitarian model, and argue that it is a more appropriate metric than multiplicative distortion when an alternative that confers significant social welfare may exist (i.e., when the stakes are high). We define a randomized analog of positional scoring rules, and present a rule which is asymptotically optimal within this class as the number of alternatives increases. We then show that the instance-optimal social choice function can be efficiently computed. Next, we take a beyond-worst-case view, bounding the additive distortion of prominent voting rules as a function of the best welfare attainable in an instance. Lastly, we evaluate the additive distortion of a range of rules on real-world election data.

1 Introduction

Distortion is a widely-used metric that captures the worst-case loss in efficiency of a social choice function (SCF) [4]. It is defined in the *implicit utilitarian* model where voters have cardinal utilities for alternatives but only report ordinal information, e.g., (partial) rankings, to the social choice function, which then outputs a distribution over winning alternatives.

Distortion evaluates SCFs according to their worst-case performance over all implicit utilities and corresponding induced rankings, where performance is measured in terms of (*utilitarian*) *social welfare*, i.e. the sum of all agents' utilities. Specifically, the distortion of a rule is the maximum ratio between the social welfare of the optimal alternative and the expected social welfare given by the rule.

While utilitarian social welfare is a defensible basis on which to evaluate social choice functions, distortion is not always the best tool for the job. In particular, we might prefer a social choice function which delivers poor multiplicative guarantees on instances where *no alternative confers significant social welfare*, so long as it performs well on instances where the potential gains are large. For example, a $1/\sqrt{m}$ -approximation is a much more tolerable loss when the maximum attainable social welfare is $O(\log n)$ (as for a symmetric profile with n alternatives) than when it is fully $\Omega(n)$.

Indeed, the canonical instance which demonstrates a $\Omega(\sqrt{m})$ lower bound on distortion for randomized social choice functions [9] allots at most a $1/\sqrt{m}$ proportion of the total utility to any alternative. In practice—for example, in political contests—we often expect that there are alternatives which confer *much* larger social welfare than the average alternative.

To address these concerns we instead study the *additive distortion* of randomized social choice functions, which may be viewed as their worst-case expected regret [10]. The additive distortion of a social choice function is the *difference* between the maximum social welfare attainable and the expected social welfare that f delivers, in the worst case over all implicit utilities. Different profiles in the implicit utilitarian model can have vastly different maximum attainable social welfare, and we posit that, in evaluating social choice functions, additive distortion appropriately prioritizes the instances in which the most utility can be gained or lost.

In its introduction to the social choice setting, distortion was compared to the distortion of metric embeddings [21]; this additive distortion is similarly analogous [18].

Although we advocate for additive distortion primarily on the above grounds, another advantage is that it remains a meaningful worst-case metric under weaker assumptions about voters' utilities. Past work on distortion in the (non-metric) implicit utilitarian model has made the assumption that all voters' utilities are *unit-sum* [21, 11, 9, 10, 7]. This is not a coincidence: with potentially apathetic voters whose utilities are instead *unit-capped*, one can show that choosing an alternative uniformly at random (incurring distortion m) is optimal, and that the distortion of any deterministic rule is infinite. Although there are compelling reasons to uniformly cap the sum of voters' utilities at one [6], the assumption that all participating voters' total utility is *equal* is often unreasonable in practice. As we will show in Section 3, additive distortion provides a discerning metric by which to evaluate SCFs in this broader context.

In this work we aim to answer the following questions:

Question 1: *What is the best additive distortion attainable for randomized social choice functions?*

Question 2: *How well do prominent social choice functions perform with respect to additive distortion, both in theory and in practice?*

Our Results In the pursuit of randomized SCFs with low additive distortion, we define a natural class of rules which we call *randomized scoring rules*, which are the natural randomized analog of positional scoring rules. A randomized scoring rule (RSR) first computes aggregate scores based on a scoring vector (as scoring rules do), and then chooses each alternative with probability proportional its score. Like scoring rules, RSRs are both intuitive and easy to compute. The two most prominent RSRs—Randomized Dictatorship, and the harmonic rule of [9]—are nearly distortion-optimal in the normalized utility and metric settings, respectively. When considered together with our results, we argue that RSRs merit wider attention in the study of distortion.

In Section 3 we address Question 1. We establish that Randomized Dictatorship (RD) has additive distortion $\frac{1}{2}(1 - 1/m) \cdot n$, and lower bound the best additive distortion obtainable by any randomized social choice function. We then present the *Best-or-Bust* (BoB) rule, which has distortion at most $\frac{11}{27} \cdot n$ and asymptotically minimizes additive distortion within the class of all randomized scoring rules. In particular, this establishes an asymptotic separation between deterministic and randomized voting rules with respect to additive distortion, even as m becomes large. We also show that the obstructions to minimizing additive distortion are information-theoretic rather than computational by presenting an instance-optimal randomized social choice function which can be computed efficiently.

In Section 4 we present an alternative metric for prioritizing the worst-case performance on instances with high attainable social welfare, which we call *promise distortion*. This is a beyond-worst-case guarantee that some alternative confers social welfare at least $\alpha \cdot n$, for some $\alpha \in [0, 1]$. We analyze the extent to which multiplicative promise distortion circumvents the $\Omega(\sqrt{m})$ lower bound of [9], relate it to additive distortion, and provide an analysis of some social choice functions with respect to both additive and multiplicative promise distortion.

We answer Question 2 in Sections 4 and 5. In Section 4 we analyze a range of prominent social choice functions through the lens of additive distortion, providing upper and lower bounds on their worst-case performance.

In Section 5, we evaluate the performance of our asymptotically optimal positional scoring rule against other scoring rules commonly used in practice, optimal randomized and deterministic algorithms for additive distortion, and an optimal randomized algorithm for (multiplicative) distortion. We observe that the optimal algorithm for multiplicative distortion is no longer optimal for additive distortion, and that the Plurality RSR performs the best on profiles encountered in practice, which suggests that, in practice, votes are far from worst-case instances.

