CS 580 (Fall 2023) Project Report
Distortion in Single-Winner Elections

Eklavya Sharma

eklavya2@illinois.edu

Abstract. In single-winner elections where voters submit rankings over candidates, *distortion* is a metric to evaluate voting rules using a utilitarian framework. I compare the distortion of the plurality voting rule to the distortion of the optimal deterministic voting rule.

Keywords: Social choice · voting · distortion.

1 Introduction

Consider a scenario where n agents want to pick an outcome from m choices, and the agents may have different preferences. This happens in many real-life scenarios:

1. Elections for selecting representatives of governing bodies.
2. Students in a course want to decide whether the final exam should be take-home or in-class.
3. A group of friends want to decide where to go for their next road trip.
4. Coordination among multiple AI agents.

How do we pick an outcome from the m choices that aligns the most with the agents' preferences? We call this problem 'single-winner election'. Henceforth, we use election terminology, i.e., there are n voters who want to elect a single winner from m candidates.

(I assume you are already familiar with the concept of distortion in single-winner elections. In Sections 1.1 and 1.2, I give a brief description of the problem statement to refresh your memory and establish notation. If you would like a more gentle introduction instead, please refer to Section D.)

1.1 Utility and Welfare

For any non-negative integer t, let $[t] := \{1, 2, \ldots, t\}$.

The j^{th} voter has a *utility function* $u_j : [m] \to \mathbb{R}$. Here $u_j(i)$ is a number that tells us how much voter j likes candidate i. We normalize utilities by enforcing $\sum_{i=1}^{m} u_j(i) = 1$. The list $U := (u_1, \ldots, u_n)$ is called the *utility profile*.

Each voter j has a weight $w_j \in \mathbb{R}_{\geq 0}$. Every weighted voting instance with rational weights can be mapped to an equivalent unweighted voting instance. Given weights $w \in [0,1]^n$ and utility profile U, the *social welfare* of candidate i is defined as

$$\mathrm{SW}_{(w,U)}(i) := \sum_{j=1}^{n} w_j u_j(i).$$

When w is clear from context, we write SW_U instead of $\mathrm{SW}_{(w,U)}$.

1.2 Ordinal Preferences and Distortion

Often, we cannot elicit exact numeric utilities from voters since it imposes a cognitive and communication burden. Hence, voting systems often ask voters to rank the candidates.

Therefore, instead of observing the utility profile U, we observe a *preference profile* $\Pi := (\pi_1, \pi_2, \ldots, \pi_m)$. Here $\pi_j : [m] \to [m]$ is a bijection where $\pi_j(i)$ is the j^{th} voter's i^{th} favorite candidate. We say that utility profile U is *consistent* with preference profile Π (denoted as $U \rhd \Pi$) if for each voter j, we have $u_j(\pi_j(1)) \geq u_j(\pi_j(2)) \geq \ldots \geq u_j(\pi_j(m))$. We receive as input the pair (w, Π) (called a *ranked voting instance*), where $w \in [0,1]^n$ is the vector of voters' weights.

Our goal is to find a candidate whose social welfare is within a small factor of the best social welfare achievable in each consistent utility profile. This factor is called *distortion*. Formally, for input (w, Π), the distortion of candidate i is defined as

$$\Delta(i, (w, \Pi)) := \sup_{U : U \rhd \Pi} \frac{\max_{i^* \in [m]} \mathrm{SW}_{(w,U)}(i^*)}{\mathrm{SW}_{(w,U)}(i)}.$$

Our aim is to find a candidate of low distortion.

We may also pick a candidate randomly, in which case we extend the definition of distortion to distributions over candidates. Formally, let $p \in [0,1]^m$ be a vector such that $\sum_{i=1}^{m} p_i = 1$. Suppose we pick candidate i with probability p_i for all $i \in [m]$. Then the expected value of our chosen candidate's social welfare is $\sum_{i=1}^{m} p_i \,\mathrm{SW}_{(w,U)}(i)$. Hence, we define the distortion of p for the input (w, Π) as

$$\Delta(p, (w, \Pi)) := \sup_{U : U \rhd \Pi} \frac{\max_{i^* \in [m]} \mathrm{SW}_{(w,U)}(i^*)}{\sum_{i \in [m]} p_i \,\mathrm{SW}_{(w,U)}(i)}.$$

An algorithm that takes input (w, Π) and outputs a candidate is called a *voting rule*. The distortion of a (randomized) voting rule on (w, Π) is the distortion of (the distribution of) the voting rule's output on (w, Π).

Conventions to ensure that suprema of ratios are well-defined:

1. Define ∞ to be larger than any real number.
2. For any $a > 0$, let $a/0 := \infty$ and $0/0 := 0$.
3. For $f : X \to \mathbb{R} \cup \{\infty\}$, if $f(x) = \infty$ for some $x \in X$, or if f is not upper-bounded by a real-number, then $\sup_{x \in X} f(x) := \infty$.

1.3 Known Results

The *plurality score* of candidate i is the total weight of voters whose most-preferred candidate is i. The *plurality voting rule* outputs a *plurality winner*, i.e, a candidate with the highest plurality score.

Theorem 1 (based on [3]). *There is a ranked voting instance (w, Π) with m candidates for which every candidate has distortion $\Omega(m^2)$. For any input, the plurality winner's distortion is at most m^2.*

Theorem 2 ([2]). *There is a ranked voting instance (w, Π) with m candidates for which every distribution p has distortion at least $\sqrt{m}/3$. There is a randomized voting rule having distortion at most $O(\sqrt{m} \log^* m)$ for every input (where $\log^* m$ is the iterated logarithm of m).*

Theorem 3 ([4]). *There is a randomized voting rule whose distortion is at most $2\sqrt{m}$ for every ranked voting instance with m candidates.*

Theorem 4 (based on [2]). *Given any ranked voting instance (w, Π) with m candidates and n voters, we can find the distribution p that minimizes the distortion by solving a linear program in $O(nm^2)$ variables and $O(nm^2)$ constraints.*

1.4 My Contributions

A cursory glance at Theorem 1 makes it seem that plurality achieves the best distortion that any deterministic voting rule can achieve. However, a closer looks reveals a caveat: there may exist inputs where some candidate has low distortion but plurality still picks a candidate with high distortion. This prompts the question: how good is plurality's distortion compared to the optimal deterministic distortion? I answer this question in the following two theorems, which I prove in Section A.

