Worst-Case Voting When the Stakes Are High

Anson Kahng[1], **Gregory Kehne**[2]
[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) (Anshelevich et al. 2021b). 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 (Boutilier et al. 2015), 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 (Boutilier et al. 2015) 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 (Caragiannis et al. 2017). 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. More concretely, consider a fixed profile of ordinal votes. Multiplicative distortion hedges against bad performance in the case of consistent utilities which assign low total welfare for all candidates, which harms its performance for consistent utilities that yield high-welfare candidates. Additive distortion, on the other hand, prioritizes good performance for this latter case.

In its introduction to the social choice setting, distortion was compared to the distortion of metric embeddings (Procaccia and Rosenschein 2006); this additive distortion is similarly analogous (Liestman and Shermer 1993).

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* (Procaccia and Rosenschein 2006; Caragiannis and Procaccia 2011; Boutilier et al. 2015; Caragiannis et al. 2017; Benadè et al. 2021). 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. However, the assumption that all participating voters' total utility is *equal* is unreasonable in many settings, and we instead uniformly cap the sum of voters' utilities at one (Aziz 2019). 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 focus on a natural class of rules known as *point voting schemes* (Barberà 1978), which are the natural randomized analog of positional scoring rules. A point voting scheme (PVS) 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, PVSs are both intuitive and easy to compute. The two most prominent PVSs—Randomized Dictatorship, and the harmonic rule of (Boutilier et al. 2015)—are nearly distortion-optimal in the normalized utility and metric settings, respectively. When considered together with our results, we argue that PVSs 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 (Boutilier et al. 2015), 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 PVS 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 (2006) in the context of deterministic single-winner social choice functions and normalized utilities. In a later paper, Caragiannis and Procaccia (2011) proved that the Plurality rule has a distortion of $O(m^2)$, and further work demonstrated that this is the best possible distortion of any deterministic voting rule (Caragiannis et al. 2017).

Beyond deterministic social choice functions, Boutilier et al. (2015) 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. (2017) 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 $\frac{1}{2}\left(1 - \frac{1}{m^2}\right) \cdot n$ 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. (2018) 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 (Skowron and Elkind 2017; Goel, Krishnaswamy, and Munagala 2017; Munagala and Wang 2019; Kempe 2020a) before its ultimate proof by Gkatzelis, Halpern, and Shah (2020). Here again randomized rules do better: Anshelevich and Postl (2017) 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 (2020b) and Gkatzelis, Halpern, and Shah (2020) each present rules attaining $3 - 2/m$, and Anshelevich and Postl (2017) and Fain et al. (2019) study variants of the randomized dictatorship mechanism. Lastly, Seddighin, Latifian, and Ghodsi (2021) 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 (Anshelevich et al. 2021b,a).

Finally, we study a class of SCFs which are the randomized analog of positional scoring rules. Young (1975) characterized deterministic scoring functions (with rounds of tiebreaking) as the SCFs which are anonymous, neutral, and consistent, and Xia and Conitzer (2008) provide a striking deterministic generalization of scoring rules. Walsh and Xia (2012) and Bentert and Skowron (2020) 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 $\Sigma := \mathcal{S}_A^n$ 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* $f : \Sigma \to A$ chooses an alternative to be the winner for this profile. Similarly, a randomized social choice function $f : \Sigma \to \Delta_A$ returns a probability distribution over winners, where Δ_A is the probability simplex over A; at election time, a winner is drawn randomly from the probability distribution $f(\sigma) \in \Delta_A$. 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

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

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

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

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 $\mathrm{sw}(a) \geq \alpha \cdot n$.*

2.1 Randomized Scoring Rules

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

Definition 2.3. *A* point voting scheme *(PVS) 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 $\mathcal{PVS}$ denote the class of all such rules.*

Example 2.4. *Consider the PVS 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, $\mathcal{PVS}$ also contains noteworthy rules. One is the harmonic scoring vector-based rule of Boutilier et al. (2015) mentioned above, which is nearly optimal for multiplicative distortion. It is given by $s = (1 + H_m/m, 1/2 + H_m/m, \ldots, 1/m + H_m/m)$, where H_m is the m^{th} harmonic number. Another is Randomized Dictatorship, given by $s = (1, 0, \ldots, 0)$. Remarkably, RD incurs $O(3 - 2/n)$ multiplicative distortion in the metric setting, which is also nearly optimal (Anshelevich and Postl 2017).

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 $P : \mathbb{R}_+^m - \mathbf{0} \to \Delta_A$ a *probabilizer*, and focus on neutral probabilizers, i.e., the P which commute with all permutations of A. Then a *generalized PVS* 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 $\mathcal{PVS}^*$ denote the class of all such SCFs. This is indeed a generalization, since any PVS given by s is a generalized PVS with the probabilizer $P(S)_a := S_a/\|S\|_1$ for all a, where $\|S\|_1 := \sum_{a\in A} S_a$. Note that $\mathcal{PVS}^*$ 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, $\mathcal{PVS}^*$ 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 (Anshelevich and Postl 2017) and the Random Oligarchy mechanism of (Fain et al. 2019).

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. That is, even when voters have uniformly capped (instead of normalized) utilities, the worst case instances for additive distortion are when all voters have utilities summing to 1. The proof (and all other omitted proofs in this paper) can be found in the appendix.

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 $m \geq 3$, $dist^+(f) \geq \frac{5}{18} \cdot n$.*

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

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

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

For deterministic rules, these symmetric instances offer even stronger lower bounds. The following claim was shown by Caragiannis et al. (2017) 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 (Caragiannis et al. 2017))**.** *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 $\text{dist}^+(f) \geq \frac{1}{2} \cdot n$. □

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 (Caragiannis et al. 2017).

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

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

For more than two alternatives, the problem of identifying optimal SCFs or even optimal PVSs 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. (2017) in the general setting of choosing k winners out of m alternatives.

Theorem 3.6 (Theorem 1 in (Caragiannis et al. 2017))**.** *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 (2017), 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 PVSs $(s, P) \in \mathcal{PVS}^*$ with $s = (1, 0, \ldots, 0)$ 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 PVSs which score beyond voters' first choices.

3.3 An Asymptotically Optimal rule in $\mathcal{PVS}$

After the success in Section 3.1, we might hope to derive optimal PVSs for $m \geq 3$ directly. Unfortunately, the natural formulations of finding such optimal PVSs 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:

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

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 $\mathcal{PVS}$ as m becomes large.

We formulate the problem of finding the optimal PVS under eq. (1) in (8) below, and prove that the scoring vector which optimizes this problem is $s^* = (25/33, 7/33, 1/33, 0, \ldots, 0)$ for all $m \geq 3$. Since it is the PVS 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 $m \geq 3$, $\mathrm{dist}^+(BoB) \leq \frac{11}{27} \cdot n$. 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 PVS for all $m \geq 3$.*

We now set about formulating the problem of finding the PVS which minimizes the right-hand side of Equation (1). For a given choice of $\alpha \in [0, 1]$ and scoring vector $s = (s_1, \ldots, s_m)$ 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 $\mathrm{sw}(a^*) = \alpha \cdot n$; we will then consider the worst-case probability that the PVS 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 $\mathrm{sw}(a^*) = n \cdot \sum_i \frac{1}{i} x_i$.

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

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

$$s.t. \qquad \sum_i \frac{1}{i} x_i = \alpha, \qquad x \in \Delta_{[m]}. \qquad (3)$$

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

$$D^+(s) := \max\ D^+(s, \alpha) \qquad (4)$$

$$s.t. \qquad 0 \leq \alpha \leq 1, \qquad (5)$$

where eq. (5) captures that $\mathrm{sw}(a^*) \leq n$, since each voter's utilities are normalized to 1.

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

$$s^* := \arg\min D^+(s) \qquad (6)$$

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

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

$$D^+(s) := \max \left(\sum_i \frac{1}{i} x_i\right) \left(\sum_i (1 - s_i) x_i\right) \qquad (7)$$

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

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

$$\mathcal{D}^+ := \min_s \max_x \left(\sum_i \frac{1}{i} x_i\right) \left(\sum_i (1 - s_i) x_i\right) \qquad (8)$$

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

Two Alternatives and the Harmonic PVS As noted in Observation 3.5, the optimal SCF when $m = 2$ is the PVS given by $s = (1, 0)$. On the other hand, for $m = 2$ the optimal scoring vector for the formulation (8) 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 PVS 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 (8), as we show in the next section. Indeed the harmonic PVS 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 (8). However it turns out that the same scoring vector s^* which (together with a suffix of trailing zeros) optimizes (8) for all $m \geq 3$ simultaneously.

