

Voting with Preference Intensities

Anson Kahng¹, Mohamad Latifian², Nisarg Shah²

¹ University of Rochester

² University of Toronto

akahng2@cs.rochester.edu, latifian@cs.toronto.edu, nisarg@cs.toronto.edu

Abstract

When an agent votes, she typically ranks the set of available alternatives. Occasionally, she may also wish to report the intensity of her preferences by indicating adjacent pairs of alternatives in her ranking between which her preference is acutely decisive; for instance, she may suggest that she likes alternative a more than b , but b *much more* than c . We design near-optimal voting rules which aggregate such *preference rankings with intensities* using the recently-popular distortion framework. We also show that traditional voting rules, which aggregate preference rankings while ignoring (or not eliciting) intensities, can incur significant welfare loss.

Introduction

Professor X wants to take her students out for a group lunch. Being a computational social choice researcher, she realizes that this is the perfect opportunity to conduct real-life voting, so she offers them the choice between four popular restaurants (a , b , c , and d) and asks the ambiguously worded — though, perhaps intentionally so — question: “Tell me your preferences over these restaurants”. One of her students, Sam, responds with the *preference ranking* $a > b > c > d$, which means he likes a the most and d the least. Meera has a slightly different preference ranking, $b > d > c > a$. Additionally, she also wants to convey that she likes b and d *much more* than a or c (she is vegetarian and the latter two lack good vegetarian options). Thus, she responds with a *preference ranking with intensities*, $b > d \gg c > a$; the ‘ \gg ’ between d and c indicates that her preference intensity drops sharply from d to c .¹

Professor X was planning to use the Borda count, a well-established voting rule for aggregating preference rankings that would give each restaurant 3, 2, 1, and 0 points each time it is ranked first, second, third, and fourth, respectively, and pick the restaurant with the most points as the winner. But now, she begins to wonder how she should aggregate preference rankings with intensities indicated by ‘ \gg ’ signs. A natural approach would be to modify Borda count so that there is a steeper drop in points awarded when encountering a ‘ \gg ’ sign. But exactly how steep should it be? In Meera’s

case ($b > d \gg c > a$), perhaps the four restaurants should be awarded 3, 2.5, 0.5, and 0 points in order. Or should it be 3, 2.1, 0.9, and 0 points? Professor X is not satisfied with this, and for a more systematic approach, she looks towards the *distortion* framework.

Introduced by Procaccia and Rosenschein (2006), the distortion framework posits that ordinal preferences reported by the agents stem from their underlying numerical utility functions over the set A of m alternatives. That is, when agent i reports that she prefers alternative a to b (denoted $a > b$), it must be the case that $u_i(a) \geq u_i(b)$ according to her underlying utility function $u_i : A \rightarrow \mathbb{R}_{\geq 0}$. The goal of the distortion framework is to seek voting rules which, subject to the available partial information about agents’ utility functions, minimize the worst-case approximation ratio between the highest possible utilitarian social welfare (sum of agent utilities) and that of the outcome picked by the rule; this quantity is termed *distortion*. For aggregating preference rankings (without intensities), the best possible distortion for deterministic and randomized voting rules is known to be $\Theta(m^2)$ (Caragiannis and Procaccia 2011; Caragiannis et al. 2017) and $\Theta(\sqrt{m})$ (Boutilier et al. 2015; Ebadian et al. 2022), respectively.

The distortion framework is extremely versatile. It can be used not only to analyze how to aggregate ranked preferences, but also to analyze how to aggregate other ballot formats such as top- t preferences (Borodin et al. 2022), threshold approval votes (Benade et al. 2021), and ranked preferences coupled with additional queries (Amanatidis et al. 2021; Brams and Sanver 2009); compare different ballot formats (Benade et al. 2021); and even design optimal ballot formats (Mandal et al. 2019; Mandal, Shah, and Woodruff 2020).

We argue that this framework also provides a systematic approach to aggregating preference rankings with intensities. Let $\alpha \in [0, 1]$ be a parameter that governs how decisive an agent must be for her to use a ‘ \gg ’ sign. When agent i reports $a > b$, we still assume that $u_i(a) \geq u_i(b)$. But when agent i reports $a \gg b$, we assume that the agent has an α -decisive preference for a over b given by $\alpha \cdot u_i(a) \geq u_i(b)$. Thus, when $\alpha = 0$, ‘ \gg ’ signs provide significant additional information since the agent’s utility must drop to 0 starting with the alternative after the first ‘ \gg ’ sign (if it exists). In contrast, when $\alpha \approx 1$, ‘ \gg ’ signs provide little additional in-

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹We use “intensity” and “decisiveness” interchangeably.

formation and are essentially equivalent to ‘>’ signs.

Given a value of α , the distortion framework allows us to identify an alternative (or a distribution over alternatives) with the the minimum distortion, when agents (voluntarily) use ‘ \succcurlyeq ’ signs to reveal α -decisive preferences. Alternatively, the designer may choose a value of α and mandate all agents to use ‘ \succcurlyeq ’ signs *whenever* they have α -decisive preferences; in this case, agent i reporting $a > b$ would translate to a stronger condition of $u_i(a) \geq u_i(b) > \alpha \cdot u_i(a)$ because we know the agent’s preference for a over b is decidedly not decisive (no pun intended). In both these cases of voluntary and mandatory intensity revelation, one can not only study optimal voting rules, but also the social welfare loss that traditional voting rules can incur by either not eliciting or ignoring the intensity information. This leads to our main research questions:

How should one aggregate preference rankings with intensities? Can a better distortion be guaranteed if all agents have decisive preferences or if the designer mandates revealing decisive preferences whenever they exist? What welfare loss might the designer incur when using a traditional voting rule that simply aggregates preference rankings without intensities?

Our Contribution

A key conceptual contribution of our work is to introduce a formal model in which agents can report α -decisive preferences using the ‘ \succcurlyeq ’ sign and the impact of this on the distortion of voting rules can be analyzed.