Theorem 5. *For any ranked voting instance (w, Π) with m candidates, let i_p be the plurality winner, let β be her plurality score, and let $\widehat{\imath}$ be some other candidate. Then*

$$\Delta(i_p, (w, \Pi)) \le 1 + \frac{m}{2\beta} \le 1 + \frac{m^2}{2}, \qquad \Delta(\widehat{\imath}, (w, \Pi)) \ge 1 + m.$$

When $\beta \ge 1/2$, Theorem 5 implies that the plurality winner has the least distortion. Moreover, we show that Theorem 5's analysis is tight.

Theorem 6. *For parameters $m \in \mathbb{Z}_{\ge 3}$, $\beta \in (\frac{1}{m}, \frac{1}{2}]$, and $\varepsilon \in (0, \beta - \frac{1}{m}]$, there is a ranked voting instance (w, Π) with m candidates such that all of the following hold:*

1. *Candidate 1 is the unique plurality winner. Her plurality score is β.*
2. *$\Delta(1, (w, \Pi)) \ge 1 + \frac{m}{2\beta}(1 - \varepsilon)$.*
3. *$\Delta(m, (w, \Pi)) = 1 + \frac{m}{1 - \varepsilon/\beta}$.*

When $\beta \to \frac{1}{m}$ and $\varepsilon \to 0$, Theorem 6 implies that $\Delta(1, (w, \Pi)) \to 1 + m^2/2$ and $\Delta(m, (w, \Pi)) \to 1 + m$.

These results rely on a characterization of the utility profiles that maximize the ratio of the social welfare of any two given candidates, which I present in Section 2. A slight variation of this characterization was independently discovered by [1] (Appendix D in their paper)[1].

This characterization also helps us get a fast algorithm for computing the distortion of any candidate in a ranked voting instance.

Theorem 7. *Given a ranked voting instance (w, Π) with m candidates and n voters, we can compute $\Delta(\widehat{\imath}, (w, \Pi))$ in $O(mn)$ time for any candidate $\widehat{\imath}$.*

Moreover, to find the candidate with the optimal distortion, we can just compute the distortion of each candidate and pick the best one. This can be done in $O(nm^2)$ time.

2 Characterizing Worst-Case Utility Profiles

I define a class of utility profiles, called *ruthless* utility profiles, and show that if a utility profile maximizes the ratio of the social welfares of any two given candidates, then it must be ruthless.

[1] I got scooped! 😭

Definition 1 (domination). *In a ranked voting instance $(w, \mathit{\Pi})$, we say that candidate i_1* dominates *candidate i_2 if for every voter j with positive weight, j prefers i_1 over i_2 (i.e., $\pi_j^{-1}(i_1) < \pi_j^{-1}(i_2)$).*

Definition 2 (ruthlessness). *Let $(w, \mathit{\Pi})$ be a ranked voting instance with m candidates and let i^* and $\widehat{\imath}$ be any two candidates. A voter j is called '*ruthless *for $(i^*, \widehat{\imath})$ in utility profile U' iff the following hold:*

1. *If i^* is j's k^{th}-favorite candidate and she prefers i^* over $\widehat{\imath}$ (i.e., $k := \pi_j^{-1}(i^*) < \pi_j^{-1}(\widehat{\imath})$), then j has utility $1/k$ for her top-k candidates and utility 0 for all other candidates (i.e., $u_j(\pi_j(i)) = 1/k$ for all $i \in [k]$).*
2. *If j prefers $\widehat{\imath}$ over i^* and $\widehat{\imath}$ is j's k^{th}-favorite candidate for $k > 1$ (i.e., $1 < k := \pi_j^{-1}(\widehat{\imath}) < \pi_j^{-1}(i^*)$), then j has utility 0 for each candidate that is not among her top $k-1$ candidates (i.e., $u_j(\pi_j(i)) = 0$ for all $i \geq k$).*
3. *If $\widehat{\imath}$ is j's favorite candidate (i.e., $\pi_j^{-1}(\widehat{\imath}) = 1$), then $u_j(i) = 1/m$ for all $i \in [m]$.*

A utility profile U is called ruthless for $(i^*, \widehat{\imath})$ iff all voters of positive weight are ruthless for $(i^*, \widehat{\imath})$ in U.

Observation 8. *For any ranked voting instance $(w, \mathit{\Pi})$ and any pair of candidates $(i^*, \widehat{\imath})$, a ruthless utility profile always exists. Furthermore, for every utility profile U that is ruthless for $(i^*, \widehat{\imath})$, we have $U \rhd \mathit{\Pi}$ and $\mathrm{SW}_U(i^*) \geq \mathrm{SW}_U(\widehat{\imath})$.*

Theorem 9. *Let $(w, \mathit{\Pi})$ be a ranked voting instance. Let i^* and $\widehat{\imath}$ be two candidates such that $\widehat{\imath}$ doesn't dominate i^*. Let $\mathcal{U}$ be the set of all utility profiles consistent with $\mathit{\Pi}$. For any utility profile U, define*

$$f(U) := \frac{\mathrm{SW}_{(w,U)}(i^*)}{\mathrm{SW}_{(w,U)}(\widehat{\imath})}.$$

Then the set of ruthless utility profiles for $(i^*, \widehat{\imath})$ is exactly the set of maximizers of f over $\mathcal{U}$. Formally, for any utility profile $\widehat{U}$, we have $f(\widehat{U}) = \sup_{U \rhd \mathit{\Pi}} f(U)$ iff $\widehat{U}$ is ruthless for $(i^*, \widehat{\imath})$.

I prove Theorem 9 in Section B. Now let us see how Theorem 9 helps us get a fast algorithm for computing the distortion for any candidate.