Lemma 3.10. *For all $m \geq 3$, the unique optimal solution to* (8) *is the scoring vector $s^* = (25/33, 7/33, 1/33, 0, \ldots, 0)$, obtaining the optimum objective value $\frac{11}{27} \approx 0.407$.*

This candidate optimizer s^* of (8) 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 (7) to a new problem $\overline{D}^+(s)$; this gives a corresponding relaxation of the outer problem (8). 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 (7):

Lemma 3.11. *For $m = 3$, the unique optimal solution to* (8) *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 (8) 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 PVS derived in Section 3.3 is asymptotically optimal within $\mathcal{PVS}$, 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

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 ∈ A do
    w_a^b ← Σ_i (σ_i^{-1})^{-1} 𝟙{b ≻_i a}
end for
w_a ← (w_a^b)_{b∈A} for each a ∈ A
p* ← arg min_p {D : w_a^a − p^T w_a ≤ D ∀a ∈ A, p ∈ Δ_A}
return p*
```

of $\text{dist}^+(f, \sigma)$ over all SCFs f (which for fixed σ are probability distributions over A). In particular,

$$\begin{aligned}\text{AddOpt}(\sigma) &:= \arg\min_f \text{dist}^+(f, \sigma) \\ &= \min_{p \in \Delta_A} \max_{u \triangleright \sigma} \left(\max_{a^* \in A} \text{sw}(a^*) - \mathbb{E}_{a \sim p}[\text{sw}(a)] \right).\end{aligned}$$

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.*

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 $u, u' \triangleright \sigma$ 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*

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

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

In the promise setting, we might also hope to circumvent the relatively low-welfare $\Omega(\sqrt{m})$ lower bound given in (Boutilier et al. 2015). Indeed, the lower bound instance in (Boutilier et al. 2015) translates directly into a lower bound on distortion with an α-promise:

Theorem 4.3. *For any randomized SCF f,*

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

A slight modification of the *Stable Lottery Rule* f_{SLR} introduced by Ebadian et al. (2022) yields a matching upper bound for all $\alpha \ge 1/\sqrt{m}$. In particular, the modified rule samples alternatives from the stable lotteries of Cheng et al. (2020), which are distributions over committees of size $2/\alpha$.

Theorem 4.4. *There is an SCF ℓ_α with $\text{dist}_\alpha(\ell_\alpha) = O\left(\frac{1}{\alpha}\right)$.*

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

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

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

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

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

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

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

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

Plurality and the Harmonic Rule are robust for $\alpha \ge 3/4$, 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,*

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

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

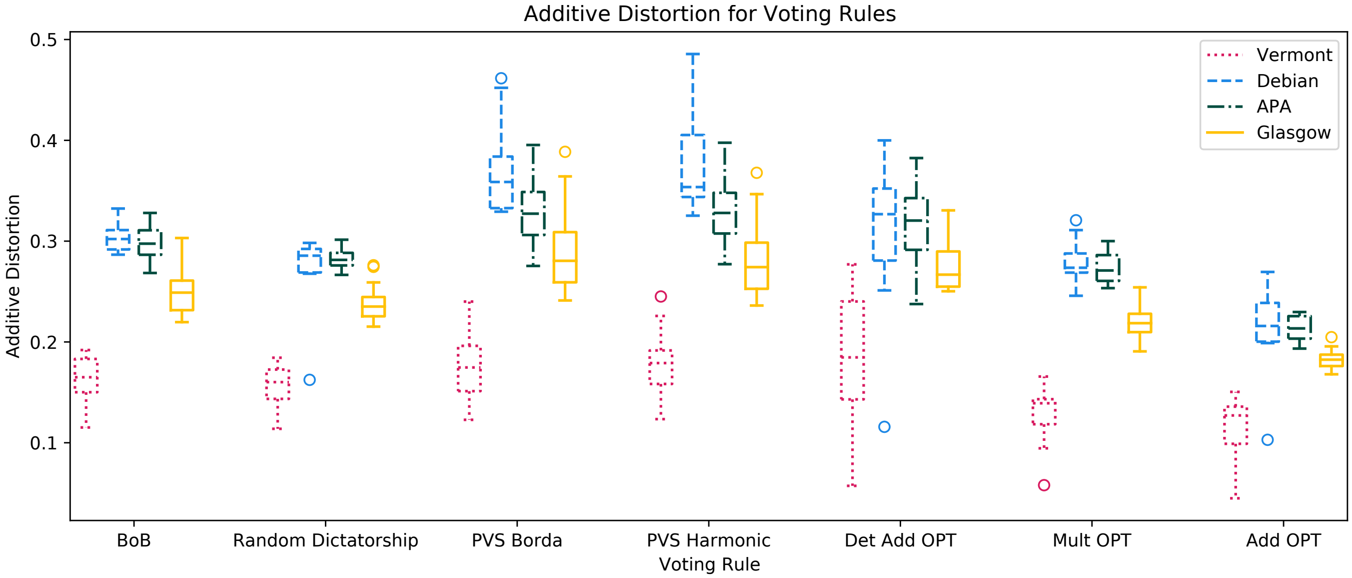

Figure 1: Additive distortion of voting rules on the Vermont, Glasgow, Debian, and APA datasets, normalized by n.

5 Experiments

We evaluated the performance of various SCFs on four datasets of election data from PrefLib (Mattei and Walsh 2013): *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); *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 $s = (1, 0, \dots)$ (Abdulkadiroğlu and Sönmez 1998); *PVS Borda* has score vector $s = (m-1, m-2, \dots, 0)$; *PVS 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 (Caragiannis et al. 2017); *Mult OPT* is the randomized rule that minimizes multiplicative distortion (Boutilier et al. 2015); 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 PVSs (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, PVS Borda, and PVS Harmonic. We find that RD consistently outperforms BoB on the four datasets, while PVS Borda and PVS Harmonic both do worse. This is surprising, since Theorem 3.9 demonstrates that BoB is asymptotically worst-case optimal among the class of all PVSs. 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. (2017) 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. 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 $\mathcal{PVS}^*$, 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.

Acknowledgments

The authors would like to thank Alex Wang for valuable discussions.

References

Abdulkadiroğlu, A.; and Sönmez, T. 1998. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica*, 66(3): 689–701.

Anshelevich, E.; Bhardwaj, O.; Elkind, E.; Postl, J.; and Skowron, P. 2018. Approximating optimal social choice under metric preferences. *Artificial Intelligence*, 264: 27–51.

Anshelevich, E.; Filos-Ratsikas, A.; Shah, N.; and Voudouris, A. A. 2021a. Distortion in social choice problems: an annotated reading list. *ACM SIGecom Exchanges*, 19(1): 12–14.

Anshelevich, E.; Filos-Ratsikas, A.; Shah, N.; and Voudouris, A. A. 2021b. Distortion in Social Choice Problems: The First 15 Years and Beyond. arXiv:2103.00911.

Anshelevich, E.; and Postl, J. 2017. Randomized social choice functions under metric preferences. *Journal of Artificial Intelligence Research (JAIR)*, 58: 797–827.

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

Barberà, S. 1978. Nice decision schemes. In *Decision Theory and Social Ethics*, 101–117.

Benadè, G.; Nath, S.; Procaccia, A. D.; and Shah, N. 2021. Preference Elicitation for Participatory Budgeting. *Management Science*, 67(5): 2813–2827.

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.

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.

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.

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

Cheng, Y.; Jiang, Z.; Munagala, K.; and Wang, K. 2020. Group Fairness in Committee Selection. *ACM Trans. Economics and Comput.*, 8(4): 23:1–23:18.

Ebadian, S.; Kahng, A.; Peters, D.; and Shah, N. 2022. Optimized Distortion and Proportional Fairness in Voting. https://www.cs.toronto.edu/~nisarg/papers/distortion.pdf. Accessed: 2022-02-22.

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.

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.

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.

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.

Kempe, D. 2020a. An analysis framework for metric voting based on LP duality. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, 2079–2086.

Kempe, D. 2020b. Communication, Distortion, and Randomness in Metric Voting. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, 2087–2094.

Liestman, A. L.; and Shermer, T. C. 1993. Additive graph spanners. *Networks*, 23(4): 343–363.

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.

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.