1.1 Related Work

Distortion was first introduced by Procaccia and Rosenschein [21] in the context of deterministic single-winner social choice functions and normalized utilities. In a later paper, Caragiannis and Procaccia [11] proved that the Plurality rule has a distortion of $O(m^2)$, and further work demon-

strated that this is the best possible distortion of any deterministic voting rule [10].

Beyond deterministic social choice functions, Boutilier et al. [9] initiated the study of average-case analysis of randomized social choice functions under distributional assumptions about utilities. They also showed an $\Omega(\sqrt{m})$ lower bound on the distortion of any randomized rule in the worst case, and introduced a pair of voting rules with distortion $O(\sqrt{m} \cdot \log^* m)$ and $O(\sqrt{m} \cdot \log m)$, the latter of which makes use of the harmonic scoring vector. Caragiannis et al. [10] introduced regret to the implicit utilitarian model of voting; the regret that they study is equivalent to additive distortion in their unit-sum utility setting. They study choosing a k -subset of alternatives when social welfare is linear in the winners. For $k = 1$ and deterministic rules, their straightforward claims apply to additive distortion also; for randomized rules their results imply a $\frac{1}{4} \cdot n$ lower bound on additive distortion and a rule with at most $\frac{1}{2}(1 - \frac{1}{m^2}) \cdot n$ additive distortion. We show better upper and lower bounds for randomized rules.

Multiplicative distortion has also received attention in the metric setting. There voters and alternatives sit in a metric space, distances are costs, and one generally aims to minimize the social cost of a chosen alternative, given only voters’ rankings. Anshelevich et al. [2] first studied metric distortion, demonstrating that the Copeland rule has a distortion of 5, in stark contrast to the bounds of the unit-sum utility setting. They also conjectured that the deterministic lower bound of 3 is tight, and many papers made progress toward this conjecture [23, 14, 20, 16] before its ultimate proof by Gkatzelis, Halpern, and Shah [13]. Here again randomized rules do better: Anshelevich and Postl [5] showed that Randomized Dictatorship has distortion at most $3 - 2/n$ and gave a lower bound of 2 on the distortion of all randomized rules in the metric setting. Kempe [17] and Gkatzelis, Halpern, and Shah [13] each present rules attaining $3 - 2/m$, and Anshelevich and Postl [5] and Fain et al. [12] study variants of the randomized dictatorship mechanism. Lastly, Seddighin, Latifian, and Ghodsi [22] studies distortion when some voters may abstain. Unfortunately additive distortion is uninteresting here because there is no (dis)utility normalization—additive distortion is made arbitrarily large by rescaling an instance. For a comprehensive survey of works concerning multiplicative distortion, see [4, 3].

Finally, we study a class of SCFs which are the randomized analog of positional scoring rules. Young [26] characterized deterministic scoring functions (with rounds of tiebreaking) as the SCFs which are anonymous, neutral, and consistent, and Xia and Conitzer [25] provide a striking deterministic generalization of scoring rules. Walsh and Xia [24] and Bentert and Skowron [8] present schemes which may be viewed as randomized generalizations of scoring rules, where deterministic rules are applied to profiles formed by subsampling voters and alternatives, respectively.

2 Setting and Definitions

Consider voters $N = [n]$ and alternatives A , with $|A| = m$. Each voter $i \in N$ has a ranking σ_i over A which is a strict total order; we say that $a \succ_i b$ for alternatives $a, b \in A$ if $\sigma_i(a) < \sigma_i(b)$. The collection of rankings $\sigma = (\sigma_i)_{i \in N}$ is a profile; let $\Sigma := S_A^n$ denote the collection of all profiles.

Voters have implicit utilities $u_i \in \mathbb{R}_+^A$ which are consistent with their rankings; that is, if $a \succ_i b$ then $u_i(a) \geq u_i(b)$. We say that $u \triangleright \sigma$ for a collection of utilities u if u_i is consistent with σ_i for all voters i . Weakening the standard unit-sum implicit utility assumption, we assume:

Assumption 2.1. *The total utility of each voter is unit-capped at $\sum_{a \in A} u_i(a) \leq 1$ for all voters i .*

Given a profile σ , a deterministic social choice function $f : \Sigma \rightarrow A$ chooses an alternative to be the winner for this profile. Similarly, a randomized social choice function $f : \Sigma \rightarrow \Delta_A$ returns a probability distribution over winners, where Δ_A is the probability simplex over A . Here SCFs are randomized unless otherwise stated.

Perhaps the most prominent class of deterministic SCFs are *scoring functions*, or (*positional*) *scoring rules* (SRs). Each SR f^s is given by a scoring vector $s \in \mathbb{R}^m$. It first assigns to each alternative $a \in A$ the aggregate score $S_a := \sum_i s_{\sigma_i^{-1}(a)}$, which is the score associated with each voter i 's ranking of a , summed over all voters. The alternative with the maximum score is then chosen. Scoring functions can handle ties either by returning the set of alternatives with maximal scores, or by using additional scoring vectors to iteratively break ties.