Theorem 7. *Given a ranked voting instance $(w, \mathit{\Pi})$ with m candidates and n voters, we can compute $\mathit{\Delta}(\widehat{\imath}, (w, \mathit{\Pi}))$ in $O(mn)$ time for any candidate $\widehat{\imath}$.*

Proof. $\widehat{\imath}$ dominates all other candidates iff $\widehat{\imath}$'s plurality score is 1. We can check this in $O(n)$ time. If $\widehat{\imath}$ dominates everyone else, then $\mathit{\Delta}(\widehat{\imath}, (w, \mathit{\Pi})) = 1$ (since $\mathrm{SW}_U(\widehat{\imath}) \ge \mathrm{SW}_U(i)$ for all $i \in [m]$ and all $U \rhd \mathit{\Pi}$).

In $O(mn)$ time, we can compute π_j^{-1} for all $j \in [n]$. $\widehat{\imath}$ dominates i iff $\pi_j^{-1}(\widehat{\imath}) < \pi_j^{-1}(i)$ for all j. We can check this in $O(n)$ time.

Let C be the candidates in $[m] \setminus \{\widehat{\imath}\}$ not dominated by $\widehat{\imath}$. Then

$$\begin{aligned}\mathit{\Delta}(\widehat{\imath}, (w, \mathit{\Pi})) &= \sup_{U : U \rhd \mathit{\Pi}} \frac{\max_{i \in [m]} \mathrm{SW}_U(i)}{\mathrm{SW}_U(\widehat{\imath})} \\ &= \max\left(1, \max_{i \in C} \sup_{U : U \rhd \mathit{\Pi}} \frac{\mathrm{SW}_U(i)}{\mathrm{SW}_U(\widehat{\imath})}\right).\end{aligned}$$

We can compute each of these $|C|$ suprema in $O(n)$ time using Theorem 9 and Definition 2. Hence, we can compute $\mathit{\Delta}(i, (w, \mathit{\Pi}))$ in $O(mn)$ time.

3 Open Problems

Although known works have characterized the worst-case distortion of (both deterministic and randomized) voting rules asymptotically as a function of the number of candidates m, it remains to be seen if we can get an exact bound on the distortion in terms of m. This is especially relevant when m is small, which is often the case.

In Section C, I explore the case $m = 2$, and show the following

1. The distortion of plurality is at most 3, and there is an instance for which no deterministic voting rule can get distortion better than 3.
2. There is a randomized voting rule having distortion at most 3/2, and there is an instance for which no randomized voting rule can get distortion better than 3/2.

References

1. Anonymous authors: Computational aspects of distortion. Submission #889 in AAMAS (2024)
2. Boutilier, C., Caragiannis, I., Haber, S., Lu, T., Procaccia, A.D., Sheffet, O.: Optimal social choice functions: A utilitarian view. In: ACM Conference on Electronic Commerce (EC). pp. 197–214 (2012). https://doi.org/10.1145/2229012.2229030
3. Caragiannis, I., Procaccia, A.D.: Voting almost maximizes social welfare despite limited communication. Artificial Intelligence **175**(9), 1655–1671 (2011). https://doi.org/10.1016/j.artint.2011.03.005
4. Ebadian, S., Kahng, A., Peters, D., Shah, N.: Optimized distortion and proportional fairness in voting. In: ACM Conference on Economics and Computation (EC). pp. 563–600. Association for Computing Machinery (2022). https://doi.org/10.1145/3490486.3538339

A The Distortion of Plurality

We first prove upper and lower bounds on the distortion of a candidate using Theorem 9.

Lemma 1. *Consider a ranked voting instance (w, Π) with m candidates. For any $i \in [m]$, let H_i be the set of voters who have candidate i as their first preference, and let h_i be the total weight of voters in H_i. Then for any two candidates i^* and $\widehat{\imath}$ such that $\widehat{\imath}$ doesn't dominate i^*,*

$$1 + \frac{mh_{i^*}}{h_{\widehat{\imath}}} \leq \sup_{U : U \rhd \Pi} \frac{\mathrm{SW}_{(w,U)}(i^*)}{\mathrm{SW}_{(w,U)}(\widehat{\imath})} \leq 1 + \frac{m}{2}\left(\frac{1 + h_{i^*}}{h_{\widehat{\imath}}} - 1\right).$$

Proof. By Theorem 9, the sup is achieved by a utility profile $\widehat{U}$ that is ruthless for $(i^*, \widehat{\imath})$. Hence,

1. $\forall j \in H_{i^*}$, $\widehat{u}_j(i^*) = 1$ and $\widehat{u}_j(\widehat{\imath}) = 0$.
2. $\forall j \in H_{\widehat{\imath}}$, $\widehat{u}_j(i^*) = \widehat{u}_j(\widehat{\imath}) = 1/m$.
3. $\forall j \in [m] \setminus (H_{i^*} \cup H_{\widehat{\imath}})$, $\widehat{u}_j(\widehat{\imath}) = 0$ and $\widehat{u}_j(i^*) \in [0, 1/2]$.

Hence, $\mathrm{SW}_{\widehat{U}}(\widehat{\imath}) = h_{\widehat{\imath}}/m$ and

$$\begin{aligned}
\mathrm{SW}_{\widehat{U}}(i^*) &\in \left[h_{i^*} + \frac{h_{\widehat{\imath}}}{m}, h_{i^*} + \frac{h_{\widehat{\imath}}}{m} + (1 - h_{i^*} - h_{\widehat{\imath}})\frac{1}{2}\right] \\
&= \left[h_{i^*} + \frac{h_{\widehat{\imath}}}{m}, \frac{h_{\widehat{\imath}}}{m} + \frac{1 + h_{i^*} - h_{\widehat{\imath}}}{2}\right].
\end{aligned}$$

Hence,

$$\frac{\mathrm{SW}_{\widehat{U}}(i^*)}{\mathrm{SW}_{\widehat{U}}(\widehat{\imath})} \in \left[1 + \frac{mh_{i^*}}{h_{\widehat{\imath}}}, 1 + \frac{m}{2}\left(\frac{1 + h_{i^*}}{h_{\widehat{\imath}}} - 1\right)\right].$$