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

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.

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.

Walsh, T.; and Xia, L. 2012. Lot-based voting rules. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 603–610.

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.

Young, H. P. 1975. Social Choice Scoring Functions. *SIAM Journal on Applied Mathematics*, 28(4): 824–838.

A Omitted Claim from Section 1

In Section 1 we claim without proof that, under the unit-capped (as contrasted with the unit-sum) assumption on agents' implicit utilities, the multiplicative distortion of choosing an alternative uniformly at random is m, and that this is tight. Letting UAR (for uniform at random) denote this rule, we provide a proof of this claim here:

Claim A.1. $\mathrm{dist}(UAR) = m$ *under the unit-capped implicit utility assumption, and this is optimal among all randomized SCFs.*

Proof. We begin by showing that $\mathrm{dist}(UAR) \leq m$ in this setting; we then show that no randomized SCF has distortion less than m. Given utility profile (u, σ), let a^* be the social-welfare-maximizing alternative. Then

$$\mathbb{E}_{a \sim f(\sigma)} = \frac{1}{m} \mathrm{sw}(a^*) + \frac{1}{m} \sum_{a \neq a^*} \mathrm{sw}(a) \leq \frac{1}{m} \mathrm{sw}(a^*).$$

this in turn implies that

$$\frac{\mathrm{sw}(a^*)}{\mathbb{E}_{a \sim f(\sigma)}} \geq \frac{\mathrm{sw}(a^*)}{\frac{1}{m} \mathrm{sw}(a^*)} = \frac{1}{m},$$

and so

$$\mathrm{dist}(UAR) = \max_{\sigma} \max_{u \triangleright \sigma} \frac{\mathrm{sw}(a^*)}{\mathbb{E}_{a \sim f(\sigma)}} \leq m.$$

On the other hand, fix a randomized SCF f and consider the symmetric profile on $m = n$ alternatives; by the pigeonhole principle there is some alternative a^* for which $\Pr[f(\sigma) = a^*] \leq \frac{1}{m}$. Next suppose that the first agent values $u_1(a^*) = 1$ and $u_1(a) = 0$ otherwise, and all other agents have no utility for all alternatives: $u_i(a) = 0$ for all $i \neq 1$ and all a. Then $\mathbb{E}_{a \sim f(\sigma}[\mathrm{sw}(a)] \leq \frac{1}{m}$ while $\mathrm{sw}(a^*) = 1$, and so $\mathrm{dist}(f) \geq m$. □

B Omitted Proofs from Section 3

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

Proof. We begin by decomposing additive distortion.

$$\begin{aligned}\mathrm{dist}^+(f) &= \max_{\sigma} \mathrm{dist}^+(f, \sigma) \\ &= \max_{\sigma} \max_{a^* \in A} \mathrm{dist}^+(f, \sigma, a^*) \\ &= \max_{\sigma} \max_{a^* \in A} \sum_{i \in [n]} dist_i^+(f, \sigma, a^*),\end{aligned}$$

where

$$\mathrm{dist}^+(f, \sigma, a^*) := \max_{u \triangleright \sigma} \left[\mathrm{sw}(a^*) - \sum_{a \in A} \Pr[f(\sigma) = a]\ \mathrm{sw}(a) \right]$$

denotes the worst-case distortion given a profile σ and a randomized SCF f with respect to alternative a^* and

$$\mathrm{dist}_i^+(f, \sigma, a^*) := \max_{u_i \triangleright \sigma_i} \left[u_i(a^*) - \sum_{a \in A} \Pr[f(\sigma) = a]\ u_i(a) \right]$$

denotes the contribution of voter i toward $\mathrm{dist}^+(f, \sigma, a^*)$.

We now show that the utilities that maximize $\mathrm{dist}_i^+(f, \sigma, a^*)$ for every voter i are of the following form. Let k_i be the position of a^* in σ_i. Then, let $u_i(a) = 1/k_i$ if $\sigma_i(a) \leq k_i$ and $u_i(a) = 0$ otherwise. Note that this utility profile satisfies $\sum_a u_i(a) = 1$.

Now, let $f(a)$ represent the probability that f chooses alternative a. Furthermore, let k be the position of a^* in σ_i, and define $q := \sum_{a : a \succ_i a^*} f(a)$, $s := f(a^*)$, and $r := \sum_{a : a^* \succ_i a} f(a)$. In other words, q represents the total probability mass assigned by f to alternatives that appear before a^* in σ_i, r represents the probability assigned by f to a^*, and s represents the probability assigned by f to alternatives that appear after a^* in σ_i. By definition, we have that $q + r + s = 1$. Now, note that we can denote the contribution of voter i to the total additive distortion as

$$\mathrm{dist}_i^+(f, \sigma, a^*) := \max_{u_i \triangleright \sigma_i} \left[u_i(a^*) - \sum_{a \in A} \Pr[f(\sigma) = a]\ u_i(a) \right]$$

$$= \max_{u_i \triangleright \sigma_i} \left[q \left(u_i(a^*) - (1/q) \sum_{a: a \succ_i a^*} u_i(a) f(a) \right) + r \left(u_i(a^*) - (1/r) u_i(a^*) f(a^*) \right) + s \left(u_i(a^*) - (1/s) \sum_{a: a^* \succ_i a} u_i(a) f(a) \right) \right]$$

$$= \max_{u_i \triangleright \sigma_i} \left[q \left(u_i(a^*) - (1/q) \sum_{a: a \succ_i a^*} u_i(a) f(a) \right) + s \left(u_i(a^*) - (1/s) \sum_{a: a^* \succ_i a} u_i(a) f(a) \right) \right]$$

because $f(a^*) = r$ by definition.

We will now argue that our choice of u_i maximizes this expression. For the first term, note that $u_i(a^*) - (1/q) \sum_{a: a \succ_i a^*} u_i(a) f(a) \leq 0$ because $(1/q) \sum_{a: a \succ_i a^*} u_i(a) f(a)$ captures the average utility of alternatives that are preferred to a^* in σ_i. Our choice of utilities results in $u_i(a^*) - (1/q) \sum_{a: a \succ_i a^*} u_i(a) f(a) = 0$, which is the maximum value this term can achieve. As for the second term, note that our choice of utilities independently maximizes $u_i(a^*)$ subject to utility constraints imposed by σ_i and simultaneously minimizes $\sum_{a: a^* \succ_i a} u_i(a) f(a)$. Because total additive distortion is additive over voters and our choice of utilities maximizes each voter's individual contribution toward additive distortion, our choice of utilities is maximizes additive distortion subject to the constraint that a^* is the true best alternative. □

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.*

Proof. If there are two alternatives then the profile σ may be parameterized by a single variable λ which denotes the proportion of $[n]$ ranking alternative a_1 before alternative a_2. Therefore in this setting each randomized social choice function is given by some $p : [0, 1] \to [0, 1]$, where $p(\lambda)$ is the probability that the rule chooses a_1 given λ.

To begin, let us fix λ and derive the optimal p. For fixed λ and p, and observing that by assumption 2.1 $\mathrm{sw}(a_2) = n - \mathrm{sw}(a_1)$, the additive distortion is given by

$$\begin{aligned} \mathrm{dist}^+(\lambda, p) &= \max_{u \triangleright \lambda} \ (\max(\mathrm{sw}(a_1), n - \mathrm{sw}(a_1)) \\ &\qquad -(p \cdot \mathrm{sw}(a_1) + (1-p) \cdot (n - \mathrm{sw}(a_1)))) \\ &= n \cdot \max_{u \triangleright \lambda} \ \big(\max(w,\ 1-w) \\ &\qquad - (p \cdot w + (1-p) \cdot (1-w)) \big), \end{aligned}$$

where $w := \mathrm{sw}(a_1)/n$ for convenience.

There are two cases to consider: when a_1 is the utility maximizer, and when a_2 is. Letting d_1^+ and d_2^+ denote the distortions in these cases, and simplifying, we have

$$\begin{aligned} d_1^+(\lambda, p) &= n \cdot \max_{u \triangleright \lambda} \ \big((1-p)(2w-1) \big), \\ d_2^+(\lambda, p) &= n \cdot \max_{u \triangleright \lambda} \ \big(p(1-2w) \big). \end{aligned}$$