As outlined above, the multiplicative distortion of a randomized SCF f is the worst-case ratio

$$\text{dist}(f) := \max_{\sigma} \max_{u \triangleright \sigma} \frac{\max_{a^* \in A} \text{sw}(a^*)}{\mathbb{E}_{a \sim f(\sigma)}[\text{sw}(a)]},$$

over all profiles σ and utility profiles u consistent with σ , where $\text{sw}(a)$ denotes the social welfare of a : $\text{sw}(a) := \sum_{i \in N} u_i(a)$. Additive distortion is the difference, rather than the ratio:

$$\text{dist}^+(f) := \max_{\sigma} \max_{u \triangleright \sigma} \left(\max_{a^* \in A} \text{sw}(a^*) - \mathbb{E}_{a \sim f(\sigma)}[\text{sw}(a)] \right).$$

For beyond-worst-case distortion, we will use the following notion of a utility promise:

Definition 2.2. *The utility profile u satisfies an α -promise on its maximum social welfare if there exists some alternative $a \in A$ for which $\text{sw}(a) \geq \alpha \cdot n$.*

2.1 Randomized Scoring Rules

Towards the goal of minimizing additive distortion, we find it compelling to define the following class of SCFs:

Definition 2.3. *A randomized scoring rule (RSR) is an SCF given by a scoring vector $s \in \mathbb{R}_+^m - \mathbf{0}$. The aggregate scores S_a are calculated in the same way as for scoring rules, and then each alternative is chosen to be the winner with probability proportional to its total score. Let \mathcal{RSR} denote the class of all such rules.*

Example 2.4. *Consider the RSR given by $s = (2, 1, 0)$ and a profile σ with $n = 3$ and $m = 3$ where two voters report $a_1 \succ a_2 \succ a_3$ and one voter reports $a_3 \succ a_2 \succ a_1$. The total score of a_1 is $2 + 2 + 0 = 4$, the total score of a_2 is $1 + 1 + 1 = 3$, and the total score of a_3 is $0 + 0 + 2 = 2$. Therefore, the probability that a_1 wins the election is $4/9$, the probability that a_2 wins the election is $1/3$, and the probability that a_3 wins the election is $2/9$.*

Just as the prominent rules Plurality, Borda Count, and Veto belong to the class of deterministic SRs, \mathcal{RSR} also contains noteworthy rules. One is the harmonic scoring vector-based rule of Boutilier et al. [9] mentioned above, which is nearly optimal for multiplicative distortion. It is given by $s = (1 + H_m/m, 1/2 + H_m/m, \dots, 1/m + H_m/m)$, where H_m is the m^{th} harmonic number. Another is Randomized Dictatorship, given by $s = (1, 0, \dots, 0)$. Remarkably, RD incurs $O(3 - 2/n)$ multiplicative distortion in the metric setting, which is also nearly optimal [5].

In principle, there are many ways in which an aggregate score vector S can be converted to a probability distribution over A . Let us call $P : \mathbb{R}_+^m - \mathbf{0} \rightarrow \Delta_A$ a *probabilizer*, and focus on neutral probabilizers, i.e., the P which commute with all permutations of A . Then a *generalized RSR* consists of a pair (s, P) of scoring vector and neutral probabilizer; given σ it first computes S according to s , then samples from the distribution $P(S)$. Let \mathcal{RSR}^* denote the class of all such SCFs. This is indeed a generalization, since any RSR given by s is a generalized RSR with the probabilizer $P(S)_a := S_a / \|S\|_1$ for all a , where $\|S\|_1 := \sum_{a \in A} S_a$. Note that \mathcal{RSR}^* also contains all (otherwise deterministic) scoring rules that break ties uniformly at random. For a given scoring vector s the scoring rule is given by (s, P) , where P returns the uniform distribution over $\arg \max_a S_a$. In fact, \mathcal{RSR}^* also generalizes the ‘‘favorite only’’ rules which have received recent attention for metric distortion; in addition to RD these include the ‘‘proportional to squares’’ mechanism studied in [5] and the Random Oligarchy mechanism of [12].

3 Additive Distortion

We begin by proving a structural lemma which establishes that, for worst-case additive distortion, voter utilities may be assumed to be normalized without loss of generality. The proof (and all other omitted proofs in this paper) can be found in the Supplementary Material.

Lemma 3.1. *For each SCF f , the utility profile that witnesses the maximum of $\text{dist}^+(f)$ is normalized, i.e., $\sum_a u_i(a) = 1$ for all voters $i \in [n]$.*

With this lemma in hand, we next show that, in the worst case, additive distortion can inevitably be quite large.

Claim 3.2. *For all SCFs f and $m \geq 3$, $\text{dist}^+(f) \geq \frac{5}{18} \cdot n$.*

Proof. We assume that $n = 3k$ for some positive integer k , take $m = 3$, and let the alternatives be a_1, a_2 , and a_3 . Consider the profile in which $n/3$ voters believe $a_1 \succ a_2 \succ a_3$, $n/3$ voters believe $a_2 \succ a_3 \succ a_1$, and $n/3$ voters believe $a_3 \succ a_1 \succ a_2$. Let p_i be the probability that f chooses a_i , and without loss of generality assume that $p_1 \geq p_2 \geq p_3$.

Now, let the first $n/3$ voters have utilities $u(a_1) = u(a_2) = u(a_3) = 1/3$; the second $n/3$ voters have utilities $u(a_2) = u(a_3) = 1/2$ and $u(a_1) = 0$; and the last $n/3$ voters have utilities $u(a_3) = 1$ and $u(a_1) = u(a_2) = 0$.

Therefore, we have

$$\begin{aligned} \text{dist}^+(f, \sigma) &\geq \max_{a^* \in A} \text{sw}(a^*) - \mathbb{E}_{a \sim f(\sigma)}[\text{sw}(a)] \\ &= \frac{11}{18} \cdot n - \left(\frac{1}{9} \cdot p_1 + \frac{5}{18} \cdot p_2 + \frac{11}{18} \cdot p_3 \right) n \\ &\geq \frac{5}{18} \cdot n. \quad (\text{because } p_1 \geq p_2 \geq p_3) \end{aligned}$$

Note that this construction straightforwardly extends to any other $m > 3$. \square

For deterministic rules, these symmetric instances offer even stronger lower bounds. The following claim was shown by Caragiannis et al. [10] in a more general setting of choosing k winners out of m alternatives; for completeness, we reproduce the example for the single-winner setting below.