□

Lemma 2. *Consider a ranked voting instance (w, Π) with m candidates. For any $i \in [m]$, let h_i be the plurality score of candidate i. Let $i^* :=$ $\mathrm{argmax}_{i \in [m] \setminus \{\widehat{\imath}\}} h_i$. Then for any candidate $\widehat{\imath}$, we get*

$$1 + \frac{mh_{i^*}}{h_{\widehat{\imath}}} \leq \Delta(\widehat{\imath}, (w, \Pi)) \leq 1 + \frac{m}{2}\left(\frac{1 + h_{i^*}}{h_{\widehat{\imath}}} - 1\right).$$

Proof. If $\widehat{\imath}$ dominates everyone else, then $h_{\widehat{\imath}} = 1$, $h_{i^*} = 0$, and $\Delta(\widehat{\imath}, (w, \Pi)) = 1$ (since $\mathrm{SW}_U(\widehat{\imath}) \geq \mathrm{SW}_U(i)$ for all $i \in [m]$ and all $U \rhd \Pi$). Hence, the lemma holds trivially.

Let C be the set of candidates in $[m] \setminus \{\widehat{\imath}\}$ not dominated by $\widehat{\imath}$. Then $C \neq \emptyset \implies h_{\widehat{\imath}} < 1 \implies h_{i^*} > 0 \implies i^* \in C$.

$$\begin{aligned}
\Delta(\widehat{\imath}, (w, \Pi)) &= \sup_{U : U \rhd \Pi} \frac{\max_{i \in [m]} \mathrm{SW}_U(i)}{\mathrm{SW}_U(\widehat{\imath})} \\
&= \max\left(1, \max_{i \in C} \sup_{U : U \rhd \Pi} \frac{\mathrm{SW}_U(i)}{\mathrm{SW}_U(\widehat{\imath})}\right) \\
&\in \left[\max_{i \in C}\left(1 + \frac{m h_i}{h_{\widehat{\imath}}}\right), \max_{i \in C}\left(1 + \frac{m}{2}\left(\frac{1 + h_i}{h_{\widehat{\imath}}} - 1\right)\right)\right] \\
&\qquad \text{(by Lemma 1)} \\
&= \left[1 + \frac{m h_{i^*}}{h_{\widehat{\imath}}}, 1 + \frac{m}{2}\left(\frac{1 + h_{i^*}}{h_{\widehat{\imath}}} - 1\right)\right]. \qquad \square
\end{aligned}$$

Theorem 5. *For any ranked voting instance (w, Π) with m candidates, let i_p be the plurality winner, let β be her plurality score, and let $\widehat{\imath}$ be some other candidate. Then*

$$\Delta(i_p, (w, \Pi)) \le 1 + \frac{m}{2\beta} \le 1 + \frac{m^2}{2}, \qquad \Delta(\widehat{\imath}, (w, \Pi)) \ge 1 + m.$$

Proof. Let i_q be the candidate with the second highest plurality score, and let her plurality score be γ. Then $\gamma \le \beta$ and $\beta \ge 1/m$. Hence, by Lemma 2, we get

$$\Delta(i_p, (w, \Pi)) \le 1 + \frac{m}{2}\left(\frac{1 + \gamma}{\beta} - 1\right) \le 1 + \frac{m}{2\beta} \le 1 + \frac{m^2}{2}.$$

$$\Delta(\widehat{\imath}, (w, \Pi)) \ge 1 + \frac{m\beta}{\gamma} \ge 1 + m. \qquad \square$$

We now show that Theorem 5's analysis is tight by considering a specific ranked voting instance.

Example 1. Let (w, Π) be a ranked voting instance with m candidates and m voters where voter j's first preference is j and the remaining candidates are ranked in descending order. Formally, for all $j \in [m]$, we have $\pi_j(1) = j$ and for all $i_1, i_2 \in [m] \setminus \{j\}$, we have $\pi_j^{-1}(i_1) < \pi_j^{-1}(i_2)$ iff $i_1 > i_2$.

Theorem 6. *For parameters $m \in \mathbb{Z}_{\ge 3}$, $\beta \in (\frac{1}{m}, \frac{1}{2}]$, and $\varepsilon \in (0, \beta - \frac{1}{m}]$, there is a ranked voting instance (w, Π) with m candidates such that all of the following hold:*

1. *Candidate 1 is the unique plurality winner. Her plurality score is β.*
2. $\Delta(1, (w, \Pi)) \ge 1 + \frac{m}{2\beta}(1 - \varepsilon)$.
3. $\Delta(m, (w, \Pi)) = 1 + \frac{m}{1-\varepsilon/\beta}$.

Proof. In Example 1, set $w_1 = \beta$, $w_m = \beta - \varepsilon$, $\delta := (1 - w_1 - w_m)/(m-2)$, and $w_i = \delta$ for all $i \in [m-1] \setminus \{1\}$. Then $w_m \ge 1/m$ and $0 < \delta < 1/m$. Hence, candidate 1 is the unique plurality winner and her plurality score is β. Also, every candidate has a positive plurality score, so no candidate dominates any other candidate.

Using Theorem 9, we get

$$\Delta(m, (w, \Pi)) = \max_{i=1}^{m} \sup_{U : U \rhd \Pi} \frac{\mathrm{SW}_U(i)}{\mathrm{SW}_U(m)} = \max_{i=1}^{m} \left(1 + \frac{m w_i}{w_m}\right) = 1 + \frac{m\beta}{\beta - \varepsilon}.$$

$$\begin{aligned} \Delta(1, (w, \Pi)) &\ge \sup_{U : U \rhd \Pi} \frac{\mathrm{SW}_U(m)}{\mathrm{SW}_U(1)} \\ &= 1 + \frac{m}{w_1}\left(w_m + \frac{(m-2)\delta}{2}\right) = 1 + \frac{m}{2\beta}(1 - \varepsilon). \qquad \square \end{aligned}$$

B Proof of Theorem 9

Theorem 9. *Let (w, Π) be a ranked voting instance. Let i^* and $\widehat{\imath}$ be two candidates such that $\widehat{\imath}$ doesn't dominate i^*. Let $\mathcal{U}$ be the set of all utility profiles consistent with Π. For any utility profile U, define*

$$f(U) := \frac{\mathrm{SW}_{(w,U)}(i^*)}{\mathrm{SW}_{(w,U)}(\widehat{\imath})}.$$

Then the set of ruthless utility profiles for $(i^*, \widehat{\imath})$ is exactly the set of maximizers of f over $\mathcal{U}$. Formally, for any utility profile $\widehat{U}$, we have $f(\widehat{U}) = \sup_{U \rhd \Pi} f(U)$ iff $\widehat{U}$ is ruthless for $(i^*, \widehat{\imath})$.