In the first case, the u maximizing the expression for a given λ puts maximal utility on a_1, and so $w = (\lambda+1)/2$. In the second, the maximizing u puts maximal weight on a_2, and so $w = \lambda/2$. Therefore

$$\begin{aligned} d_1^+(\lambda, p) &= n \cdot (1-p)\lambda, \\ d_2^+(\lambda, p) &= n \cdot p(1-\lambda), \end{aligned}$$

and so

$$\begin{aligned} d^+(\lambda, p) &= \max(d_1^+(\lambda, p), \\ d_2^+(\lambda, p)) &= n \cdot \max((1-p)\lambda,\ p(1-\lambda)). \end{aligned}$$

Since the first term is monotonically decreasing in p and the second is monotonically increasing, this is minimized when they are equal, giving $p = \lambda$. □

Theorem 3.6 (Theorem 1 in (Caragiannis et al. 2017))**.** *Plurality is an optimal deterministic SCF, with additive distortion $\frac{1}{2} \cdot n$.*

Proof. Claim 3.3 establishes a lower bound of $\frac{1}{2} \cdot n$ for all deterministic SCFs, so it remains to show that the additive distortion incurred by Plurality is never worse.

For a given instance let a^* be the social welfare maximizing alternative, and suppose without loss of generality that Plurality chooses $a \neq a^*$. There are three types of voters to consider: voters ranking a first, voters ranking a^* first, and voters ranking something else first. Denote these voter types A, B, and C, and suppose that an α, β, γ proportion of the voters are of each type, respectively. Note that $\alpha + \beta + \gamma = 1$ and that $\alpha \geq \beta$, since Plurality chose a over a^*. Also observe that $u_A(a^*) \leq u_A(a)$

for all voters of type A, since they prefer a to a^*, and that $u_C(a^*) \leq 1/2$ for similar reasons. Therefore the contributions to the additive distortion from each group are

$$\begin{aligned} d_A^+ &= \alpha\, n \cdot \max_{u_A} [u_A(a^*) - u_A(a)] = 0 \\ d_B^+ &= \beta\, n \cdot \max_{u_B} [u_B(a^*) - u_B(a)] = \beta\, n \\ d_C^+ &\leq \gamma\, n \cdot \max_{u_C} [u_C(a^*) - u_C(a)] = \frac{1}{2}\gamma\, n, \end{aligned}$$

with inequality for d_C^+ because a^* may not be ranked second for all type C voters. These together yield

$$\text{dist}^+(\mathit{Plurality}) = d_A^+ + d_B^+ + d_C^+ \leq (\beta + \gamma/2)n \leq \frac{1}{2}n,$$

since $\beta \leq \alpha$. □

Proof of Theorem 3.7. To start, let β_a be the proportion of voters ranking a first, and let a^* be the social welfare maximizing alternative. Note that RD chooses an alternative according to the distribution β. Since additive distortion is linear in the voters, it is maximized when it is maximized for all voters individually, and each voter i contributes $u_i(a^*) - \mathbb{E}_{a\sim\beta}[u_i(a)]$.

For an voter ranking a^* first, we may therefore assume that they value $u(a^*) = 1$ and all other alternatives confer no utility. For voters i ranking $a \neq a^*$ first, their contribution to the distortion may be rewritten as

$$u(a^*)(1 - \beta_{a^*}) - \beta_a u_i(a) - \sum_{a' \neq a, a^*} \beta_{a'} u_i(a').$$

Given that this voter values a above a^*, the ranking and utility vector which simultaneously maximizes this first term and minimizes the others assigns $u_i(a) = u_i(a^*) = 1/2$, and $u_i(a') = 0$ for all other a'. Combining these assumptions and simplifying, we obtain

$$\begin{aligned} \text{dist}^+(RD) &= \max_{\beta} \left(\text{sw}(a^*) - \sum_a n \cdot \beta_a \ \text{sw}(a) \right) \\ &= \max_{\beta}\ n \cdot \left(\frac{1}{2} \left(1 - \sum_a \beta_a^2 \right) \right) = \frac{1}{2}\left(1 - \frac{1}{m}\right) \cdot n \end{aligned}$$

where the last inequality follows from the $\ell_1 - \ell_2$ inequality, which implies that the maximum is obtained when $\beta_a = 1/m$ for all $a \in A$. □

Proof of Claim 3.8. Consider the tight instance given in the proof of theorem 3.7, where an equal number of voters rank each alternative first, and all voters not ranking a^* first rank a^* second. Therefore the scores derived from $s = (1, 0, \ldots, 0)$ are equal; $S_a = S_{a'}$ for all $a, a' \in A$. (Note that non-neutral "favorite-only" mechanisms do only worse than neutral mechanisms on this instance). Given this uniform aggregate score vector S, every neutral mechanism f returns the uniform distribution over A. This bounds its additive distortion as

$$\text{dist}^+(f) \geq \frac{m+1}{2m} - \frac{1}{2m}\left(\frac{m-1}{m} + \frac{m+1}{m}\right) = \frac{1}{2}\left(1 - \frac{1}{m}\right) \cdot n,$$

as desired. □

Lemma 3.11. *For $m = 3$, the unique optimal solution to* (8) *is the scoring vector $s^* = (25/33, 7/33, 1/33)$, obtaining objective value $\frac{11}{27}$.*

Proof of Lemma 3.11. Consider the restricted inner problem $\overline{D}^+(s)$ given by

$$\begin{aligned} \overline{D}^+(s) := \max &\left(\sum_i \frac{1}{i} x_i \right) \left(\sum_i (1 - s_i) x_i \right) \\ s.t. \qquad x &= \begin{cases} (\beta, 1-\beta, 0) & \text{for } \beta \in [0,1] \quad \text{or} \\ (\gamma, 0, 1-\gamma) & \text{for } \gamma \in [0,1] \end{cases} \end{aligned} \tag{9}$$

this is a restriction of (7) since the x which are feasible for (9) are also feasible for (7).

Writing $s = (s_1, s_2, 1 - s_1 - s_2)$, the objective of (9) for each form of x is

$$\max_{\beta\in[0,1]} \frac{1}{2}(\beta+1)\left(1-(s_1-s_2)\beta - s_2\right), \tag{10}$$

$$\max_{\gamma\in[0,1]} \frac{1}{3}(2\gamma+1)\left(s_1+s_2+\gamma(1-2s_1-s_2)\right). \tag{11}$$

If $\overline{D}^+(s) \leq 11/27$ then both (10) $\leq 11/27$ and (11) $\leq 11/27$. Taking derivatives, we find that the β which maximizes the argument of (10) in terms of s_1 and s_2 is either $\beta = 0$, $\beta = 1$, or $\beta = \frac{1-s_1}{2(s_1-s_2)}$ (provided that $0 \leq \frac{1-s_1}{2(s_1-s_2)} \leq 1$). Requiring that $\overline{D}^+(s) \leq 11/27$ in each of these cases yields the inequalities

$$1 - s_1 \leq 11/27, \tag{12}$$

$$\frac{1-s_2}{2} \leq 11/27, \tag{13}$$

$$\frac{(1+s_1-2s_2)^2}{8(s_1-s_2)} \leq 11/27. \tag{14}$$

These first two constraints imply that $\frac{1-s_1}{2(s_1-s_2)} \geq 0$; we therefore require (14) when $\frac{1-s_1}{2(s_1-s_2)} \leq 1$.

The second partial formulation (11) for the critical values $\gamma = 0, \gamma = 1$, and $\gamma = \frac{1+s_2}{4(2s_1+s_2-1)}$ similarly yields the inequalities

$$\frac{s_1+s_2}{3} \leq 11/27 \tag{15}$$

$$\frac{(4s_1+3s_2-1)^2}{24(2s_1+s_2-1)} \leq 11/27 \tag{16}$$

as well as (12) again. The inequality (15) is trivial, since we assume that $w \geq 0$ and $\sum_i s_i = 1$. We again require (16) so long as $0 \leq \frac{1+s_2}{4(2s_1+s_2-1)} \leq 1$; fortunately $0 \leq \frac{1+s_2}{4(2s_1+s_2-1)}$ follows from (12), and $\frac{1+s_2}{4(2s_1+s_2-1)} \leq 1$ follows from (12) and (13) together.

There are now two cases to consider. First, we argue that the inequalities $\frac{1-s_1}{2(s_1-s_2)} \geq 1$ and (16) are not simultaneously satisfiable; we may therefore assume that $\frac{1-s_1}{2(s_1-s_2)} \leq 1$ and consequently that (14) holds. The first inequality is the half-plane

$$3s_1 - 2s_2 \leq 1$$