We begin our analysis with the *top decisiveness* setting, in which every agent is assumed to be α -decisive between their top two choices. For this setting, which has been considered in the literature under a different model, we identify asymptotically optimal distortion bounds for both deterministic and randomized voting rules. We also extend this to *uniform decisiveness*, in which every agent is assumed to be α -decisive between every pair of adjacent alternatives in her preference ranking, and provide tight distortion bounds for deterministic rules.

In general when agents can have α -decisive preferences at arbitrary locations in their preference rankings, we introduce and study the *price of ignoring the intensities (POII)*, which captures the efficiency loss of traditional voting rules that do not elicit preference intensities. When agents can voluntarily report their α -decisive preferences (but are not required to do so), we show that both deterministic and randomized rules can suffer from a significant POII. When agents must mandatorily report their α -decisive preferences, the POII can be even higher, and we prove a better lower bound for deterministic rules.

In case of mandatory reporting, by proving non-trivial lower bounds, we raise the interesting open question of whether one can establish better distortion bounds for deterministic rules by carefully choosing the value of α , but show that this cannot be done for randomized rules.

Related Work

Our work builds on the distortion framework, specifically the utilitarian distortion framework in which agents have normalized utilities for alternatives (Procaccia and Rosenschein 2006). There is a closely related framework of metric distortion (Anshelevich et al. 2018; Anshelevich and Postl 2017; Abramowitz, Anshelevich, and Zhu 2019), in which agents and alternatives are embedded in an underlying metric space, agents have costs for the alternatives measured by the distance between them, and agents rank the alternatives by these costs. While constant distortion of 3 can already be guaranteed in this framework by aggregating ranked preferences using a deterministic rule (Gkatzelis, Halpern, and Shah 2020; Kizilkaya and Kempe 2022), Anshelevich and Postl (2017) introduced the idea of deriving improved distortion bounds when the underlying metric space is α -decisive, that is, when each agent’s distance to her closest alternative is at most α times her distance to her next-closest alternative. This has induced a significant body of follow-up work (Gross, Anshelevich, and Xia 2017; Ghodsi, Latifian, and Seddighin 2019; Gkatzelis, Halpern, and Shah 2020; Anagnostides, Fotakis, and Patsilinos 2022). The top-decisiveness setting from our paper is precisely the counterpart of this in the utilitarian distortion framework, where each agent’s utility for her second-best alternative is assumed to be at most α times her utility for her best alternative. Our uniform decisiveness setting takes this one step further, where the agent is decisive between every adjacent pair of alternatives in her preference ranking, not only between her top two alternatives.

For our setting in which agents voluntarily reveal their decisiveness, improved distortion bounds cannot be guaranteed because in the worst case, none of the agents may reveal any decisiveness. For this setting, we study the price of ignoring the intensities, which is the maximum factor by which the distortion can increase due to not taking into account the voluntarily revealed decisiveness information; here, the worst case is when agents do reveal substantial useful information. This is similar to the concept of competitive ratio from online algorithms (Borodin and El-Yaniv 2005), except the ignorance in our case is rectifiable.

Our setting in which the designer chooses a value of α and forces agents to reveal whenever they have α -decisive preferences is equivalent to querying agents for additional information beyond their ranked preferences, which has been studied by Amanatidis et al. (2021). Specifically, our setting is a special case of their framework, in which $m - 1$ comparison queries are made to each agent, one for every adjacent pair of alternatives in her preference ranking. While we show a lower bound of $\Omega(m)$ for our setting (and leave open the exciting question of whether this is tight), they show that constant distortion can be achieved by using $O(\log^2 m)$ arbitrary comparison queries per agent. Mandal et al. (2019) and Mandal, Shah, and Woodruff (2020) extend this idea to allow arbitrary queries and search over the space of all possible ballot formats to find ones that provide the optimal trade-off between distortion and the number of bits of information elicited from each agent. We refer the interested reader to

the survey by Anshelevich et al. (2021) for further related work on distortion.

Preliminaries

For $t \in \mathbb{N}$, define $[t] = \{1, 2, \dots, t\}$, and for a set X , define $\Delta(X)$ to be the set of all distributions over X . Let $N = [n]$ be a set of n agents, and $A = \{a_1, a_2, \dots, a_m\}$ be a set of m alternatives.

Preference rankings with intensities. Each agent i has ordinal preferences over the alternatives given by a *preference ranking with intensities*, denoted $\sigma_i = (\pi_i, \bowtie_i)$, where $\pi_i : [m] \rightarrow A$ is a one-to-one function that determines a ranking over the alternatives, and $\bowtie_i : [m-1] \rightarrow \{', >', \gg'\}$ is a function that determines the intensities: $\bowtie_i(j) = '>'$ indicates a preference for $\pi_i(j)$ over $\pi_i(j+1)$, whereas $\bowtie_i(j) = \gg'$ indicates a strong preference for $\pi_i(j)$ over $\pi_i(j+1)$. Let $\vec{\sigma} = (\sigma_1, \dots, \sigma_n)$ denote the ordinal preference profile, and $\mathcal{S}(A)$ be the set of all total orderings with intensities over A .

Voting rule. A (randomized) voting rule $f : \mathcal{S}(A)^n \rightarrow \Delta(A)$ maps an ordinal preference profile to a distribution over the alternatives. We say that f is deterministic if it always outputs a distribution with singleton support; in this case, we use $f(\vec{\sigma})$ to directly denote the single alternative in the support of the distribution chosen by f on $\vec{\sigma}$.

Utilitarian framework. We assume that the ordinal preferences of the agents stem from underlying cardinal preferences. Specifically, each agent i has an underlying utility function $u_i : A \rightarrow \mathbb{R}_{\geq 0}$, where $u_i(a)$ denotes the utility of agent i for alternative a . For a distribution $x \in \Delta(A)$, we write $u_i(x) = \mathbb{E}_{a \sim x} u_i(a)$. Following the literature (Aziz 2020), we work with unit-sum utilities, where $\sum_{a \in A} u_i(a) = 1$ for each $i \in N$.