Assume without loss of generality that $w_j > 0$ for each voter $j \in [n]$.

Observation 14. *If U_1 and U_2 are ruthless utility profiles for $(i^*, \widehat{\imath})$, then $\mathrm{SW}_{U_1}(i) = \mathrm{SW}_{U_2}(i)$ for all $i \in \{i^*, \widehat{\imath}\}$, so $f(U_1) = f(U_2)$. Hence, either all ruthless profiles are maximizers of f, or none of them are.*

Let $\widehat{U} := (\widehat{u}_1, \ldots, \widehat{u}_n)$ be any ruthless utility profile. Since $\widehat{\imath}$ doesn't dominate i^*, for some voter $\ell \in [n]$, we get $\widehat{u}_\ell(i^*) > 0$. Hence, $\mathrm{SW}_{\widehat{U}}(i^*) > 0$. If $\mathrm{SW}_{\widehat{U}}(\widehat{\imath}) = 0$, then $f(\widehat{U}) = \infty$, so we are done. Now assume $\mathrm{SW}_{\widehat{U}}(\widehat{\imath}) >$

0. Then $f(\widehat{U}) \neq \infty$. Now it is enough to prove that f attains a finite maximum in $\mathcal{U}$, and the maximum is ruthless.

Since $\mathrm{SW}_{\widehat{U}}(\widehat{\imath}) > 0$, some voter j must have $\widehat{u}_j(\widehat{\imath}) > 0$. This can happen iff $\widehat{\imath}$ is j's favorite candidate. Let $\mathcal{U} := \{U : U \rhd \mathit{\Pi}\}$. Then for any utility profile $U \in \mathcal{U}$, we have

$$\mathrm{SW}_U(\widehat{\imath}) \geq w_j u_j(\widehat{\imath}) = w_j \max_{i=1}^{m} u_j(i) \geq \frac{w_j}{m} \sum_{i=1}^{m} u_j(i) = \frac{w_j}{m} > 0.$$

$\mathrm{SW}_U(\widehat{\imath})$ and $\mathrm{SW}_U(i^*)$ are linear in U, and hence, they are continuous functions. Since $\mathrm{SW}_U(\widehat{\imath}) > 0$ for all $U \in \mathcal{U}$, and the ratio of continuous functions is continuous, we get that f is continuous in $\mathcal{U}$. $\mathcal{U}$ is a bounded polyhedron, since $\mathcal{U} \subseteq [0,1]^{n \times m}$ and the constraint $U \rhd \mathit{\Pi}$ can be represented using $O(mn)$ linear constraints. Hence, by Weierstrass' extreme value theorem, we get that f attains a real-valued maximum on $\mathcal{U}$.

Let $U^* := (u_1^*, \ldots, u_n^*) \in \mathcal{U}$ be a maximum of f. We now use the local optimality of U^* to show that U^* must be ruthless. We do this using a case analysis over different types of voters.

Lemma 3. *If a voter j prefers i^* over $\widehat{\imath}$ (i.e., $\pi_j^{-1}(i^*) < \pi_j^{-1}(\widehat{\imath})$), then j is ruthless for $(i^*, \widehat{\imath})$ in U^*.*

Proof. Let $\widehat{U}$ be a ruthless utility profile. Suppose i^* is the k^{th}-favorite candidate of voter j. Let $\widetilde{U}$ be a utility profile where $\widetilde{u}_j := \widehat{u}_j$ and $\widetilde{u}_t := u_t^*$ for all $t \in [n] \setminus \{j\}$. Then $u_j^*(i^*) \leq 1/k$, $\widetilde{u}_j(i^*) = 1/k$, and $\widetilde{u}_j(\widehat{\imath}) = 0$. Then $\mathrm{SW}_{\widetilde{U}}(i^*) \geq \mathrm{SW}_{U^*}(i^*)$ and $\mathrm{SW}_{\widetilde{U}}(\widehat{\imath}) < \mathrm{SW}_{U^*}(\widehat{\imath})$. Hence, $f(\widetilde{U}) \geq f(U^*)$. Moreover, if $u_j^*(\widehat{\imath}) > 0$ or $u_j^*(i^*) < 1/k$, then $f(\widetilde{U}) > f(U^*)$, which is a contradiction, since U^* maximizes f. Hence, $u_j^*(\widehat{\imath}) = 0$ and $u_j^*(i^*) = 1/k$, which implies that j is ruthless in U^*. □

Lemma 4. *If a voter j prefers $\widehat{\imath}$ over i^* (i.e., $\pi_j^{-1}(\widehat{\imath}) < \pi_j^{-1}(i^*)$), then $u_j^*(\widehat{\imath}) = u_j^*(i^*)$.*