(since (12) implies that $s_1 > s_2$), and so in order to argue incompatibility it suffices to consider the optimization problem, where the constraint is a rearrangement of (16):

$$\begin{aligned} &\min_s \ 3s_1 - 2s_2 \\ s.t. \qquad &(4s_1+3s_2-1)^2 - \frac{88}{9}(2s_1+s_2-1) \leq 0. \end{aligned}$$

A minimum objective value greater than one will demonstrate unsatisfiability. By the change of variable $x = 4s_1 + 3s_2 - 1$ and $y = 3s_1 - 2s_2$, this becomes

$$\begin{aligned} &\min \ y \\ s.t. \qquad &\frac{153}{176}x^2 - \frac{7}{2}x + 5 \leq y, \end{aligned}$$

which has a minimum value of $226/153 > 1$.

We have shown that (14) and (16) must hold if the restricted problem (9) is to have a solution with objective at most $11/27$. We finally argue that s^* is the unique s satisfying both (14) and (16). As above, (12) implies that the denominators of (14) and (16) are positive, and we may therefore rewrite them as

$$(1+s_1-2s_2)^2 - \frac{88}{27}(s_1-s_2) \leq 0 \tag{17}$$

$$(4s_1+3s_2-1)^2 - \frac{88}{9}(2s_1+s_2-1) \leq 0. \tag{18}$$

We will show that (17) and (18) are simultaneously satisfied at the single point $(s_1^*, s_2^*) = (25/33, 7/33)$ by way of the line ℓ given by $2s_1 + 7s_2 = 3$, which separates their feasible regions. Through the changes of variables $r = -s_1 + 2s_2$, $t = 2s_1 + s_2$, and $u = 4s_1 + 3s_2$, $v = -3s_1 + 4s_2$, (17) and (18) become

$$\frac{3}{88}(45 - 2r + 45r^2) \leq t \tag{19}$$

$$\frac{1}{176}(-2425 + 1418u - 225u^2) \geq v, \tag{20}$$

while ℓ takes the form $12r + 11t = 15$ and $29u + 22v = 75$.

First, if (17) and $2s_1 + 7s_2 \leq 3$ then equivalently (19) and $t \leq 15/11 - 12/11 \cdot s$. These imply that

$$\begin{gathered}
\frac{3}{88}(45 - 2r + 45r^2) \leq 15/11 - 12/11 \cdot r \\
\frac{15}{88}(1 + 3r)^2 \leq 0 \\
r = -1/3 \qquad t = 19/11 \\
s_1 = 25/33 \qquad s_2 = 7/33.
\end{gathered}$$

Similarly, if (18) and $2s_1 + 7s_2 \geq 3$ then equivalently (20) and $v \geq 75/22 - 29/22 \cdot u$. These imply that

$$\begin{gathered}
\frac{1}{176}(-2425 + 1418u - 225u^2) \geq 75/22 - 29/22 \cdot u \\
-\frac{25}{176}(11 - 3u)^2 \geq 0 \\
u = 11/3 \qquad v = -47/33 \\
s_1 = 25/33 \qquad s_2 = 7/33.
\end{gathered}$$

This demonstrates that the feasible regions of (17) and (18) are contained within each of the two half-planes bordered by ℓ. Since both (17) and (18) intersect with ℓ only at s^*, we conclude that s^* is the unique point satisfying both (17) and (18).

Given that s^* is the unique minimizer of (9), it remains to argue that it is the unique optimizer of (8) for $m = 3$. Writing $x = (x_1, x_2, 1 - x_1 - x_2)$, the objective of $D^+(s^*)$ takes the explicit form

$$\begin{aligned}
D^+(s^*) &= \max_x \ \frac{1}{99}(2 + 4x_1 + x_2)(16 - 12x_1 - 3x_2) \\
&= \max_x \ \frac{11}{27} - \frac{1}{33}\left(4x_1 + x_2 - \frac{5}{3}\right)^2,
\end{aligned}$$

which demonstrates that no x obtains a larger objective on s^* even when we relax the inner problem from (9) to (7) when $m = 3$. □

Lemma 3.10. *For all* $m \geq 3$*, the unique optimal solution to* (8) *is the scoring vector* $s^* = (25/33, 7/33, 1/33, 0, \ldots, 0)$*, obtaining the optimum objective value* $\frac{11}{27} \approx 0.407$.

Proof. Given the scoring vector $s^* = (25/33, 7/33, 1/33, 0, \ldots, 0)$, we argue that $D^+(s^*) \leq 11/27$ for all $m > 3$. Recall that the objective of (7) is

$$\max_x \left(\sum_i \frac{1}{i} x_i\right)\left(\sum_i (1 - s_i) x_i\right).$$

For any x with $x_i > 0$ for any $i \geq 4$, the vector $x' := (x_1, x_2, x_3, x'_4, 0, \ldots, 0)$ for $x'_4 := \sum_{i \geq 4} x_i$ will only increase this objective, since the first factor in this product increases from x to x', and the second factor remains unchanged. Therefore we may assume without loss of generality that $m = 4$.

From here, we next show that we may in fact restrict our attention to x of the form $x = (x_1, x_2, x_3, 0, \ldots, 0)$. We will take a somewhat convoluted approach, which nevertheless seems to be the best available. To see this, consider the explicit form which the objective of (8) takes for s^* and $m = 4$. Substituting $x_4 = 1 - (x_1, +x_2 + x_3)$, it takes the form

$$O_4(x) := -\frac{1}{396}(3 + 9x_1 + 3x_2 + x_3)(-33 + 25x_1 + 7x_2 + x_3).$$

We may ask what $x \in \Delta_4$ maximizes this objective, which may be formulated as the following program:

$$\begin{aligned}
&\max \ O_4(x) && (21) \\
s.t. \quad & x_1 + x_2 + x_3 \leq 1 \quad (\lambda) && (22) \\
& x_1, x_2, x_3 \geq 0. \quad (\mu_i) && (23)
\end{aligned}$$

We would like to establish that $x_4 = 0$, which is equivalent to constraint (22) being tight at optimality.

Since any active constraints are independent and linear, any x^* maximizing (21) must satisfy the KKT conditions (note that it must obtain a maximum, since it is a continuous function on a compact set). The KKT conditions of (21) are the linear inequalities

$$-\frac{\partial O_4(x)}{\partial x_i} - \mu_i + \lambda = 0 \qquad i \in \{1, 2, 3\}$$

for $\mu_i \geq 0$, with strict equality for constraint i if $x_i^* > 0$. If (22) is loose (i.e. $x_4 > 0$) then $\lambda = 0$ by complementary slackness; since $\mu_i \geq 0$, this becomes

$$\begin{aligned} \frac{1}{198}\left(111 - 225x_1 - 69x_2 - 17x_3\right) &\leq 0 \qquad (C_1)\\ \frac{1}{198}\left(39 - 69x_1 - 21x_2 - 5x_3\right) &\leq 0 \qquad (C_2)\\ \frac{1}{198}\left(15 - 17x_1 - 5x_2 - xx_3\right) &\leq 0 \qquad (C_3), \end{aligned}$$

with constraint (C_i) tight if $x_i > 0$ at optimality. By checking the (x_1, x_2, x_3) which are feasible for both (21) and these linear constraints, we may determine that the only candidate optima of (21) for which $x_4 > 0$ have $x_1 = x_2 = 0$ or $x_1 = x_3 = 0$ or $x_2 = x_3 = 0$. Calculating the optima for these single-variable cases, we may confirm that all feasible candidate optima have objective at most $1/4 < 11/27$. Since we know that $11/27$ is attainable for this problem, it follows that we may assume that $x_4 = 0$ at optimality.

By showing that without loss of generality the x maximizing the inner problem (7) are of the form $x = (x_1, x_2, x_3, 0, \ldots, 0)$, we now know by Lemma 3.11 that s^* attains an objective of at most $11/27$ for (8) for any $m \geq 4$.

To confirm that s^* is indeed the optimal solution for all $m \geq 4$, we must finally argue that for any given $m \geq 4$ no s can do better—i.e. attain a strictly smaller value of $D^+(s)$. Informally this is because for any s obtaining some objective d on (8), its prefixes must obtain d or better on the lower-dimensional problems. This is because the maximizing x may place all of its weight on some prefix of the coordinates. Formally, suppose that s is an optimal solution to (8) for some $m \geq 4$. Let $s_{:3} := (s_1, s_2, s_3, 0, \ldots, 0)$, and let $\overline{s_{:3}}$ denote $s_{:3}/\|s_{:3}\|_1$ (and if $\|s_{:3}\|_1 = 0$ then s does truly terribly, and isn't worth worrying about). It is clear that for fixed $x = (x_1, x_2, x_3, 0, \ldots, 0)$ the objective of (8) is weakly greater at $s_{:3}$ than it is at $\overline{s_{:3}}$; since x is feasible to begin with and $\overline{s_{:3}}$ is feasible for (8) when $m = 3$, it then follows from Lemma 3.11 that the objective value of (8) at s is at least $11/27$. □