Let $\vec{u} = (u_1, \dots, u_n)$ be the utility profile. Let $\alpha \in [0, 1]$ be a parameter which may be chosen either by the designer or implicitly by the agents. We say that utility profile \vec{u} is α -consistent with preference profile $\vec{\sigma}$, denoted $\vec{u} \triangleright_{\alpha} \vec{\sigma}$, if, for all $i \in N$ and $j \in [m-1]$, we have $u_i(\pi_i(j)) \geq u_i(\pi_i(j+1))$ if $\bowtie_i(j) = '>'$ and $\alpha \cdot u_i(\pi_i(j)) \geq u_i(\pi_i(j+1))$ if $\bowtie_i(j) = \gg'$.

Given a utility profile \vec{u} , the social welfare of distribution $x \in \Delta(A)$ is $\text{sw}(x, \vec{u}) = \sum_{i \in N} u_i(x)$. The distortion of x over \vec{u} is defined as

$$\text{dist}(x, \vec{u}) = \frac{\max_{a \in A} \text{sw}(a, \vec{u})}{\text{sw}(x, \vec{u})},$$

and the α -distortion of x on a preference profile $\vec{\sigma}$ is defined as $\text{dist}_{\alpha}(x, \vec{\sigma}) = \sup_{\vec{u} \triangleright_{\alpha} \vec{\sigma}} \text{dist}(x, \vec{u})$. Throughout the paper we might use candidate a when a distribution $x \in \Delta(A)$ is expected. In this case by a we mean a distribution that gives probability 1 to a and zero to the rest of the candidates.

Intuitively, distortion quantifies the efficiency of a distribution x for a preference profile $\vec{\sigma}$ by taking the worst-case ratio between the social welfare of the optimal alternative and the (expected) social welfare of x over all utility profiles that are α -consistent with the preference profile, which, in this case, includes preferences intensities.

We also define the α -distortion of a voting rule f as

$$\text{dist}_{\alpha}(f) = \max_{\vec{\sigma} \in \mathcal{S}(A)^n} \text{dist}_{\alpha}(f(\vec{\sigma}), \vec{\sigma}),$$

or as the worst case α -distortion of the output of f over all preference profiles.

Top Decisiveness

Without assuming that agents have at least somewhat decisive preferences, one cannot hope to improve distortion bounds ($O(m^2)$ for deterministic rules and $O(\sqrt{m})$ for randomized rules (Ebadian et al. 2022)) because, in the worst case, none of the agents may use the \gg' sign. Hence, in the related literature on metric distortion, a number of papers have analyzed distortion bounds subject to a natural restriction on agent preferences: all agents are assumed to be α -decisive in their preference for their best alternative over their second-best alternative, for a fixed α (Anshelevich and Postl 2017; Gross, Anshelevich, and Xia 2017; Ghodsi, Latifian, and Seddighin 2019; Gkatzelis, Halpern, and Shah 2020; Anagnostides, Fotakis, and Patsilinas 2022).

We study such preferences in the utilitarian framework, and rename them as top α -decisive to indicate that the decisiveness is at the top of the preference ranking. While identifying asymptotically optimal distortion bounds for top α -decisive preferences is still an open question in the metric distortion framework (Gkatzelis, Halpern, and Shah 2020), we are able to obtain tight distortion bounds for this in the utilitarian framework.

Definition 1 (Top α -Decisiveness). Let $\alpha \in [0, 1]$. We say that agent i with utility function u_i is top α -decisive if there exists an alternative a^* such that $u_i(a) \leq \alpha \cdot u_i(a^*)$ for all $a \neq a^*$. To restate this, define \bowtie^{top} to be a function where $\bowtie^{\text{top}}(1) = \gg'$ and $\bowtie^{\text{top}}(j) = '>'$ for $1 < j < m$, and $\mathcal{S}_d(A) \subset \mathcal{S}(A)$ to be the set of preference rankings with intensities containing $\sigma = (\pi, \bowtie^{\text{top}})$ for all preference rankings π . Then, agent i is top α -decisive if there exists $\sigma_i \in \mathcal{S}_d(A)$ such that $u_i \triangleright_{\alpha} \sigma_i$. A preference profile $\vec{\sigma}$ is top α -decisive if $\vec{\sigma} \in \mathcal{S}_d(A)^n$, and a utility profile \vec{u} is top α -decisive if $\vec{u} \triangleright_{\alpha} \vec{\sigma}$ for some $\vec{\sigma} \in \mathcal{S}_d(A)^n$.

Definition 2 (Top α -decisive distortion of rule f). The distortion of a voting rule f with respect to top α -decisive preferences is defined as

$$\text{dist}_{\alpha}^{\text{top}}(f) = \max_{\vec{\sigma} \in \mathcal{S}_d(A)^n} \text{dist}_{\alpha}(f(\vec{\sigma}), \vec{\sigma}).$$

First, we show matching lower (Theorem 1) and upper (Theorem 2) bounds on the best possible distortion of deterministic rules with respect to top α -decisive preferences.

Theorem 1. For any $\alpha \in [0, 1]$, the distortion of every deterministic voting rule with respect to top α -decisive utilitarian spaces is $\Omega(\alpha^2 m^2 + 1)$.

The proof of this result and all other missing proofs are available in the full version.²

Next, we show that the simple plurality rule provides a matching upper bound for all α , despite ignoring preference

²Full version: <https://www.cs.toronto.edu/~nisarg/papers/intensities.pdf>

intensities and being oblivious to the value of α . This is perhaps not so surprising: plurality is known to provide asymptotically the best distortion for aggregating (not necessarily decisive) ranked preferences (Caragiannis and Procaccia 2011; Caragiannis et al. 2017), and assuming top decisiveness only increases the importance of focusing on the top choices of the agents, which plurality does.

Theorem 2. For any $\alpha \in [0, 1]$, the distortion of plurality with respect to top α -decisive preferences is $O(\alpha^2 m^2 + 1)$.

Proof. Let f_{plu} denote the plurality rule. Consider a top α -decisive preference profile $\vec{\sigma}$ and a utility profile \vec{u} such that $\vec{u} \triangleright_{\alpha} \vec{\sigma}$. Let a^p be the plurality winner and $a^* \in \arg \max_{a \in A} \text{sw}(a, \vec{u})$ be an optimal alternative. Let N^{a^p} and N^{a^*} denote the sets of agents who have a^p and a^* as their top choices in $\vec{\sigma}$, respectively.