Claim 3.3 (Theorem 1 in [10]). *For all deterministic SCFs f and $m \geq 2$, $\text{dist}^+(f) \geq \frac{1}{2} \cdot n$.*

Proof. Let $m = 2$ and consider the profile σ with voters equally divided between $a_1 \succ a_2$ and $a_2 \succ a_1$. Suppose that f chooses a_2 . If the first group has utilities $u(a_1) = 1$, $u(a_2) = 0$ and the second has $u(a_1) = 1/2$, $u(a_2) = 1/2$, then we have

$$\text{dist}^+(f, \sigma) \geq \text{sw}(a_1) - \text{sw}(a_2) = \frac{1}{2} \cdot n.$$

This again extends to $m \geq 3$; for $m = 3$ the instance demonstrating Claim 3.2 also gives $\text{dist}^+(f) \geq \frac{1}{2} \cdot n$. \square

3.1 Two Alternatives

As a warm-up, we begin with the case when there are $m = 2$ alternatives. Here we may compute the optimal randomized SCF directly.

Claim 3.4. *For $m = 2$ alternatives, the optimal SCF chooses each $a \in A$ with probability proportional to the number of voters ranking a first.*

Note that since this is the optimal SCF, choosing an equally divided profile of voters yields a lower bound of $\text{dist}^+(f) \geq 1/4$ for all SCFs f , recovering that of [10].

It is also noteworthy that this rule is in \mathcal{RSR} :

Observation 3.5. *For $m = 2$ the optimal randomized rule belongs to \mathcal{RSR} , given by scoring vector $s^* = (1, 0)$.*

For more than two alternatives, the problem of identifying optimal SCFs or even optimal RSRs becomes difficult.

3.2 Plurality and RD

When there are two alternatives, it is intuitive that the best deterministic rule should choose the alternative most frequently ranked first. In the class of deterministic rules, it turns out that this is always the best possible, as shown by Caragiannis et al. [10] in the general setting of choosing k winners out of m alternatives.

Theorem 3.6 (Theorem 1 in [10]). *Plurality is an optimal deterministic SCF, with additive distortion $\frac{1}{2} \cdot n$.*

The randomized analog to Plurality is Randomized Dictatorship, and Section 3.1 revealed that RD is the optimal SCF in the two alternative setting, attaining additive distortion $\frac{1}{4} \cdot n$ and significantly outperforming Plurality. One might reasonably hope that RD continues to significantly outperform Plurality for $m \geq 3$. However, we show that this is not the case:

Theorem 3.7. *RD has additive distortion $\frac{1}{2} \left(1 - \frac{1}{m}\right) \cdot n$.*

In fact, we must incorporate more than just voters' first choices in order to asymptotically improve upon $\frac{1}{2} \cdot n$. In the spirit of Gross, Anshelevich, and Xia [15], who give a lower bound of $3 - 2/m$ on the distortion of favorite-only mechanisms in the metric setting, the proof of Theorem 3.7 can be modified in order to show that:

Claim 3.8. *All generalized RSRs $(s, P) \in \mathcal{RSR}^*$ with $s = (1, 0, \dots, 0)$ have additive distortion at least $\frac{1}{2} \left(1 - \frac{1}{m}\right) \cdot n$.*

Since RD is optimal within the class of favorite-only mechanisms, we continue the search for better rules among RSRs which score beyond voters' first choices.

3.3 An Asymptotically Optimal rule in \mathcal{RSR}

After the success in Section 3.1, we might hope to derive optimal RSRs for $m \geq 3$ directly. Unfortunately, the natural formulations of finding such optimal RSRs are nonconvex max-min optimization problems which we have been unable to solve. In order to render this problem tractable, we let a^* denote the alternative which maximizes social welfare, and we ignore the social welfare derived by choosing any alternative besides a^* . This provides an upper bound on the additive distortion of a given rule. We call this the *best-or-bust bound*, and we will use it repeatedly:

$$\text{dist}^+(f) \leq \text{sw}(a^*) (1 - \Pr[f(\sigma) = a^*]). \quad (1)$$

Informally speaking, this bound is apt because in the worst case and for large m , the non- a^* alternatives may evenly divide the remaining utility of the voters. In this case, the social welfare attained by choosing an alternative other than a^* is approximately $\frac{n - \text{sw}(a^*)}{m}$, and so (1) is asymptotically tight for \mathcal{RSR} as m becomes large.

We formulate the problem of finding the optimal RSR under eq. (1) in (9) below, and prove that the scoring vector which optimizes this problem is $s^* = (25/33, 7/33, 1/33, 0, \dots, 0)$ for all $m \geq 3$. Since it is the RSR which minimizes the upper bound eq. (1), we call this the *Best-or-Bust* (BoB) rule.

This in turn implies the following theorem:

Theorem 3.9. *For all $m \geq 3$, $\text{dist}^+(BoB) \leq \frac{11}{27} \cdot n$. It is furthermore a $\left(1 - \frac{16}{27} \frac{1}{m-1}\right)^{-1} \leq \left(1 + \frac{1}{m-1}\right)$ -approximation to the optimal RSR for all $m \geq 3$.*

We now set about formulating the problem of finding the RSR which minimizes the right-hand side of Equation (1). For a given choice of $\alpha \in [0, 1]$ and scoring vector $s = (s_1, \dots, s_m)$ for which $\|s\|_1 = 1$, we may parameterize the

solutions according to the optimum social welfare $\alpha \cdot n$ attainable. Let a^* be the alternative for which $\text{sw}(a^*) = \alpha \cdot n$; we will then consider the worst-case probability that the RSR f^s selects a^* .