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 PVS for all $m \geq 3$.*

Proof of Theorem 3.9. Let f^* denote this PVS with score s^*. We note first that f^* has additive distortion $d^+(f^*) \leq \frac{11}{27} \cdot n$ for all $m \geq 3$, by the best-or-bust bound eq. (1). The additive distortion of a rule f is given by

$$\text{dist}^+(f) = \max_{\sigma} \max_{u \triangleright \sigma} \left(\text{sw}(a^*) - \sum_{a \in A} \Pr[f(\sigma) = a]\ \text{sw}(a)\right),$$

while we have argued above that the objective of (7) takes the form

$$D^+(s) = \max_{\sigma} \max_{u \triangleright \sigma} \left(\text{sw}(a^*) - \Pr[f^s(\sigma) = a^*]\ \text{sw}(a^*)\right) \tag{24}$$

for f^s an PVS, and where a^* denotes the social-welfare-maximizing alternative given u. Therefore it is clear that for any fixed PVS f^s the objective (24) is larger than the additive distortion of the rule. In particular, this is also true for the rule f^* which minimizes (24) over all s.

Next we argue that all PVSs f^s satisfy $\text{dist}^+(f^s) \geq \left(1 - \frac{16}{27}\frac{1}{m-1}\right)\frac{11}{27} \cdot n$. To see this, fix m and choose any scoring vector s with corresponding rule f^s. For every (rational) x as in (7), fix some alternative a^* and consider a profile (u, σ) for which

- an x_i proportion of voters rank a^* in position i,
- for each i, an equal number of voters who rank a^* in position i rank each alternative $a \neq a^*$ in each position $i' \neq i$, and
- u maximizes $\text{sw}(a^*)$ subject to consistency with σ.

Since utilities may tie, this third condition is well-defined and determines all utilities. It also guarantees that a^* is a (not necessarily unique) maximizer of social welfare for (u, σ). Finally, for the purposes of additive distortion it is without loss of generality to assume u of this form given a^* and σ, since increasing $\text{sw}(a^*)$ for fixed σ only increases the maximized inner term $\text{sw}(a^*) - \mathbb{E}_{a \sim f(\sigma)}[\text{sw}(a)]$.

Let $\Pr[a] := \Pr[f^s(\sigma) = a]$ for convenience; then for each x we have

$$d^+(f^s, \sigma) \geq \mathrm{sw}(a^*) - \sum_a \Pr[a]\ \mathrm{sw}(a),$$

and regrouping and taking the maximum over all profiles σ derived from x,

$$\begin{aligned} d^+(f^s) &\geq \max_x \left(\mathrm{sw}(a^*)(1 - \Pr[a^*]) - \sum_{a \neq a^*} \Pr[a]\ \mathrm{sw}(a) \right) \\ &= \max_x \left(\mathrm{sw}(a^*)(1 - \Pr[a^*]) - (1 - \Pr[a^*]) \frac{n - \mathrm{sw}(a^*)}{m-1} \right) \qquad (25) \\ &= \max_x \left(\mathrm{sw}(a^*)(1 - \Pr[a^*]) \left(1 + \frac{1}{m-1} \right) - (1 - \Pr[a^*]) \frac{n}{m-1} \right), \qquad (26) \end{aligned}$$

where (25) follows from the construction of σ and u given x, so that the probability of $f^s(\sigma) = a$ is equal for all $a \neq a^*$, together with the fact that the total utility is n. This first term is precisely the objective of (8), and so by Lemma 3.10 we have

$$\begin{aligned} &\geq \max_x \left(\left(1 + \frac{1}{m-1} \right) \frac{11}{27} \cdot n - (1 - \Pr[a^*]) \frac{n}{m-1} \right) \\ &\geq \left(1 + \frac{1}{m-1} \right) \frac{11}{27} \cdot n - \frac{n}{m-1} \\ &= \left(1 - \frac{16}{27} \frac{1}{m-1} \right) \frac{11}{27} \cdot n. \end{aligned}$$

Since this last term is an upper bound on the additive distortion of the rule for s^*, and since this holds for all positional scoring rules, we may take the right-hand side to be the optimal PSR for m, and then invert to obtain that f^* is a $\left(1 + \frac{1}{m-1}\right)$-approximation to the optimal PVS, for all $m \geq 3$.

□

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

Proof. We first leverage the approach in Lemma 3.1 in order to drastically simplify the optimization problem by restricting the space of worst-case utilities we must consider.

The worst-case additive distortion on a profile σ can be written as

$$\begin{aligned} \mathrm{dist}^+(\sigma) &= \min_f \mathrm{dist}^+(f, \sigma) \\ &= \min_f \max_{a^* \in A} \mathrm{dist}^+(f, \sigma, a^*). \end{aligned}$$

In other words, this can be then decomposed into independently finding the best f for m different optimization problems each corresponding to the case that a particular $a^* \in A$ is the true best alternative and then taking the minimum over these f solutions. However, we showed in Lemma 3.1 that no matter the choice of f, for each a^* we can explicitly write down the utilities consistent with σ that maximize $\mathrm{sw}(a^*) - \mathbb{E}_{a \sim f(\sigma)}[\mathrm{sw}(a)]$ for any f, which drastically simplifies the problem of finding instance-optimal solutions.

In particular, given a profile σ, for each $a \in A$, compute u^a as defined in Lemma 3.1 as follows. For each voter i, let k_i be the position of a in σ_i. Then, for all $a' \in A$, let $u_i^a(a') = 1/k_i$ if $\sigma_i(a') \leq k_i$ and $u_i^a(a') = 0$ otherwise. Next, compute $\mathrm{sw}(a, a') := \mathrm{sw}_{u^a}(a')$ for all $a, a' \in A$. Now, letting $w_a := (\mathrm{sw}(a, a_1), \ldots, \mathrm{sw}(a, a_n))$, solve the following linear program to minimize distortion over all vectors of probabilities p, which correspond to social choice functions f.

$$\begin{aligned} &\min D \\ &\mathrm{sw}(a, a) - p^T w_a \leq D && \forall a \in A \\ &\sum_{a \in A} p_a = 1 \\ &p_a \geq 0 && \forall a \in A \end{aligned}$$

□

C Omitted Proofs from Section 4

Claim 4.2. *For any randomized SCF f,*

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

Proof. We begin with the first implication. If $d^+(f) \leq \beta \cdot n$, then by definition for all $(\sigma, u \triangleright \sigma)$ we have that $\max_a \mathrm{sw}(a) - \mathbb{E}[f(\sigma)] \leq \beta \cdot n$. Therefore

$$\begin{aligned} d_\alpha(f) &:= \max_\sigma \max_{\substack{u \triangleright \sigma \\ u \in U_\alpha}} \frac{\max_a \mathrm{sw}(a)}{\mathbb{E}_{a\sim f(\sigma)}[\mathrm{sw}(a)]} \\ &\leq \max_\sigma \max_{\substack{u \triangleright \sigma \\ u \in U_\alpha}} \frac{\max_a \mathrm{sw}(a)}{\max_a \mathrm{sw}(a) - \beta \cdot n} \\ &\leq \frac{\alpha \cdot n}{\alpha \cdot n - \beta \cdot n} \\ &= \frac{\alpha}{\alpha - \beta}, \end{aligned}$$

where the last inequality follows because $\alpha \geq \gamma$ by the definition of d^+.

Next, if $d_\alpha(f) \leq \gamma$, then depending on (σ, u) there are two cases: first, if for a given (σ, u) it is the case that $\max_a \mathrm{sw}(a) < \alpha \cdot n$, and so $\max_a \mathrm{sw}(a) - \mathbb{E}_{a\sim f(\sigma)}[\mathrm{sw}(a)] \leq \alpha \cdot n$ also. Otherwise, by the definition of d_α we have that $\max_a \mathrm{sw}(a) \leq \gamma \cdot \mathbb{E}_{a\sim f(\sigma)}[\mathrm{sw}(a)]$, and so