Then, due to the utility functions being unit-sum and top α -decisive, we have

$$\begin{aligned} \text{sw}(a^p, \vec{u}) &\geq |N^{a^p}| \cdot \frac{1}{\alpha(m-1) + 1}, \\ \text{sw}(a^*, \vec{u}) &\leq |N^{a^*}| \cdot 1 + (n - |N^{a^*}|) \cdot \frac{\alpha}{\alpha + 1}. \end{aligned}$$

Using the fact that $|N^{a^p}| \geq \max(|N^{a^*}|, n/m)$, we have

$$\begin{aligned} \text{dist}(f_{\text{plu}}(\vec{\sigma}), \vec{u}) &= \frac{\text{sw}(a^*, \vec{u})}{\text{sw}(a^p, \vec{u})} \\ &\leq \alpha(m-1) + 1 + \frac{\alpha(m-1)}{\alpha + 1} \cdot (\alpha(m-1) + 1) \\ &= O(m^2 \alpha^2 + 1). \quad \square \end{aligned}$$

Note that the best distortion bound of deterministic rules drops from $\Theta(m^2)$ at $\alpha = 1$ (which is the traditional setting) to $\Theta(1)$ at $\alpha = \Theta(1/m)$, and then stays $\Theta(1)$.

As is often the case with distortion-based analysis, randomized rules offer significantly better guarantees. We next present matching lower (Theorem 3) and upper (Theorem 4) bounds for the distortion of randomized voting rules with respect to top α -decisive preferences, with some minor results along the way.

Theorem 3. The distortion of every randomized voting rule with respect to top α -decisive utilitarian spaces is $\Omega(\frac{\alpha m + 1}{\alpha \sqrt{m + 1}})$.

Given our results for deterministic rules, one may wonder whether randomized rules that are known to be (near-)optimal for aggregating ranked preferences ($\alpha = 1$) may also happen to be (near-)optimal for aggregating top α -decisive preferences for all $\alpha \in [0, 1]$, despite being oblivious to preference intensities and the value of α . For aggregating ranked preferences, the harmonic rule (Boutlier et al. 2015) provides $O(\sqrt{m} \log m)$ distortion while the stable lottery rule (Ebadian et al. 2022) provides the optimal $O(\sqrt{m})$ distortion. While these distortion upper bounds trivially hold for aggregating top α -decisive preferences, unfortunately the distortion of these rules does not seem to improve when $\alpha < 1$.

Fortunately, we are able to modify the stable lottery rule to design a new randomized voting rule, which pays

increased attention to agents' top choices and provably achieves asymptotically optimal distortion with respect to top α -decisive preferences, while still being oblivious to preference intensities as well as the value of α .

Definition 3. We define f_{dec} to be a randomized voting rule that, with probability $\frac{1}{2}$, runs the stable lottery rule (f_{slr}); with probability $\frac{1}{4}$, selects an alternative uniformly at random from the set of alternatives which are the top choice of at least one agent; and with the remaining probability $\frac{1}{4}$, selects an alternative with the maximum plurality score (i.e., is the top-choice of the most agents).

Theorem 4. For every $\alpha \in [0, 1]$, the distortion of the (randomized) rule f_{dec} with respect to top α -decisive preferences is $O(\frac{\alpha m + 1}{\alpha \sqrt{m + 1}})$.

Proof. Consider any top α -decisive preference profile $\vec{\sigma} = (\vec{\pi}, \boxtimes)$, and let $\vec{u} \triangleright_{\alpha} \vec{\sigma}$ be the underlying utility profile. Let $a^* \in \arg \max_{a \in A} \text{sw}(a, \vec{u})$ be an optimal alternative. Define N^a to be the set of agents who have alternative a as their top choice, $N^{-a} = N \setminus N^a$, and \mathcal{T} to be the set of the alternatives that are the top choice of at least one agent. For each $i \in N^{-a^*}$ we have $u_i(\pi_i(1)) \geq \frac{u_i(a^*)}{\alpha}$, so we have

$$\begin{aligned} \sum_{a \in \mathcal{T}} \text{sw}(a, \vec{u}) &\geq \sum_{i \in N} u_i(\pi_i(1)) \geq \sum_{i \in N^{-a^*}} u_i(\pi_i(1)) \\ &\geq \sum_{i \in N^{-a^*}} \frac{u_i(a^*)}{\alpha}. \end{aligned} \quad (1)$$

The plurality winner a^p must be the top choice of at least $|N^{a^*}|$ agents, each of whom must have utility at least $\frac{1}{\alpha(m-1)+1} \geq \frac{1}{\alpha m + 1}$ due to having a unit-sum and top α -decisive utility function. Thus, we have

$$\text{sw}(a^p, \vec{u}) \geq \frac{|N^{a^*}|}{\alpha m + 1}. \quad (2)$$

Putting Equations (1) and (2) together with the fact that the distortion of f_{slr} is at most $2\sqrt{m}$ (Ebadian et al. 2022), we have:

$$\begin{aligned} \mathbb{E}[\text{sw}(f_{\text{dec}}(\vec{\sigma}), \vec{u})] &\geq \frac{1}{4} \left(\sum_{i \in N^{-a^*}} \frac{u_i(a^*)}{\alpha m} + \frac{|N^{a^*}|}{\alpha m + 1} + \frac{\text{sw}(a^*, \vec{u})}{\sqrt{m}} \right) \\ &\geq \frac{1}{4} \left(\sum_{i \in N^{-a^*}} \frac{u_i(a^*)}{\alpha m + 1} + \frac{|N^{a^*}|}{\alpha m + 1} + \frac{\text{sw}(a^*, \vec{u})}{\sqrt{m}} \right) \\ &\geq \frac{1}{4} \left(\frac{\text{sw}(a^*, \vec{u})}{\alpha m + 1} + \frac{\text{sw}(a^*, \vec{u})}{\sqrt{m}} \right) \\ &\geq \frac{\text{sw}(a^*, \vec{u})}{4} \cdot \frac{\alpha m + \sqrt{m} + 1}{\alpha m \sqrt{m} + \sqrt{m}} \geq \text{sw}(a^*, \vec{u}) \cdot \frac{\alpha \sqrt{m} + 1}{4(\alpha m + 1)}, \end{aligned}$$

and hence,

$$\begin{aligned} \text{dist}(f_{\text{dec}}(\vec{\sigma}), \vec{u}) &= \frac{\text{sw}(a^*, \vec{u})}{\mathbb{E}[\text{sw}(f_{\text{dec}}(\vec{\sigma}), \vec{u})]} \\ &\leq \frac{4(\alpha m + 1)}{\alpha \sqrt{m} + 1} = O\left(\frac{\alpha m + 1}{\alpha \sqrt{m} + 1}\right). \quad \square \end{aligned}$$

Note that the distortion stays $\Theta(\sqrt{m})$ from $\alpha = 1$ to $\alpha = \Theta(1/\sqrt{m})$, then drops to $\Theta(1)$ by $\alpha = \Theta(1/m)$, and then stays $\Theta(1)$ as α drops further.