To this end, let x_i denote the proportion of voters $[n]$ who rank a^* i^{th} . Note that since rankings are assumed to be complete, $\|x\|_1 = 1$. Since f^s is a randomized scoring rule, and the probability of f^s choosing a^* is less than 1, in the worst case a^* has maximum utility possible given its vector of ranking proportions x . Therefore we may assume that $\text{sw}(a^*) = n \cdot \sum_i \frac{1}{i} x_i$.

We may then identify the worst-case best-or-bust bound attained by s for given α by solving the linear program

$$D^+(s, \alpha) := \max \alpha - \alpha \sum_i s_i x_i \quad (2)$$

$$\text{s.t.} \quad \sum_i \frac{1}{i} x_i = \alpha \quad (3)$$

$$x \in \Delta_{[m]}. \quad (4)$$

The objective (2) is (up to scaling by n) equal to the best-or-bust bound, since $\text{sw}(a^*) = \alpha \cdot n$ and we have that $\text{sw}(a^*) \Pr[f^s(\sigma) = a^*] = \frac{\alpha \cdot n}{\sum_i s_i} \sum_i s_i x_i = \alpha \cdot n \sum_i s_i x_i$, since $\|s\|_1 = 1$ by assumption. By optimizing over α as well, we may similarly characterize $\text{dist}^+(f^s)$ as the optimal value of a quadratic program:

$$D^+(s) := \max D^+(s, \alpha) \quad (5)$$

$$\text{s.t.} \quad 0 \leq \alpha \leq 1, \quad (6)$$

where eq. (6) captures that $\text{sw}(a^*) \leq n$, since each voter's utilities are normalized to 1.

We might then hope to derive the optimal RSR directly, by solving $s^* := \arg \min_s D^+(s)$. This takes the following form:

$$s^* := \arg \min D^+(s) \quad (7)$$

$$\text{s.t.} \quad s \in \Delta_{[m]}.$$

Finally note that $\alpha = \sum_i x_i$; therefore constraints (3) and (4) imply (6). We may also replace α with $\sum_i \frac{1}{i} x_i$. Taken together, these let us rewrite (5) as follows:

$$D^+(s) := \max \left(\sum_i \frac{1}{i} x_i \right) \left(\sum_i (1 - s_i) x_i \right) \quad (8)$$

$$\text{s.t.} \quad x \in \Delta_{[m]}.$$

The general problem for which we hope to find optimal s^* is then

$$D^+ := \min_s \max_x \left(\sum_i \frac{1}{i} x_i \right) \left(\sum_i (1 - s_i) x_i \right) \quad (9)$$

$$\text{s.t.} \quad s, x \in \Delta_{[m]}.$$

Two Alternatives and the Harmonic RSR As noted in Observation 3.5, the optimal SCF when $m = 2$ is the RSR given by $s = (1, 0)$. On the other hand, for $m = 2$ the optimal scoring vector for the formulation (9) is $s_2^* = (2/3, 1/3)$.

This illustrates that the formulation above (and the best-or-bust bound) indeed only asymptotically capture the problem of identifying the optimal RSR for each m .

Incidentally, this $s_2^* = (2/3, 1/3)$ coincides with the harmonic scoring vector for $m = 2$. However this coincidence does not continue: for $m = 3$ the harmonic scoring vector $s = (6/11, 3/11, 2/11)$ does not coincide with the optimal score for (9), as we show in the next section. Indeed the harmonic RSR incurs an additive distortion of at least $(1 - H_m^{-1}) \cdot n$, which is witnessed by the profile in which all voters value the same alternative with utility 1. Since $H_m^{-1} = o(1)$, this incurs additive distortion which is asymptotically the worst possible.

Three Or More Alternatives One might expect that for each m there is a distinct randomized scoring rule with scoring vector s_m^* which optimizes (9). However it turns out that the same scoring vector s^* which (together with a suffix of trailing zeros) optimizes (9) for all $m \geq 3$ simultaneously.

Lemma 3.10. *For all $m \geq 3$, the unique optimal solution to (9) is the scoring vector $s^* = (25/33, 7/33, 1/33, 0, \dots, 0)$, obtaining the optimum objective value $\frac{11}{27} \approx 0.407$.*

This candidate optimizer s^* of (9) was first identified via computer-assisted search. We now prove that it is optimal.

The proof proceeds in two stages. We begin by restricting the inner problem (8) to a new problem $\overline{D}^+(s)$; this gives a corresponding relaxation of the outer problem (9). We then argue that, for $m = 3$, if $\overline{D}^+(s) \leq \frac{11}{27}$ then $s = s^*$. Given this s^* , we demonstrate that the objective does not increase when we move from the restricted inner problem to the general inner problem (8):

Lemma 3.11. *For $m = 3$, the unique optimal solution to (9) is the scoring vector $s^* = (25/33, 7/33, 1/33)$, obtaining objective value $\frac{11}{27}$.*

We finally show that this s^* does not incur a larger objective even for $m > 3$, and that for fixed $m > 3$ no other s can do better; this demonstrates that s^* optimizes (9) for all $m \geq 3$ simultaneously, proving Lemma 3.10. The proof of Theorem 3.9 then follows.

3.4 An Additive Distortion Instance-Optimal SCF

Although the RSR derived in Section 3.3 is asymptotically optimal within \mathcal{RSR} , we do not anticipate that it is optimal among all SCFs, even asymptotically. In pursuit of better rules, we turn to instance-optimal SCFs.