$$\begin{aligned} \max_a \mathrm{sw}(a) - \mathbb{E}_{a\sim f(\sigma)}[\mathrm{sw}(a)] &\leq \max_a \mathrm{sw}(a)(1 - 1/\gamma) \\ &\leq (1 - 1/\gamma) \cdot n. \end{aligned}$$

Combining these cases yields $\mathrm{dist}^+(f) \leq \max(\alpha, 1 - 1/\gamma) \cdot n$. □

Theorem 4.3. *For any randomized SCF f,*

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

Proof. Suppose without loss of generality that $k := 1/\alpha \in \mathbb{N}$ and that k and $\sqrt{m}$ divide n. Consider an instance for which the voters $[n]$ are arranged into t equal groups, each sharing a utility function. The first group has $u_i(a) = 1$ and $u_i(a') = 0$ for all $a' \neq a$. All remaining groups are indifferent to all alternatives, with ties broken in such a way that groups' utility profiles are rotationally symmetric for the first k alternatives.

Then $\mathrm{sw}(a) = \frac{n}{t} + \frac{n(t-1)}{t}\frac{1}{m}$, while $\mathrm{sw}(a') = \frac{n(t-1)}{t}\frac{1}{m}$ for all other alternatives. By the α-guarantee, we must have that $\mathrm{sw}(a) \geq \alpha \cdot n$, which is satisfied when $t \leq k$.

Given these symmetric rankings, the rule f can do no better than choosing uniformly over the first t alternatives, for an expected social welfare of $\mathrm{sw}(f(\mathcal{P})) = \frac{1}{t}(\frac{n}{t} + \frac{n(t-1)}{t}\frac{1}{m}) + \frac{t-1}{t}\frac{n(t-1)}{t}\frac{1}{m}$. Therefore the randomized distortion is

$$\begin{aligned} \mathrm{dist}(f) &= \frac{\frac{n}{t} + \frac{n(t-1)}{t}\frac{1}{m}}{\frac{1}{t}(\frac{n}{t} + \frac{n(t-1)}{t}\frac{1}{m}) + \frac{t-1}{t}\frac{n(t-1)}{t}\frac{1}{m}} \\ &= \begin{cases} \Omega(t) & t = O(\sqrt{m}) \\ \Omega(\frac{m}{t}) & t = \Omega(\sqrt{m}). \end{cases} \end{aligned}$$

For fixed m, this is maximized by choosing $t = \Theta(\sqrt{m})$, proving the distortion lower bound of (Boutilier et al. 2015). Since requiring $t \leq 1/\alpha$ is sufficient to satisfy the α-guarantee, it follows that

$$\mathrm{dist}_\alpha(f) = \Omega\left(\min\left\{\sqrt{m}, {}^1\!/\!_\alpha\right\}\right).$$ □

Theorem 4.4. *There is an SCF ℓ_α with* $\mathrm{dist}_\alpha(\ell_\alpha) = O\left(\frac{1}{\alpha}\right)$.

We take as our point of departure the analysis that Ebadian et al. (2022) give for their Stable Lottery Rule f_{SLR}. For fixed $k \in \mathbb{N}$, a *stable lottery* is a distribution $\mathcal{X}$ over committees $X \subseteq A$, of size $|X| = k$ for which the expected number of agents preferring any fixed alternative a^* to any $a \in X$ is small. In particular, for a fixed preference profile σ, the lottery $\mathcal{X}$ is *stable* if for all $a^* \in A$, $\Pr_{i\in N} \Pr_{X\sim\mathcal{X}}[a^* \succ_i X] \leq \frac{1}{k}$, where $a^* \succ_i X$ denotes that i ranks a^* ahead of all $a \in X$.

Such stable lotteries are shown to exist for all σ and k in Theorem 1 of (Cheng et al. 2020).

We will now define the rule ℓ_α.

Definition C.1. *For α with $1/\alpha \in \mathbb{N}$, the randomized SCF ℓ_α identifies some lottery $\mathcal{X}$ over committees of size $k = 2/\alpha$ which is* stable *for the input profile σ. It then samples a committee $X \sim \mathcal{X}$ and finally returns an alternative $a \sim X$ drawn uniformly at random.*

Although we do not emphasize the efficient computability of ℓ_α in theorem 4.4, (Cheng et al. 2020) show that stable lotteries sufficient for our purposes can be calculated in polynomial time, and so ℓ_α can be efficiently implemented.

Proof. To begin, assume that $1/\alpha \in \{2, 3, 4, \ldots\}$; all $\alpha \in [0, 1]$ are within a constant factor of such an $\alpha' \leq \alpha$, and in such cases an α-guarantee implies an α'-guarantee, and so this is without loss of generality.

As usual, let a^* denote the social-welfare-maximizing alternative. For a fixed committee X, let V_X denote the voters $i \in N$ for which $a^* \succ_i X$, and let $\overline{V}_X = N \setminus V_X$ denote its complement. By the definition of social welfare, for all X we have that

$$sw(a^*) = \sum_{i \in N} u_i(a^*) = \sum_{i \in V_X} u_i(a^*) + \sum_{i \in \overline{V}_X} u_i(a^*),$$

and so

$$sw(a^*) = \mathbb{E}_{X \sim \mathcal{X}} \left[\sum_{i \in V_X} u_i(a^*) + \sum_{i \in \overline{V}_X} u_i(a^*) \right]. \tag{27}$$

Since $u_i(a^*) \leq 1$, by the stability of $\mathcal{X}$ we have that

$$\mathbb{E}_{X \sim \mathcal{X}} \left[\sum_{i \in V_X} u_i(a^*) \right] \leq \mathbb{E}_{X \sim \mathcal{X}} \left[|V_X| \right] \leq \frac{n}{k}. \tag{28}$$

Using the α-promise that $sw(a^*) \geq \alpha \cdot n = 2\frac{n}{k}$, we find that

$$\begin{aligned}
\frac{1}{2} \cdot sw(a^*) &\leq sw(a^*) - \frac{n}{k} && (29)\\
&\leq sw(a^*) - \mathbb{E}_{X \sim \mathcal{X}} \left[\sum_{i \in V_X} u_i(a^*) \right] && (30)\\
&= \mathbb{E}_{X \sim \mathcal{X}} \left[\sum_{i \in \overline{V}_X} u_i(a^*) \right] && (31)\\
&\leq \mathbb{E}_{X \sim \mathcal{X}} \left[\sum_{i \in \overline{V}_X} k \cdot \mathbb{E}_{a \sim X} u_i(a) \right] && (32)\\
&\leq k \cdot \mathbb{E}_{X \sim \mathcal{X}} \left[\sum_{i \in N} \mathbb{E}_{a \sim X} u_i(a) \right] && (33)\\
&= \frac{2}{\alpha} \cdot \mathbb{E}_{a \sim \ell_\alpha(\sigma)} \left[sw(a) \right]. && (34)
\end{aligned}$$

Here (30) follows from (28), (31) follows from (27), (32) follows because if $i \in \overline{V}_X$ then i has at least as much utility for some $a \in X$ as for a^*, and this a has probability $1/k$ of being sampled from X, and (34) follows from the definitions of k and ℓ_α.

In conclusion, we have shown that $sw(a^*) \leq \frac{4}{\alpha} \mathbb{E}_{a \sim \ell_\alpha(\sigma)} [sw(a)]$, and therefore $\mathrm{dist}_\alpha(\ell_\alpha) = O(\frac{1}{\alpha})$. □

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

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

Proof. Suppose for now that $\mathrm{sw}(a^*) = \alpha$ exactly.

Next, we examine the first case. When $\alpha \geq 3/4$, this means that at least $n/2$ voters rank a^* first, and therefore, if we break ties in our favor, a^* is the winner under Plurality and the additive distortion is 0.

In the second case, consider a profile in which a $\frac{1}{2} - \epsilon$ fraction of voters rank a^* first (and value it at 1) and the other $\frac{1}{2} + \epsilon$ fraction of voters ranks an alternative $a \neq a^*$ first. This means that a is the winner under Plurality.

Now, let V_a be the set of voters who rank $a \neq a^*$ first. Of these voters, let a β fraction of them rank $a \succ a^* \succ \dots$, where they value a^* and a at $1/2$. The other $1-\beta$ fraction of voters in V_a are indifferent between all alternatives and value each of them at $1/m$. It is easy to verify that setting

$$\beta = \frac{\alpha - \frac{1}{2} - \frac{1}{2m}}{\frac{1}{2} - \frac{1}{m}}$$

suffices to ensure that a^* confers $\alpha \cdot n$ utility:

$$\frac{1}{2} + \frac{1}{2} \cdot \beta + \left(\frac{1}{2} - \beta\right) \cdot \frac{1}{m} = \alpha.$$

Plugging this value of β into the expression for additive distortion yields