Uniform Decisiveness

In some applications, it may be natural for agents to be α -decisive not only between their top two choices, but between every pair of adjacent alternatives in their ranking. We initiate the study of this natural restriction on agent preferences. We say that agent i is uniform α -decisive if $\bowtie_i = (\succ, \dots, \succ)$ in her ordinal preferences $\sigma_i = (\pi_i, \bowtie_i)$. We then define uniform α -decisive preference profiles and utility profiles, as well as the distortion of a voting rule with respect to uniform α -decisive preferences similarly to the case of top α -decisiveness.

While any distortion upper bounds with respect to top α -decisive preferences continue to hold with respect to uniform α -decisive preferences, one may hope to find improved distortion bounds now that the agents are more decisive. Indeed, we are able to provide improved matching lower (Theorem 5) and upper (Theorem 6) bounds for deterministic rules. For randomized rules, we present a lower bound (Theorem 7), but leave open the question of whether an upper bound better than the one in Theorem 4 can be derived.

Theorem 5. For every $\alpha \in [0, 1]$, the distortion of every deterministic voting rule with respect to uniform α -decisive preferences is $\Omega\left(\frac{(m\alpha+1)(1-\alpha^m)}{1-\alpha}\right)$.

Theorem 6. For every $\alpha \in [0, 1]$, the distortion of the (deterministic) plurality rule with respect to uniform α -decisive preferences is $O\left(\frac{(m\alpha+1)(1-\alpha^m)}{1-\alpha}\right)$.

Theorem 7. For every $\alpha \in [0, 1]$, the distortion of every randomized voting rule with respect to uniform α -decisive preferences is $\Omega\left(\min\left(\sqrt{m}, \frac{1-\alpha^m}{1-\alpha}\right)\right)$.

Voluntary Reporting of Intensities

We now relax the assumption that all agents have similarly decisive preferences. Instead, we study a setting in which agents may have decisive preferences at arbitrary positions in their preference rankings, and even when they do, they may choose to not reveal them and use the ' $>$ ' sign instead of the ' \succ ' sign. As argued before, one cannot hope to derive improved distortion bounds in this case because it is possible that none of the agents use ' \succ ' sign anywhere in the preference profile.³

Instead, we focus on the loss of efficiency that a traditional voting rule can incur in the worst case, measured by increased α -distortion, due to aggregating only preference rankings and not preference rankings with intensities. This is meaningful as the worst case now occurs when the agents

³One can still seek to computationally find the distortion-optimal distribution on any given (not necessarily worst-case) preference profile. It is easy to see that this requires only a slight modification in the linear program Boutilier et al. (2015), where $u_i(a) \geq u_i(b)$ is replaced by $\alpha \cdot u_i(a) \geq u_i(b)$ whenever agent i reports $a \succ b$ (in adjacent positions in her preference ranking).

could have revealed α -decisive preferences by using the ' \succ ' sign, but the traditional voting rule either did not elicit such intensities or chose to ignore them. We initiate the study of this efficiency loss, which we term the *price of ignoring the intensities*.

Definition 4 (Intensity-aware optimal). Let $\text{opt}_\alpha^{\text{aw}}(\vec{\sigma})$ be a distribution x over A that minimizes the worst-case distortion over the utility profiles that are α -consistent with $\vec{\sigma}$, i.e.,

$$\text{opt}_\alpha^{\text{aw}}(\vec{\sigma}) = \arg \min_{x \in \Delta(A)} \text{dist}_\alpha(x, \vec{\sigma}).$$

Definition 5 (Price of ignoring the intensities (POII)). We define the price of ignoring the intensities (POII) of a distribution $x \in \Delta(A)$ on a preference profile $\vec{\sigma}$ as the ratio between the α -distortion of x and that of the intensity-aware optimal distribution:

$$\text{Poll}(x, \vec{\sigma}, \alpha) = \frac{\text{dist}_\alpha(x, \vec{\sigma})}{\text{dist}_\alpha(\text{opt}_\alpha^{\text{aw}}(\vec{\sigma}), \vec{\sigma})}.$$

When x is chosen based only on the ranked preference profile (without intensities) $\vec{\pi}$, its POII on $\vec{\pi}$ is defined as $\text{Poll}(x, \vec{\pi}, \alpha) = \max_{\vec{\sigma} \triangleright \vec{\pi}} \text{Poll}(x, \vec{\sigma}, \alpha)$, where we use $\vec{\sigma} \triangleright \vec{\pi}$ to denote that $\vec{\sigma} = (\vec{\pi}, \vec{\bowtie})$ for some $\vec{\bowtie}$. This allows us to define both the *intensity-oblivious optimal* distribution on $\vec{\pi}$ as $\text{opt}_\alpha^{\text{ob}}(\vec{\pi}) = \arg \min_{x \in \Delta(A)} \text{Poll}(x, \vec{\pi}, \alpha)$ and the POII on $\vec{\pi}$ as $\text{Poll}(\vec{\pi}, \alpha) = \text{Poll}(\text{opt}_\alpha^{\text{ob}}(\vec{\pi}), \vec{\pi}, \alpha)$. We are interested in the worst case of this over all $\vec{\pi}$, termed the *POII for α -decisive preferences*: $\text{Poll}(\alpha) = \max_{\vec{\pi}} \text{Poll}(\vec{\pi}, \alpha)$.