Proof. Assume $u_j^*(\widehat{\imath}) \neq u_j^*(i^*)$. Let K be the candidates who j prefers at least as much as i^* and at most as much as $\widehat{\imath}$, i.e., $K := \{i \in [m] : \pi_j^{-1}(\widehat{\imath}) \leq \pi_j^{-1}(i) \leq \pi_j^{-1}(i^*)\}$. Note that $\widehat{\imath}, i^* \in K$. Let α be the average utility for candidates in K, i.e.,

$$\alpha := \frac{1}{|K|} \sum_{i \in K} u_j^*(i).$$

Let $\widetilde{U}$ be a utility profile where $\widetilde{u}_j(i) := \alpha$ for $i \in K$, and $\widetilde{u}_j(i) := u_j^*(i)$ if $i \notin K$, and $\widetilde{u}_t := u_t^*$ for $t \in [n] \setminus \{j\}$. Then $\widetilde{U} \in \mathcal{U}$. Furthermore, we have $u_j^*(\widehat{\imath}) > \alpha > u_j^*(i^*)$. Hence, $\mathrm{SW}_{\widetilde{U}}(\widehat{\imath}) < \mathrm{SW}_{U^*}(\widehat{\imath})$ and $\mathrm{SW}_{\widetilde{U}}(i^*) > \mathrm{SW}_{U^*}(i^*)$, so $f(\widetilde{U}) > f(U^*)$. This is a contradiction, since U^* maximizes f over $\mathcal{U}$. Hence, $u_j^*(\widehat{\imath}) = u_j^*(i^*)$. □

Lemma 5. $\mathrm{SW}_{U^*}(i^*) > \mathrm{SW}_{U^*}(\widehat{\imath})$.

Proof. Since $\widehat{\imath}$ doesn't dominate i^*, there is a voter $\ell \in [n]$ who prefers i^* over $\widehat{\imath}$. Hence, by Lemma 3, we get $u_\ell^*(i^*) > u_\ell^*(\widehat{\imath}) = 0$. By Lemmas 3 and 4, we get $u_j^*(i^*) \geq u_j^*(\widehat{\imath})$ for every voter j. Hence, $\mathrm{SW}_{U^*}(i^*) > \mathrm{SW}_{U^*}(\widehat{\imath})$. □

Lemma 6. *If a voter j prefers $\widehat{\imath}$ over i^* (i.e., $\pi_j^{-1}(\widehat{\imath}) < \pi_j^{-1}(i^*)$), then j is ruthless for $(i^*, \widehat{\imath})$ in U^*.*

Proof. Let $\alpha := u_j^*(i^*) = u_j^*(\widehat{\imath})$ (by Lemma 4). Suppose there is a utility profile $\widetilde{U}$ where $\widetilde{u}_t = u_t^*$ for all $t \in [n] \setminus \{j\}$. We will decide how exactly to set $\widetilde{u}_j$ later. For now, we only impose the conditions $\widetilde{U} \rhd \Pi$ and $\widetilde{u}_j(\widehat{\imath}) = \widetilde{u}_j(i^*) = x$ for some real number x whose value we will decide later. Then $\mathrm{SW}_{\widetilde{U}}(i) = \mathrm{SW}_{U^*}(i) + w_j(x - \alpha)$ for $i \in \{i^*, \widehat{\imath}\}$. Hence,

$$f(\widetilde{U}) = \frac{\mathrm{SW}_{U^*}(i^*) + w_j(x-\alpha)}{\mathrm{SW}_{U^*}(\widehat{\imath}) + w_j(x-\alpha)} = 1 + \frac{\mathrm{SW}_{U^*}(i^*) - \mathrm{SW}_{U^*}(\widehat{\imath})}{(\mathrm{SW}_{U^*}(\widehat{\imath}) - w_j\alpha) + w_j x}.$$

Since $\widehat{\imath}$ doesn't dominate i^*, there is a voter $\ell \in [n]$ who prefers i^* over $\widehat{\imath}$. Hence, by Lemma 3, we get $\mathrm{SW}_{U^*}(\widehat{\imath}) - w_j\alpha \geq u_\ell^*(i^*) > 0$. By Lemma 5, $f(\widetilde{U})$ is a strictly decreasing function of x. If we set $\widetilde{u}_j = u_j^*$, we get $\widetilde{U} = U^*$ and $x = \alpha$. On the other hand, if we set $\widetilde{u}_j$ such that $x < \alpha$, then we get $f(\widetilde{U}) > f(U^*)$. This contradicts the maximality of U^* for f. Hence, $\alpha \leq x$.

Suppose $\widehat{\imath}$ is the k^{th}-favorite candidate of voter j. We have two cases.
Case 1: $k = 1$:
Then $\alpha = u_j^*(\widehat{\imath}) \geq 1/m$. Set $\widetilde{u}_j(i) = 1/m$ for all i. Hence, $\alpha \leq x = 1/m$. Hence, $\alpha = 1/m$. Hence, $u_j^*(i) = 1/m$ for all $i \in [m]$.
Case 2: $k > 1$: Set $\widetilde{u}_j(\pi_j(1)) = 1$ and $\widetilde{u}_j(i) = 0$ for all other i. Then $x = 0$, so $\alpha = 0$. Hence, $u_j^*(\pi_j(i)) = 0$ for all $i \geq k$.

In both cases, we see that voter j is ruthless for $(i^*, \widehat{\imath})$ in U^*. □

Lemmas 3 and 6 imply that U^* is ruthless for $(i^*, \widehat{\imath})$, which completes the proof of Theorem 9.

C Two Candidates

Consider a ranked voting instance (w, Π) with 2 candidates. We can assume without loss of generality that $n = 2$ and $\pi_i(1) = i$ for $i \in [2]$.

Lemma 7. *Let $p \in [0,1]^2$ such that $p_1 + p_2 = 1$. Then*

$$\begin{aligned}\Delta(p,(w,\Pi))^{-1} &= \min_{i\in[2]}\left(p_i + (1-p_i)\frac{1-w_i}{1+w_i}\right)\\ &= \min\left(\frac{1-w_1}{1+w_1} + \frac{2w_1}{1+w_1}p_1, 1 - \frac{2-2w_1}{2-w_1}p_1\right).\end{aligned}$$

Proof. For any utility profile U, we have $\mathrm{SW}_U(1) + \mathrm{SW}_U(2) = 1$.