The instance-optimal SCF from the perspective of additive distortion, for any given profile σ , mimics the minimizer of $\text{dist}^+(f, \sigma)$ over all SCFs f (which for fixed σ are probability distributions over A). In particular,

$$\text{AddOpt}(\sigma) := \arg \min_f \text{dist}^+(f, \sigma)$$

$$= \min_{p \in \Delta_A} \max_{u \triangleright \sigma} \left(\max_{a^* \in A} \text{sw}(a^*) - \mathbb{E}_{a \sim p} [\text{sw}(a)] \right).$$

We make use of Lemma 3.1 to show the following, which we empirically test in Section 5:

Theorem 3.12. *For any profile σ , Algorithm 1 computes the distribution over A which minimizes (expected) additive distortion in polynomial time.*

Algorithm 1: ADDITIVEOPTIMAL

Input: Ranking $\sigma \in \mathcal{S}_A^n$

Output: Distribution $p^* \in \Delta_m$ minimizing $\text{dist}^+(p, \sigma)$

for $a, b \in A$ **do**

$w_a^b \leftarrow \sum_i (\sigma_i^{-1})^{-1} \mathbb{1}\{b \succ_i a\}$

end for

$w_a \leftarrow (w_a^b)_{b \in A}$ for each $a \in A$

$p^* \leftarrow \arg \min_p \{D : w_a^a - p^T w_a \leq D \forall a \in A, p \in \Delta_A\}$

return p^*

4 Distortion With a Promise

We began by motivating additive distortion based on the observation that traditional distortion may not be the best metric when the maximum social welfare attainable is *potentially* quite large. For a given profile σ , additive distortion provides a soft sort of guarantee with respect to maximum attainable welfare, in the following sense: for $u, u' \succ \sigma$ where the maximum attainable welfare is higher under u than u' , additive distortion measures the extent to which a SCF provides simultaneous guarantees for both utility profiles *simultaneously*, requiring additively better guarantees for u .

In this section we instead suppose we are *promised* that there exists an alternative with high social welfare, and ask about distortion subject to this promise. We define α -promise distortion as the distortion over all profiles (σ, u) for which u satisfies the α -promise of Definition 2.2:

Definition 4.1. *For $\alpha \in [0, 1]$, the α -promise distortion of a rule f is given by*

$$\text{dist}_\alpha(f) := \max_\sigma \max_{\substack{u \succ \sigma \\ u \in U_\alpha}} \text{dist}(f, \sigma),$$

where U_α is the collection of u satisfying the α -promise.

Since α -promise multiplicative distortion and additive distortion both address the high-stakes setting, our first result interrelates the two:

Claim 4.2. *For any randomized SCF f ,*

- *If $\text{dist}^+(f) \leq \beta \cdot n$, then $\text{dist}_\alpha(f) \leq \frac{\alpha}{\alpha - \beta}$.*
- *If $\text{dist}_\alpha(f) \leq \gamma$, then $\text{dist}^+(f) \leq \max(\alpha \cdot n, n - n/\gamma)$.*

In the promise setting, we might also hope to circumvent the relatively low-welfare $\Omega(\sqrt{m})$ lower bound given in [9]. Indeed, the lower bound instance in [9] translates directly into a lower bound on distortion with an α -promise:

Theorem 4.3. *For any randomized SCF f ,*

$$\text{dist}_\alpha(f) = \Omega(\min\{\sqrt{m}, 1/\alpha\}).$$

For upper bounds, a modification of the analysis of the *harmonic scoring function* h presented in [9] yields

Theorem 4.4. $\text{dist}_\alpha(h) = O\left(\frac{\log m}{\alpha}\right)$.

By Theorem 4.3 this is tight up to a $\log m$ factor for all $\alpha \geq 1/\sqrt{m}$.

4.1 Additive Distortion With a Promise

We now turn to α -promise *additive distortion*, which is defined analogously to Definition 4.1. In this subsection, we are focused on the *robustness* of each rule, where we ask how the additive distortion guarantees degrade with the promise α . Intuitively, this asks ‘‘How well do these rules perform when the winner is clear?’’ We consider $\alpha \geq 1/2$; for all $\alpha < 1/2$ we know the additive distortion is at most α .

We begin with three deterministic scoring rules:

- The *Plurality Rule* (f_{Plur}) is a deterministic scoring rule with score vector $s = (1, 0, \dots, 0)$.
- The *Harmonic Rule* (f_{Harm}) is a deterministic positional scoring rule with score vector $s = (1, 1/2, \dots, 1/m)$.
- The *Borda Rule* (f_{Borda}) is a deterministic positional scoring rule with score vector $s = (m-1, m-2, \dots, 0)$.

We begin by showing that Plurality and the Harmonic Rule are robust for $\alpha \geq 3/4$, but once $\alpha < 3/4$ their additive distortion becomes as bad as the worst case:

Claim 4.5. *For the Plurality Rule (f_{Plur}),*

$$\text{dist}_\alpha^+(f_{Plur}) = \begin{cases} 0 & \text{for } \alpha \geq 3/4 \\ 1/2 & \text{for } \alpha < 3/4. \end{cases}$$

Claim 4.6. *For the Harmonic rule (f_{Harm}),*

$$\text{dist}_\alpha^+(f_{Harm}) \begin{cases} = 0 & \text{for } \alpha \geq 3/4 \\ \geq 1/2 & \text{for } \alpha < 3/4. \end{cases}$$

Claim 4.7. *For the Borda rule (f_{Borda}),*

$$\text{dist}_\alpha^+(f_{Borda}) \begin{cases} = 0 & \text{for } \alpha \geq \frac{m-1}{m} \\ \geq \frac{m-1}{m} - \frac{1}{m^2} & \text{for } \alpha < \frac{m-1}{m}. \end{cases}$$

Plurality and the Harmonic Rule are robust for $\alpha \geq 3/4$, which is the largest possible interval of α on which any SCF can guarantee an α -promise additive distortion of 0. For smaller α the situation for the Borda Rule is much worse. In particular, Borda ceases to be robust as soon as α dips below $\frac{m-1}{m}$. Lastly, we consider Randomized Dictatorship:

Claim 4.8. *For Randomized Dictatorship,*

$$\text{dist}_\alpha^+(RD) = \begin{cases} 2\alpha(1-\alpha) - \frac{2(1-\alpha)^2}{m-1} & \text{for } \alpha \geq \frac{1}{2}\left(1 + \frac{1}{m}\right) \\ \frac{1}{2}\left(1 - \frac{1}{m}\right) & \text{for } \alpha < \frac{1}{2}\left(1 + \frac{1}{m}\right). \end{cases}$$

In particular, as we might expect for randomized rules, additive distortion decays smoothly towards 0 as $\alpha \rightarrow 1$.