We observe the following lemma, which provides a way to derive a lower bound on the price of ignorance.

Lemma 1. For any ranked preference profile (without intensities) $\vec{\pi}$, preference profile $\vec{\sigma} \triangleright \vec{\pi}$, and distribution $x \in \Delta(A)$, we have:

$$\text{Poll}(\vec{\pi}, \alpha) \geq \frac{\text{dist}_\alpha(\text{opt}_\alpha^{\text{ob}}(\vec{\pi}), \vec{\sigma})}{\text{dist}_\alpha(x, \vec{\sigma})}.$$

Proof. By the definitions,

$$\begin{aligned} \text{Poll}(\vec{\pi}, \alpha) &\geq \text{Poll}(\text{opt}_\alpha^{\text{ob}}(\vec{\pi}), \vec{\sigma}, \alpha) = \frac{\text{dist}_\alpha(\text{opt}_\alpha^{\text{ob}}(\vec{\pi}), \vec{\sigma})}{\text{dist}_\alpha(\text{opt}_\alpha^{\text{aw}}(\vec{\sigma}), \vec{\sigma})} \\ &\geq \frac{\text{dist}_\alpha(\text{opt}_\alpha^{\text{ob}}(\vec{\pi}), \vec{\sigma})}{\text{dist}_\alpha(x, \vec{\sigma})}. \quad \square \end{aligned}$$

Theorem 8. For any $\alpha \in [0, 1]$, the price of ignoring the intensities for α -decisive preferences is $\text{Poll}(\alpha) = \Omega\left(\frac{\sqrt{m}(1-\alpha)}{1-\alpha^m} + 1\right)$.

Proof. For ease of exposition, assume \sqrt{m} divides n . Partition the agents into \sqrt{m} equal-sized subsets $N_1, \dots, N_{\sqrt{m}}$, and consider ranked preference profile (without intensities) $\vec{\pi}$, where members of N_j rank a_j first and the rest of the alternatives in a cyclic order.⁴

⁴All we need is that for $j > 1$, each alternative appears in the j -th position in the preference rankings of at most n/m agents.

Fix an intensity-oblivious optimal distribution $\text{opt}_\alpha^{\text{ob}}(\vec{\pi})$. Without loss of generality, assume that $\text{opt}_\alpha^{\text{ob}}(\vec{\pi})$ places the lowest probability on a_1 among the alternatives $a_1, \dots, a_{\sqrt{m}}$ (thus, this probability is at most $1/\sqrt{m}$).

Now, define intensities in a way that for $i \in N_1$, we have $\bowtie_i = (\gg, \dots, \gg)$, whereas for all other agents $i' \notin N_1$, we have $\bowtie_{i'} = (>, \dots, >)$. Consider the preference profile $\vec{\sigma} = (\vec{\pi}, \vec{\bowtie})$. We show that $\text{opt}_\alpha^{\text{ob}}(\vec{\pi})$ has a significant POII on $\vec{\sigma}$. We use Lemma 1 to derive this in two steps: proving a lower bound on $\text{dist}_\alpha(\text{opt}_\alpha^{\text{ob}}(\vec{\pi}), \vec{\sigma})$ and proving an upper bound on $\text{dist}_\alpha(x, \vec{\sigma})$ for some distribution $x \in \Delta(A)$.

Step 1. For proving a lower bound on $\text{dist}_\alpha(\text{opt}_\alpha^{\text{ob}}(\vec{\pi}), \vec{\sigma})$, consider the utility profile \vec{u}^* in which all members of N_1 have utility 1 for a_1 and zero for all other alternatives, and all other agents have utility $1/m$ for every alternative. Since $\vec{u}^* \triangleright_\alpha \vec{\sigma}$, we can see that

$$\text{dist}_\alpha(\text{opt}_\alpha^{\text{ob}}(\vec{\pi}), \vec{\sigma}) \geq \text{dist}_\alpha(\text{opt}_\alpha^{\text{ob}}(\vec{\pi}), \vec{u}^*) \geq \frac{\sqrt{m}}{2}. \quad (3)$$

We omit the detailed calculation as it is identical to the one given in Boutilier et al. (2015) for proving a distortion lower bound on randomized voting rules for aggregating ranked preferences (without intensities).⁵

Step 2. On the other hand, consider the distribution x that places probability 1 on a_1 . To prove an upper bound on $\text{dist}_\alpha(x, \vec{\sigma}) = \text{dist}_\alpha(a_1, \vec{\sigma})$, consider any utility profile \vec{u} such that $\vec{u} \triangleright_\alpha \vec{\sigma}$. For each $i \in N_1$, we can see that $u_i(a_1) \geq \frac{1-\alpha}{1-\alpha^m}$ due to unit-sum and uniform α -decisive utility function u_i . Hence,

$$\text{sw}(a_1, \vec{u}) \geq \frac{n}{\sqrt{m}} \cdot \frac{1-\alpha}{1-\alpha^m}. \quad (4)$$

Also, every other alternative a appears as the top choice of at most n/\sqrt{m} agents (who each have utility at most 1 for it), and for $1 < j \leq m$, it appears in the j -th position in the preference rankings of at most n/m of the agents (who each have utility at most $1/j$ for it due to their utility functions being unit-sum). Hence, we have

$$\text{sw}(a, \vec{u}) \leq \frac{n}{\sqrt{m}} \cdot 1 + \frac{n}{m} \cdot H_m \leq \frac{2n}{\sqrt{m}}, \quad (5)$$

where $H_m = \sum_{i=1}^m 1/i$ is the m -th harmonic number. From Equations (4) and (5), we have that

$$\text{dist}_\alpha(a_1, \vec{\sigma}) \leq \frac{2(1-\alpha^m)}{1-\alpha}.$$

Using Lemma 1 and Equation (3), we have

$$\begin{aligned} \text{Poll}(\alpha) &\geq \text{Poll}(\vec{\pi}, \alpha) \geq \frac{\text{dist}_\alpha(\text{opt}_\alpha^{\text{ob}}(\vec{\pi}), \vec{\sigma})}{\text{dist}_\alpha(a_1, \vec{\sigma})} \\ &\geq \frac{(1-\alpha)\sqrt{m}}{4(1-\alpha^m)} = \Omega\left(\frac{(1-\alpha)\sqrt{m}}{1-\alpha^m} + 1\right). \end{aligned}$$

The last step follows because the POII is always at least 1. \square

⁵While they derive a lower bound of $\sqrt{m}/3$, it is not difficult to see that a careful calculation in their analysis yields $\sqrt{m}/2$.