$$\begin{aligned}\Delta(p,(w,\Pi))^{-1} &= \inf_{U:U\rhd\Pi} \frac{\sum_{i\in[2]} p_i \,\mathrm{SW}_U(i)}{\max_{i\in[2]} \mathrm{SW}_U(i)}\\ &= \min_{i\in[2]} \inf_{U:U\rhd\Pi} \frac{\sum_{j\in[2]} p_j \,\mathrm{SW}_U(j)}{\mathrm{SW}_U(i)}\\ &= \min_{i\in[2]} \inf_{U:U\rhd\Pi} \frac{p_i \,\mathrm{SW}_U(i) + (1-p_i)(1-\mathrm{SW}_U(i))}{\mathrm{SW}_U(i)}\\ &= \min_{i\in[2]} \inf_{U:U\rhd\Pi} \left(p_i + (1-p_i)\left(\frac{1}{SW_U(i)} - 1\right)\right)\\ &= \min_{i\in[2]} \left(p_i + (1-p_i)\left(\frac{1}{\sup_{U:U\rhd\Pi} SW_U(i)} - 1\right)\right).\end{aligned}$$

For $i \in [2]$ and $U \rhd \Pi$, we have $u_i(i) \in [1/2, 1]$ and $u_{3-i}(i) \in [0, 1/2]$. Hence,

$$\begin{aligned}\sup_{U:U\rhd\Pi} \mathrm{SW}_U(i) &= \sup_{U\rhd\Pi}(w_i u_i(i) + (1-w_i)u_{3-i}(i))\\ &= w_i + (1-w_i)\frac{1}{2} = \frac{1+w_i}{2}.\end{aligned}$$

Hence,

$$\begin{aligned}\Delta(p,(w,\Pi))^{-1} &= \min_{i\in[2]}\left(p_i + (1-p_i)\left(\frac{2}{1+w_i} - 1\right)\right)\\ &= \min\left(p_1 + (1-p_1)\frac{1-w_1}{1+w_1}, (1-p_1) + p_1\frac{w_1}{2-w_1}\right)\\ &= \min\left(\frac{1-w_1}{1+w_1} + \frac{2w_1}{1+w_1}p_1, 1 - \frac{2-2w_1}{2-w_1}p_1\right).\end{aligned}$$

□

Corollary 1. *For any $i \in [2]$, we get $\Delta(i,(w,\Pi)) = 2/w_i - 1$.*

Corollary 1 tells us that the plurality winner achieves optimal deterministic distortion. Since the plurality winner has weight at least $1/2$, her distortion must be at most 3. Moreover, for $w_1 = w_2 = 1/2$, we get that both candidates have distortion 3.

Lemma 8. *The optimal randomized distortion is achieved by* $p^* \in [0,1]^2$, *where*

$$p_1^* = \frac{w_1(2-w_1)}{1+2w_1w_2}, \qquad \Delta(p^*, (w, \Pi)) = 1 + 2w_1w_2.$$

Proof. By Lemma 7, we get

$$\Delta(p, (w, \Pi))^{-1} = \min\left(\frac{1-w_1}{1+w_1} + \frac{2w_1}{1+w_1}p_1, 1 - \frac{2-2w_1}{2-w_1}p_1\right).$$

To minimize distortion, we need to maximize $\Delta(p, (w, \Pi))^{-1}$, which is the minimum of an increasing function and a decreasing function of p_1. Let us compute their point of intersection.

$$\begin{aligned}
&\frac{1-w_1}{1+w_1} + \frac{2w_1}{1+w_1}p_1 = 1 - \frac{2-2w_1}{2-w_1}p_1 \\
\iff\ & p_1\left(\frac{2w_1}{1+w_1} + \frac{2(1-w_1)}{2-w_1}\right) = 1 - \frac{1-w_1}{1+w_1} \\
\iff\ & 2p_1\frac{1+2w_1w_2}{(1+w_1)(2-w_1)} = \frac{2w_1}{1+w_1} \\
\iff\ & p_1 = \frac{w_1(2-w_1)}{1+2w_1w_2}.
\end{aligned}$$

Since $(1+2w_1w_2) - w_1(2-w_1) = 1 - w_1^2 \geq 0$, we get that $p_1 \in [0,1]$. Hence, $\Delta(p, (w, \Pi))^{-1}$ is maximized at $p_1 = w_1(2-w_1)/(2+w_1w_2)$. Moreover,

$$\Delta(p, (w, \Pi))^{-1} = \frac{1-w_1}{1+w_1} + \frac{2w_1}{1+w_1}\frac{w_1(2-w_1)}{1+2w_1(1-w_1)} = \frac{1}{1+2w_1w_2}. \qquad \square$$

By Lemma 8, the optimal randomized distortion is at most $1+2w_1(1-w_1) \leq 3/2$. This bound is tight for $w_1 = w_2 = 1/2$.

D Introduction to Distortion

We consider the single-winner election problem, i.e., n voters want to elect a single winner from m candidates. For any non-negative integer t, let $[t] := \{1, 2, \ldots, t\}$.

D.1 Utility and Welfare

We can take a utilitarian approach to the single-winner election problem: assume the j^{th} voter has a utility function $u_j : [m] \to \mathbb{R}$. Here $u_j(i)$ is a number that tells us how much voter j likes candidate i, or how much voter j benefits from candidate i's election.

Each voter j has a weight $w_j \in \mathbb{R}_{\geq 0}$ which is a measure of the importance of voter j compared to the other voters. Assume without loss of generality that $\sum_{j=1}^{n} w_j = 1$. Usually we assume that all voters are equal (called *unweighted voting*), and so their weights are the same, but there are situations where we may want to use weighted voting, e.g., when stakeholders of a company want to make a joint decision, they may want to give more weight to voters with higher stake in the company.

Studying weighted voting also gives us insight into unweighted voting. This is because if all the weights are rational, and each voter j has weight n_j/D for $n_j \in \mathbb{Z}$, then we can replace each voter j with n_j clones of that voter, each having weight $1/D$ to get an equivalent unweighted voting instance. Since $\mathbb{Q}$ is dense in $\mathbb{R}$, having rational weights is not a restrictive assumption.

Any single-winner election can be fully described by the pair (w, U), where $w \in [0, 1]^n$ is the vector of voters' weights and U is the voters' *utility profile*, i.e., U is a list $(u_1, u_2, \ldots, u_n)$, where $u_j : [m] \to \mathbb{R}$ is voter j's utility function. We often represent U as an n-by-m matrix of real numbers, where the entry in the j^{th} row and i^{th} column is $u_j(i)$.