5 Experiments

We evaluated the performance of various SCFs on four datasets of election data from PrefLib [19]:

- *Vermont* consists of data from public office elections in 2014 (15 different races, with 3 to 6 candidates and 532 to 1960 voters per race);
- *Glasgow* consists of data from the 2007 Glasgow City Council elections (21 wards, with 8 to 13 candidates and 5199 to 12744 voters per ward);

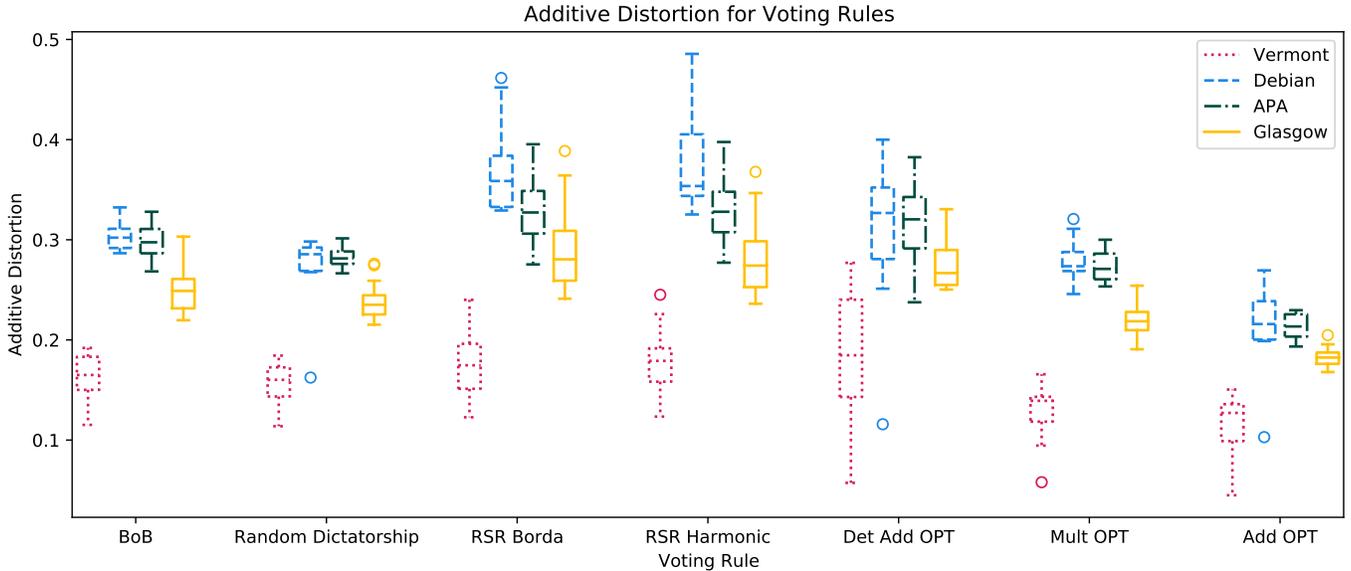


Figure 1: Additive distortion of voting rules on the Vermont, Glasgow, Debian, and APA datasets, normalized by n .

- *Debian* consists of votes for the Debian logo (8 elections, with 4 to 9 alternatives and 142 to 504 voters per election); and
- *APA* consists of election data from the American Psychological Association between 1998 and 2009 (12 elections, with 5 alternatives and 13318 and 20239 voters).

We also considered seven SCFs. Four of them are randomized scoring rules: *Randomized Dictatorship* has score vector $s = (1, 0, \dots)$ [1]; *RSR Borda* has score vector $s = (m - 1, m - 2, \dots, 0)$; *RSR Harmonic* has score vector $s = (1, 1/2, \dots, 1/m)$; and *BoB* has score vector $s = (25/33, 7/33, 1/33, 0, \dots)$. The other three are instance-optimal rules: *Det Add OPT* is the deterministic rule that minimizes additive distortion [10]; *Mult OPT* is the randomized rule that minimizes multiplicative distortion [9]; and *Add OPT* is the randomized rule that minimizes additive distortion based on Theorem 3.12.

Notably, all data was presented as a complete ranking that allowed ties between alternatives. Therefore in computing the rules, we split weight equally in the RSRs (i.e., if k alternatives were tied, they split the total score that the rule allocates over those k positions) and enforced the constraint that the implicit utility assigned to all tied alternatives is equal.

The additive distortions of each voting rule for each dataset are depicted in Figure 1. BoB generally outperforms Det Add OPT on all datasets, meaning that it results in lower additive distortion than *any* deterministic rule, which is why we compare its performance to the other randomized scoring rules RD, RSR Borda, and RSR Harmonic. We find that RD consistently outperforms BoB on the four datasets, while RSR Borda and RSR Harmonic both do worse. This is surprising, since Theorem 3.9 demonstrates that BoB is asymptotically worst-case optimal among the class of all RSRs. This suggests that real-life instances may not resemble worst-case additive distortion instances, and

that more “imbalanced” randomized positional scoring rules (with more precipitous drop-offs in scores after the first position) result in lower additive distortion in practice.

Notably, Caragiannis et al. [10] performed experiments in which Det Add OPT performed the best of the (deterministic) rules that they tested; the fact that both BoB and RD outperform Det Add OPT in terms of worst-case additive distortion is surprising and encouraging.