You can see that when α is a constant less than 1, $\text{Poll}(\alpha) = \Theta(\sqrt{m})$. Note that we can not hope for a higher POII since the stable lottery rule (Ebadian et al. 2022) already guarantees $O(\sqrt{m})$ distortion while being intensity oblivious. That means we have a rule that could not suffer from a multiplicative increase in distortion by an $\omega(\sqrt{m})$ factor compared to the intensity-aware optimal distribution in hindsight. That said, for other regimes of α , it remains to be seen whether the bound from Theorem 8 is tight.

The POII can similarly be defined when restricting to deterministic choices. Here, we prove the following lower bound, which is not necessarily tight even when α is a constant less than 1, resulting in another open question.

Theorem 9. For any $\alpha \in [0, 1]$, the price of ignoring the intensities of deterministic rules for α -decisive preferences is $\Omega\left(\frac{m(1-\alpha)}{1-\alpha^m} + 1\right)$.

Mandatory Reporting of Intensities

To this point, we have studied the setting where agents may *choose* to report decisive preferences. Crucially, they are not *required* to report all (or indeed any) α -decisive preferences. One natural question is whether a designer can do significantly better by choosing a value of α and requiring all agents to report all locations in their preference ranking where their preference between adjacent alternatives is α -decisive. In this case, when agent i reports $a > b$ for adjacent alternatives a and b in her preference ranking, we know not only that $u_i(a) \geq u_i(b)$ (agent i prefers a to b), but also that $\alpha \cdot u_i(a) \leq u_i(b)$ (the preference is *not* α -decisive).

We say that utility function u is *strictly α -consistent* with preference ordering $\sigma = (\pi, \bowtie)$, and denote it with $u \triangleright_\alpha^+ \sigma$, if u is α -consistent with σ , and for each $j \in [m-1]$ where $\bowtie(j) = '>'$, we have $\alpha \cdot u(\pi(j)) \leq u(\pi(j+1))$. A utility profile \vec{u} is strictly α -consistent with a preference profile $\vec{\sigma}$ if $u_i \triangleright_\alpha^+ \sigma_i$ for each agent i .

In this section, we assume that agents' utility functions are strictly α -consistent with the preference ordering they submit to the voting rule. Our goal is to see if we can get improved distortion bounds in this case.

Definition 6 (Strict Distortion). The strict α -distortion of distribution $x \in \Delta(A)$ on preference profile $\vec{\sigma}$ is defined as:

$$\text{dist}_\alpha^S(x, \vec{\sigma}) = \sup_{\vec{u} \triangleright_\alpha^+ \vec{\sigma}} \text{dist}(x, \vec{u}).$$

We prove a lower bound on the best possible strict distortion of any deterministic rule as a function of α . Our bound is $\Omega(m^2)$ at $\alpha = \Theta(1)$ and $\alpha = O(1/m)$, showing that mandatory reporting of intensities cannot help reduce distortion in these regimes. However, our bound is the weakest ($\Omega(m^{5/4})$) at $\alpha = \Theta(1/\sqrt[4]{m})$, raising the interesting possibility of significantly improving upon the $O(m^2)$ distortion by choosing the right value of α and mandating agents to report α -decisive preferences.

Theorem 10. For any $\alpha \in [0, 1]$, the strict α -distortion of every deterministic voting rule f satisfies:

$$\text{dist}_\alpha^S(f) = \Omega\left(\max\left(m^2\alpha^3, \frac{m^2}{m\alpha + 1}\right)\right).$$




Figure 1: A conceptual plot of the bounds in Theorem 10.

The bound achieves its weakest value of $\Omega(m^{5/4})$ at $m = \Theta(1/\sqrt[4]{m})$ (see Figure 1).

Proof sketch. We give two bounds that complement each other as shown in Figure 1.

For the first bound (the gray line in Figure 1) consider a profile in which agents all have the same alternative as their second choice but are split evenly among favorite alternatives. All agents are decisive only between the first and second, and second and third positions, i.e., $\bowtie_i = (\gg, \gg, >, \dots, >)$.

For the second bound (the black line in Figure 1), consider a profile in which agents are divided into groups. All the agents have a_m as their second choice, each group of agents ranks a particular subset of alternatives first, and the rest in an arbitrary order. For any deterministic voting rule f , we can case on whether f chooses a_m or not. \square

While it may be possible to improve the distortion of deterministic rules via mandatory reporting of intensities, we prove that this is decidedly not the case for randomized rules: regardless of the value of $\alpha \in [0, 1]$, every randomized rule has strict distortion $\Omega(\sqrt{m})$, meaning that the stable lottery rule of Ebadian et al. (2022) achieves the optimal strict distortion of $O(\sqrt{m})$ for all $\alpha \in [0, 1]$. The result is proved by deriving two separate lower bounds, which establish the desired implication in complementary regions of α , as depicted in Figure 2.

Theorem 11. For every $\alpha \in [0, 1]$, every (randomized) voting rule f has strict α -distortion $\text{dist}_\alpha^S(f) = \Omega(\sqrt{m})$.

Proof sketch. We consider the classic lower bound instance from Boutilier et al. (2015), where \sqrt{m} of the alternatives appear as the top choices of n/\sqrt{m} agents each. We bound the strict distortion of this instance with respect to two different intensity profiles: in the first profile, each agent i reports $\bowtie_i = (>, \dots, >)$, and in the second profile, each agent i reports $\bowtie_i = (\gg, >, \dots, >)$.