In a utility profile U on m candidates, the *social welfare* of candidate i is defined as

$$\mathrm{SW}_{(w,U)}(i) := \sum_{j=1}^{n} w_j u_j(i).$$

When w is clear from context, we write SW_U instead of $\mathrm{SW}_{(w,U)}$. Intuitively, $\mathrm{SW}_U(i)$ is the net gain to society if candidate i is elected as winner. Hence, a natural objective is to elect the candidate who maximizes the social welfare, known as the *social-welfare-optimal* (SWO) candidate. Given (w, U) as input, finding the SWO candidate is trivial and can be done in $\Theta(mn)$ time.

Often, we constrain voters' utility functions in some way to ensure that they can't freely scale them up to increase their relative importance. In this document, we consider normalized utilities, i.e., utility functions $u : [m] \to \mathbb{R}$ where $u(i) \geq 0$ for all $i \in [m]$ and $\sum_{i=1}^{m} u(i) = 1$. (Other classes of utility functions are also commonly studied, like unit-range functions, binary functions, and metric utilities.)

D.2 Ordinal Preferences and Distortion

In most real-world scenarios, it is hard for voters to accurately convey their utilities exactly (based on studies in behavioral economics and psychology) (I don't have direct citations for this, but [2] claim this in their paper). On the other hand, it's easy for voters to compare candidates to each other. Hence, voting systems often ask voters to rank the candidates.

Therefore, instead of observing the utility profile U, we observe a *preference profile* $\mathit{\Pi} := (\pi_1, \pi_2, \ldots, \pi_m)$. Here $\pi_j : [m] \to [m]$ is a bijection where $\pi_j(i)$ is the j^{th} voter's i^{th} favorite candidate. We often represent $\mathit{\Pi}$ as an n-by-m matrix, where the j^{th} row and i^{th} column is $\pi_j(i)$. We say that utility profile U is *consistent* with preference profile $\mathit{\Pi}$ (denoted as $U \rhd \mathit{\Pi}$) if for each voter j, we have $u_j(\pi_j(1)) \geq u_j(\pi_j(2)) \geq \ldots \geq u_j(\pi_j(m))$.

We receive as input the pair $(w, \mathit{\Pi})$ (called a *ranked voting instance*), where $w \in [0,1]^n$ is the vector of voters' weights. The voters' utility profile U (which is consistent with $\mathit{\Pi}$) is hidden from us. We would like to pick a candidate i such that $\mathrm{SW}_{(w,U)}(i)$ is maximized. However, as the following example shows, finding the SWO candidate is impossible without knowing U.

Example 2. Consider a ranked voting instance $(w, \mathit{\Pi})$, where $w_1 = w_2 = 1/2$ and $\mathit{\Pi} = \left(\begin{smallmatrix} 1 & 2 \\ 2 & 1 \end{smallmatrix}\right)$ (i.e., 2 voters and 2 candidates, voter 1 prefers candidate 1, voter 2 prefers candidate 2). Consider utility profiles U_1 and U_2:

$$U_1 := \begin{pmatrix} \frac{1}{2} + \varepsilon & \frac{1}{2} - \varepsilon \\ 0 & 1 \end{pmatrix} \qquad\qquad U_2 := \begin{pmatrix} 1 & 0 \\ \frac{1}{2} - \varepsilon & \frac{1}{2} + \varepsilon \end{pmatrix}.$$

Here $0 < \varepsilon \ll 1/2$. Then both U_1 and U_2 are consistent with $\mathit{\Pi}$ and $\mathrm{SW}_{U_1}(1) = \mathrm{SW}_{U_2}(2) = (1+2\varepsilon)/4$ and $\mathrm{SW}_{U_1}(2) = \mathrm{SW}_{U_2}(1) = (3-2\varepsilon)/4$. Hence, whichever candidate we pick, there is a consistent utility profile for which the candidate we picked has social welfare $\approx 1/4$ and the other candidate has social welfare $\approx 3/4$.

Since finding the SWO candidate is impossible without knowing U, we try finding a candidate whose social welfare is within a small factor of the SWO candidate's social welfare for all consistent utility profiles. This factor is called *distortion.* Formally, for input $(w, \mathit{\Pi})$, the distortion of candidate i is defined as

$$\mathit{\Delta}(i, (w, \mathit{\Pi})) := \sup_{U : U \rhd \mathit{\Pi}} \frac{\max_{i^* \in [m]} \mathrm{SW}_{(w,U)}(i^*)}{\mathrm{SW}_{(w,U)}(i)}.$$

Our aim is to find a candidate of low distortion.

We may also pick a candidate randomly, in which case we extend the definition of distortion to distributions over candidates. Formally, let $p \in [0,1]^m$ be a vector such that $\sum_{i=1}^m p_i = 1$. Suppose we pick candidate i with probability p_i for all $i \in [m]$. Then the expected value of our chosen candidate's social welfare is $\sum_{i=1}^m p_i \operatorname{SW}_{(w,U)}(i)$. Hence, we define the distortion of p for the input $(w, \mathit{\Pi})$ as

$$\mathit{\Delta}(p, (w, \mathit{\Pi})) := \sup_{U : U \rhd \mathit{\Pi}} \frac{\max_{i^* \in [m]} \operatorname{SW}_{(w,U)}(i^*)}{\sum_{i \in [m]} p_i \operatorname{SW}_{(w,U)}(i)}.$$

An algorithm that takes input $(w, \mathit{\Pi})$ and outputs a candidate is called a *voting rule*. The distortion of a (randomized) voting rule on $(w, \mathit{\Pi})$ is the distortion of (the distribution of) the voting rule's output on $(w, \mathit{\Pi})$.

Conventions to ensure that suprema of ratios are well-defined:

1. Define ∞ to be larger than any real number.
2. For any $a > 0$, let $a/0 := \infty$ and $0/0 := 0$.
3. For $f : X \to \mathbb{R} \cup \{\infty\}$, if $f(x) = \infty$ for some $x \in X$, or if f is not upper-bounded by a real-number, then $\sup_{x \in X} f(x) := \infty$.