Additionally, there is a separation between the performance of Add OPT and Mult OPT (particularly for the Debian and APA datasets), which suggests that existing distortion-optimal rules do not optimize for additive distortion. Despite this separation, Mult OPT often outperforms the randomized positional scoring rules we implemented.

Furthermore, note that Add OPT significantly outperforms all rules on all elections. Encouragingly, calculating Add OPT is extremely efficient due to Theorem 3.12, and we expect that this approach is scalable to much larger elections. In comparison, Mult OPT took on the order of thousands of times longer than the others we tested.

6 Discussion

There are many exciting directions for future work. Most immediately, it would be nice to close the gap between our upper and lower bounds of $\frac{5}{18} \cdot n$ and $\frac{11}{27} \cdot n$ for randomized rules in the limit as m becomes large. It would also be interesting to explore the additive distortion guarantees of more rules (especially randomized rules) in the α -promise setting. We believe that is also worth further exploring the class of rules \mathcal{RSR}^* , since it features rules that perform remarkably well with respect to additive and multiplicative distortion in a range of settings. Finally, it would be interesting to characterize the instances on which multiplicative and additive distortion come apart; this could help to determine which distortion is the right fit in various settings.

References

- [1] Abdulkadiroğlu, A.; and Sönmez, T. 1998. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica*, 66(3): 689–701.
- [2] Anshelevich, E.; Bhardwaj, O.; Elkind, E.; Postl, J.; and Skowron, P. 2018. Approximating optimal social choice under metric preferences. *Artificial Intelligence*, 264: 27–51.
- [3] Anshelevich, E.; Filos-Ratsikas, A.; Shah, N.; and Voudouris, A. A. 2021. Distortion in social choice problems: an annotated reading list. *ACM SIGecom Exchanges*, 19(1): 12–14.
- [4] Anshelevich, E.; Filos-Ratsikas, A.; Shah, N.; and Voudouris, A. A. 2021. Distortion in Social Choice Problems: The First 15 Years and Beyond. arXiv:2103.00911.
- [5] Anshelevich, E.; and Postl, J. 2017. Randomized social choice functions under metric preferences. *Journal of Artificial Intelligence Research (JAIR)*, 58: 797–827.
- [6] Aziz, H. 2019. Justifications of welfare guarantees under normalized utilities. *ACM SIGecom Exchanges*, 17(2): 71–75.
- [7] Benadè, G.; Nath, S.; Procaccia, A. D.; and Shah, N. 2021. Preference Elicitation for Participatory Budgeting. *Management Science*, 67(5): 2813–2827.
- [8] Bentert, M.; and Skowron, P. 2020. Comparing Election Methods Where Each Voter Ranks Only Few Candidates. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, 2218–2225.
- [9] 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.
- [10] Caragiannis, I.; Nath, S.; Procaccia, A. D.; and Shah, N. 2017. Subset selection via implicit utilitarian voting. *Journal of Artificial Intelligence Research (JAIR)*, 58: 123–152.
- [11] Caragiannis, I.; and Procaccia, A. D. 2011. Voting almost maximizes social welfare despite limited communication. *Artificial Intelligence*, 175(9): 1655–1671.
- [12] Fain, B.; Goel, A.; Munagala, K.; and Prabhu, N. 2019. Random Dictators with a Random Referee: Constant Sample Complexity Mechanisms for Social Choice. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, 1893–1900.
- [13] Gkatzelis, V.; Halpern, D.; and Shah, N. 2020. Resolving the optimal metric distortion conjecture. In *Proceedings of the 61st Annual Symposium on Foundations of Computer Science (FOCS)*, 1427–1438.
- [14] Goel, A.; Krishnaswamy, A. K.; and Munagala, K. 2017. Metric distortion of social choice rules: Lower bounds and fairness properties. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, 287–304.
- [15] Gross, S.; Anshelevich, E.; and Xia, L. 2017. Vote Until Two of You Agree: Mechanisms with Small Distortion and Sample Complexity. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, 544–550.
- [16] Kempe, D. 2020. An analysis framework for metric voting based on LP duality. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, 2079–2086.
- [17] Kempe, D. 2020. Communication, Distortion, and Randomness in Metric Voting. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, 2087–2094.
- [18] Liestman, A. L.; and Shermer, T. C. 1993. Additive graph spanners. *Networks*, 23(4): 343–363.
- [19] Mattei, N.; and Walsh, T. 2013. Preflib: A library for preferences <http://www.preflib.org>. In *International Conference on Algorithmic Decision Theory*, 259–270. Springer.
- [20] Munagala, K.; and Wang, K. 2019. Improved metric distortion for deterministic social choice rules. In *Proceedings of the 20th ACM Conference on Economics and Computation (EC)*, 245–262.
- [21] Procaccia, A. D.; and Rosenschein, J. S. 2006. The distortion of cardinal preferences in voting. In *International Workshop on Cooperative Information Agents*, 317–331.
- [22] Seddighin, M.; Latifian, M.; and Ghodsi, M. 2021. On the Distortion Value of Elections with Abstention. *Journal of Artificial Intelligence Research (JAIR)*, 70: 567–595.
- [23] Skowron, P. K.; and Elkind, E. 2017. Social choice under metric preferences: Scoring rules and STV. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, 706–712.
- [24] Walsh, T.; and Xia, L. 2012. Lot-based voting rules. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 603–610.
- [25] Xia, L.; and Conitzer, V. 2008. Generalized scoring rules and the frequency of coalitional manipulability. In *Proceedings of the 9th ACM Conference on Electronic Commerce (EC)*, 109–118.
- [26] Young, H. P. 1975. Social Choice Scoring Functions. *SIAM Journal on Applied Mathematics*, 28(4): 824–838.