POII with Mandatory Reporting of Intensities

One can define the price of ignoring the intensities in this case like in the case of voluntary reporting of intensities. However, in the case of mandatory reporting of intensities, an intensity-oblivious rule has potentially more to lose: had




Figure 2: A conceptual plot of the bounds in Theorem 11.

it elicited intensities, it could have obtained additional information not only when a voter used the ‘ \gg ’ sign (which indicates the preference is decisive), but also when the voter used the ‘ $>$ ’ sign (which indicates the preference is *not* decisive).

We are in fact able to improve upon our lower bound from Theorem 9. While that lower bound started at $\Omega(m)$ at $\alpha = 0$ and decreased to the trivial bound of $\Omega(1)$ as $\alpha \rightarrow 1$, our new bound starts at $\Omega(m)$ at $\alpha = 0$, but increases to $\Omega(m^2)$ as $\alpha \rightarrow 1$. Note that $\Omega(m^2)$ is the highest possible POII because plurality achieves $O(m^2)$ distortion in an intensity-oblivious manner, so it cannot suffer from a multiplicative increase in distortion by an $\omega(m^2)$ factor due to the lack of intensity elicitation.

Theorem 12. For every $\alpha \in [0, 1]$, the price of ignoring the intensities of deterministic rules with mandatory reporting of α -decisive preferences is $\Omega(\frac{m(1-\alpha^m)}{1-\alpha})$.

Discussion

Our work uncovers several exciting open questions. Perhaps the most compelling of them is whether one can improve the distortion of deterministic rules to $O(m)$ by choosing a value of $\alpha \in [0, 1]$ and mandating all agents to reveal whenever their preferences are α -decisive. A compelling candidate is $\alpha = \Theta(1/\sqrt{m})$ because this is where our lower bound from Theorem 10 achieves its weakest value of $\Omega(m)$. Other open questions include settling the optimal distortion of randomized rules subject to uniform decisiveness (see Theorem 7 for our lower bound), the POII when agents reveal their α -decisive preferences voluntarily (see Theorems 8 and 9 for our lower bounds), and the POII when agents reveal their α -decisive preferences mandatorily (see Theorem 12).

In the related literature on metric distortion, top decisiveness (i.e., decisiveness between the best and second-best alternatives) has been explored extensively, but assuming other forms of decisiveness (such as uniform decisiveness) and eliciting decisive preferences deserves further exploration. Exploring the impact of decisive agent preferences on other important desiderata such as fairness (e.g., proportional fairness (Ebadian et al. 2022)) and strategyproofness is also an exciting direction for the future.

Acknowledgments

Latifian and Shah were partially supported by an NSERC Discovery Grant.

References

- Abramowitz, B.; Anshelevich, E.; and Zhu, W. 2019. Awareness of voter passion greatly improves the distortion of metric social choice. In *International Conference on Web and Internet Economics*, 3–16. Springer.
- Amanatidis, G.; Birmapas, G.; Filos-Ratsikas, A.; and Voudouris, A. A. 2021. Peeking behind the ordinal curtain: Improving distortion via cardinal queries. *Artificial Intelligence*, 296: 103488.
- Anagnostides, I.; Fotakis, D.; and Patsilinos, P. 2022. Metric-distortion bounds under limited information. *Journal of Artificial Intelligence Research*, 74: 1449–1483.
- 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. 2021. Distortion in Social Choice Problems: The First 15 Years and Beyond. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, 4294–4301. Survey Track.
- Anshelevich, E.; and Postl, J. 2017. Randomized social choice functions under metric preferences. *Journal of Artificial Intelligence Research*, 58: 797–827.
- Aziz, H. 2020. Justifications of welfare guarantees under normalized utilities. *ACM SIGecom Exchanges*, 17(2): 71–75.
- Benade, G.; Nath, S.; Procaccia, A. D.; and Shah, N. 2021. Preference elicitation for participatory budgeting. *Management Science*, 67(5): 2813–2827.
- Borodin, A.; and El-Yaniv, R. 2005. *Online computation and competitive analysis*. Cambridge University Press.
- Borodin, A.; Halpern, D.; Latifian, M.; and Shah, N. 2022. Distortion in voting with top-t preferences. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*, 116–122.
- Boutilier, C.; Caragiannis, I.; Haber, S.; Lu, T.; Procaccia, A. D.; and Sheffet, O. 2015. Optimal social choice functions: A utilitarian view. *Artificial Intelligence*, 227: 190–213.
- Brams, S. J.; and Sanver, M. R. 2009. Voting systems that combine approval and preference. In *The mathematics of preference, choice and order*, 215–237. Springer.
- Caragiannis, I.; Nath, S.; Procaccia, A. D.; and Shah, N. 2017. Subset selection via implicit utilitarian voting. *Journal of Artificial Intelligence Research*, 58: 123–152.
- Caragiannis, I.; and Procaccia, A. D. 2011. Voting almost maximizes social welfare despite limited communication. *Artificial Intelligence*, 175(9-10): 1655–1671.
- Ebadian, S.; Kahng, A.; Peters, D.; and Shah, N. 2022. Optimized Distortion and Proportional Fairness in Voting. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC)*, 563–600.
- Ghodsi, M.; Latifian, M.; and Seddighin, M. 2019. On the distortion value of the elections with abstention. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, volume 33, 1981–1988.
- Gkatzelis, V.; Halpern, D.; and Shah, N. 2020. Resolving the optimal metric distortion conjecture. In *Proceedings of the 61st Symposium on Foundations of Computer Science (FOCS)*, 1427–1438.
- 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.
- Kizilkaya, F. E.; and Kempe, D. 2022. Plurality Veto: A Simple Voting Rule Achieving Optimal Metric Distortion. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*, 349–355.
- Mandal, D.; Procaccia, A. D.; Shah, N.; and Woodruff, D. P. 2019. Efficient and thrifty voting by any means necessary. In *Proceedings of the 33rd Annual Conference on Neural Information Processing Systems (NeurIPS)*, 7180–7191.
- Mandal, D.; Shah, N.; and Woodruff, D. P. 2020. Optimal communication-distortion tradeoff in voting. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*, 795–813.
- Procaccia, A. D.; and Rosenschein, J. S. 2006. The distortion of cardinal preferences in voting. In *Proceedings of the 10th International Workshop on Cooperative Information Agents (CIA)*, 317–331.