$$\begin{aligned} \text{dist}_\alpha^+ &= \alpha - \left(\frac{1}{2} \cdot \beta + \left(\frac{1}{2} - \beta\right)\frac{1}{m}\right) \\ &= \alpha - \left(\frac{1}{2} \cdot \frac{\alpha - \frac{1}{2} - \frac{1}{2m}}{\frac{1}{2} - \frac{1}{m}} + \left(\frac{1}{2} - \frac{\alpha - \frac{1}{2} - \frac{1}{2m}}{\frac{1}{2} - \frac{1}{m}}\right)\frac{1}{m}\right) \\ &= 1/2. \qquad \text{(after simplifying)} \end{aligned}$$

It remains only to relax our original assumption that $\text{sw}(a^*) = \alpha$ exactly. Since the above expression is decreasing monotonically in $\text{sw}(a^*)$, for $\text{sw}(a^*) \geq 1/2$ we may assume that $\text{sw}(a^*) = \alpha$ as above.

Finally, by Theorem 3.6, this bound is tight. □

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

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

Proof. Suppose that $\text{sw}(a^*) = \alpha$.

In the first case, we must show that for any $\alpha \geq 3/4$, a^* will be selected by f_{Harm}. Let y_j represent the proportion of voters who rank a^* in position j. Note that the maximum score of any other alternative $a \neq a^*$ is $(1 - y_1) + \frac{1}{2} \cdot y_1$ because each voter who does not rank a^* first can rank a first, and each voter who ranks a^* first can rank a second. Furthermore, let $y_1^{\downarrow}(\alpha)$ represent the minimum value of y_1 such that the α promise holds, i.e., it must be the case that $\sum_j y_j \cdot \frac{1}{j} \geq \alpha$ because the maximum utility a^* can receive in position j is exactly $1/j$. Note that this value of $y_1^{\downarrow}(\alpha)$ also maximizes a's score. However, for all $\alpha \geq 3/4$, we can see that $y_1^{\downarrow}(\alpha) \geq 1/2$, meaning that $s(a) \leq \left(1 - \frac{1}{2}\right) + \frac{1}{2} \cdot \frac{1}{2} \leq \frac{3}{4}$, whereas $s(a^*) = \sum_j y_j \cdot \frac{1}{j} \geq \alpha$ due to the α promise guarantee. Therefore, for all $\alpha \in [3/4, 1]$, f_{Harm} will select a^* (breaking ties in our favor when $\alpha = 3/4$).

For the case of $2/3 \leq \alpha < 3/4$, consider the profile in which a $2(\alpha - 1/2)$ fraction of voters rank a^* first (with utility 1) and another alternative a in second (with utility 0), and a $2(1-\alpha)$ fraction of voters rank a^* second (with utility $1/2$) and a first (with utility $1/2$). Note that $u(a^*) = 2(\alpha - 1/2) \cdot n + \frac{1}{2} \cdot 2(1-\alpha) \cdot n = \alpha \cdot n$ and $u(a) = \frac{2(1-\alpha)}{2} \cdot n = (1-\alpha) \cdot n$.

With respect to scores, $s_{Harm}(a) = \frac{1}{2} \cdot 2(\alpha - 1/2) + 2(1-\alpha) = 3/2 - \alpha$, and $s_{Harm}(a^*) = 2(\alpha - 1/2) + \frac{1}{2} \cdot 2(1-\alpha) = \alpha$. For all $\alpha \in [2/3, 3/4)$, we indeed see that $s_{Harm}(a) > s_{Harm}(a^*)$ and therefore a is chosen, yielding an additive distortion of $\alpha - (1-\alpha) = 2\alpha - 1$.

For the case of $1/2 \leq \alpha < 2/3$, consider the profile in which an α fraction of voters rank a^* first (with utility 1) and a second (with utility 0), and the remaining $1-\alpha$ fraction of voters ranks a first (with utility $\frac{1}{m-1}$) and a^* last (with utility 0). Note that $u(a^*) = \alpha \cdot n$ and $u(a) = \frac{1-\alpha}{m-1} \cdot n$.

Turning to scores, $s_{Harm}(a) = \alpha \cdot \frac{1}{2} + (1-\alpha) = 1 - \frac{\alpha}{2}$, whereas $s_{Harm}(a^*) = \alpha + \frac{1-\alpha}{m}$. It is easy to verify that $s_{Harm}(a) > s_{Harm}(a^*)$ for all $\alpha \in [1/2, 2/3)$, so a will be selected, yielding an additive distortion of $\alpha - \frac{1-\alpha}{m-1}$.

Regardless of m, for $\alpha \in [1/2, 3/4]$ the two-alternative profile above with a utility promise of $3/4 \cdot n$ gives an additive distortion of $1/2 \cdot n$. □

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

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

Proof. Suppose that $\text{sw}(a^*) = \alpha$.

For the purposes of this proof, let $\alpha \cdot n$ be the exact social welfare conferred by the social-welfare-maximizing alternative. The stated bound for α-promise additive distortion follows by taking the maximum over all social welfare values above the promise.

When $\alpha \geq \frac{m-1}{m}$, the profile that minimizes the score of a^* consists of a $\frac{m-1}{m}$ fraction of voters who rank a^* first and assign it utility 1 and a $\frac{1}{m}$ fraction of voters who rank a^* last and assign it utility 0. The best any other alternative $a \neq a^*$ can do is be ranked second for a $\frac{m-1}{m}$ fraction of voters and first for a $\frac{1}{m}$ fraction of voters.

Therefore, $s_{Borda}(a^*) \geq \frac{m-1}{m} \cdot (m-1) + \frac{1}{m} \cdot (m-2)$, and $s_{Borda}(a) \leq \frac{m-1}{m} \cdot (m-2) + \frac{1}{m} \cdot (m-1)$. It is easy to verify that $s_{Borda}(a^*) > s_{Borda}(a)$.

In the second case, note that for $\alpha \in [1/2, (m-1)/m)$, one profile that achieves this α guarantee consists of an α fraction of voters who rank a^* first and assign it utility 1 and a $1-\alpha$ fraction of voters who rank a^* last and assign it utility 0. In this profile, $s_{Borda}(a^*) = \alpha \cdot (m-1)$. However, consider an alternative $a \neq a^*$ that is ranked second whenever a^* is ranked first and ranked first whenever a^* is ranked last. We can see that $s_{Borda}(a) = \alpha \cdot (m-2) + (1-\alpha) \cdot (m-1)$, which is greater than $s_{Borda}(a^*)$ for all $\alpha < \frac{m-1}{m}$, so the Borda rule chooses a.

Note that a may be valued at $1/m$ every time it is ranked first and at 0 every time it is ranked last. Therefore, the additive distortion is $\alpha - (1-\alpha) \cdot \frac{1}{m}$, which approaches α as m increases. □

Claim 4.8. *For Randomized Dictatorship,*

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

Proof. To begin we assume that $\mathrm{sw}(a^*) = \alpha$ exactly. In order to maximize additive distortion, by now-familiar arguments may we assume without loss of generality that all agents ranking a^* first value it at utility 1, and all not ranking a^* first value it second at utility $1/2$. Letting β_a denote the proportion of agents ranking $a \in A$ first, we have

$$\mathrm{sw}(a^*) = \frac{1}{2}(1+\beta_{a^*}), \qquad \mathrm{sw}(a) = \frac{\beta_a}{2} \qquad \text{for } a \neq a^*.$$

Omitting the factor of n and noting that $\beta_{a^*} = 2\alpha - 1$ (for $\alpha \geq 1/2$), we have

$$\begin{aligned} \operatorname{dist}^+_\alpha(RD) &= \max_\beta \left(\alpha - \sum_a \beta_a \, \mathrm{sw}(a)\right) \\ &= \max_\beta \left(\alpha - \left(\beta_{a^*}\alpha + \sum_{a \neq a^*} \beta_a \, \mathrm{sw}(a)\right)\right) \\ &= \alpha - \left(\beta_{a^*}\alpha + \frac{m-1}{2}\left(\frac{1-\beta_{a^*}}{m-1}\right)^2\right) \\ &= 2\alpha(1-\alpha) - \frac{2(1-\alpha)^2}{m-1}. \end{aligned}$$

It remains only to relax our original assumption that $\mathrm{sw}(a^*) = \alpha$ exactly. Therefore to find the α-promise additive distortion, we maximize over all $\alpha' \geq \alpha$ and obtain the stated bound. □