

On the Distortion of Voting with Multiple Representative Candidates*

Yu Cheng
Duke University

Shaddin Dughmi
University of Southern California

David Kempe
University of Southern California

Abstract

We study positional voting rules when candidates and voters are embedded in a common metric space, and cardinal preferences are naturally given by distances in the metric space. In a positional voting rule, each candidate receives a score from each ballot based on the ballot’s rank order; the candidate with the highest total score wins the election. The *cost* of a candidate is his sum of distances to all voters, and the distortion of an election is the ratio between the cost of the elected candidate and the cost of the optimum candidate. We consider the case when candidates are *representative* of the population, in the sense that they are drawn i.i.d. from the population of the voters, and analyze the expected distortion of positional voting rules.

Our main result is a clean and tight characterization of positional voting rules that have constant expected distortion (independent of the number of candidates and the metric space). Our characterization result immediately implies constant expected distortion for Borda Count and elections in which each voter approves a constant *fraction* of all candidates. On the other hand, we obtain super-constant expected distortion for Plurality, Veto, and approving a constant *number* of candidates. These results contrast with previous results on voting with metric preferences: When the candidates are chosen adversarially, all of the preceding voting rules have distortion linear in the number of candidates or voters. Thus, the model of representative candidates allows us to distinguish voting rules which seem equally bad in the worst case.

1 Introduction

In light of the classic impossibility results for axiomatic approaches to social choice (Arrow 1951) and voting (Gibbard 1973; Satterthwaite 1975), a fruitful approach has been to treat voting as an implicit optimization problem of finding the “best” candidate for the population in aggregate (Boutilier et al. 2015; Caragiannis and Procaccia 2011; Procaccia 2010; Procaccia and Rosenschein 2006). Using this approach, voting systems can be compared based on how much they *distort* the outcome, in the sense of leading to the election of suboptimal candidates. A particularly natural optimization objective is the *sum of distances* between

voters and the chosen candidate in a suitable metric space (Anshelevich 2016; Anshelevich, Bhardwaj, and Postl 2015; Anshelevich and Postl 2016; Goel, Krishnaswamy, and Munagala 2017). The underlying assumption is that the closer a candidate is to a voter, the more similar their positions on key questions are. Because proximity implies that the voter would benefit from the candidate’s election, voters will rank candidates by increasing distance, a model known as *single-peaked preferences* (Black 1948; Downs 1957; Black 1958; Moulin 1980; Merrill and Grofman 1999; Barberà, Gul, and Stacchetti 1993; Richards, Richards, and McKay 1998; Barberà 2001).

Even in the absence of strategic voting, voting systems can lead to high distortion in this setting, because they typically allow only for communication of *ordinal* preferences¹, i.e., rankings of candidates (Boutilier and Rosenschein 2016). In a beautiful piece of recent work, (Anshelevich, Bhardwaj, and Postl 2015) showed that this approach can draw very clear distinctions between voting systems: some voting systems (in particular, Copeland and related systems) have distortion bounded by a small constant, while most others (including Plurality, Veto, *k*-approval, and Borda Count) have unbounded distortion, growing linearly in the number of voters or candidates.

The examples giving bad distortion typically have the property that the candidates are not “representative” of the voters. (Anshelevich, Bhardwaj, and Postl 2015) showed more positive results when there are no near-ties for first place in any voter’s ranking. (Cheng, Dughmi, and Kempe 2017) proposed instead a model of representativeness in which the candidates are drawn randomly from the population of voters; under this model, they showed smaller constant distortion bounds than the worst-case bounds for majority voting with $n = 2$ candidates. (Cheng, Dughmi, and Kempe 2017) left as an open question the analysis of the distortion of voting systems for $n \geq 3$ representative candidates.

In the present work, we study the distortion of *positional voting systems* with $n \geq 3$ representative candidates. Informally (formal definitions of all concepts are given in Sec-

*Due to space constraints, this version omits several important proofs. The full version of this paper can be found on arXiv. Copyright © 2018, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹Of course, it is also highly questionable that voters would be able to quantify distances in a metric space sufficiently accurately, in particular given that the metric space is primarily a modeling tool rather than an actual concrete object.

tion 2), a positional voting system is one in which each voter writes down an ordering of candidates, and the system assigns a score to each candidate based solely on his² position in the voter’s ordering. The map from positions to scores is known as the *scoring rule* of the voting system, and for n candidates is a function $g_n : \{0, \dots, n-1\} \rightarrow \mathbb{R}_{\geq 0}$. The total score of a candidate is the sum of scores he obtains from all voters, and the winner is the candidate with maximum total score. The most well-known explicitly positional voting system is Borda Count (de Borda 1784), in which $g_n(i) = n - i$ for all i . Many other systems are naturally cast in this framework, including Plurality (in which voters give 1 point to their first choice only) and Veto (in which voters give 1 point to all but their last choice).

In analyzing positional voting systems, we assume that voters are not strategic, i.e., they report their true ranking of candidates based on proximity in the metric space. This is in keeping with the line of work on analyzing the distortion of social choice functions, and avoids issues of game-theoretic modeling and equilibrium existence or selection (see, e.g., (Feldman, Fiat, and Golomb 2015)) which are not our focus.

As our main contribution, we characterize when a positional voting system is guaranteed to have constant distortion, regardless of the underlying metric space of voters and candidates, and regardless of the number n of candidates that are drawn from the voter distribution. The characterization relies almost entirely on the “limit voting system.” By normalizing both the scores and the candidate index to lie in $[0, 1]$ (we associate the i^{th} out of n candidates with his quantile $\frac{i}{n-1} \in [0, 1]$), we can take a suitable limit g of the scoring functions g_n as $n \rightarrow \infty$.

Our main result (Corollary 3.2 in Section 3) states the following: (1) If g is not constant on the open interval $(0, 1)$, then the voting system has constant distortion. (2) If g is a constant other than 1 on the open interval $(0, 1)$, then the voting system does not have constant distortion. The only remaining case is when $g \equiv 1$ on $(0, 1)$. In that case, the rate of convergence of g_n to g matters, and a precise characterization is given by Theorem 3.1.

As direct applications of our main result, we obtain that Borda Count and k -approval for $k = \Theta(n)$ representative candidates have constant distortion; on the other hand, Plurality, Veto, the Nauru Dowdall method (see Section 2), and k -approval for $k = O(1)$ have super-constant distortion. In fact, it is easy to adapt the proof of Theorem 3.1 to show that the distortion of Plurality, Veto, and $O(1)$ -approval, even with representative candidates, is $\Omega(n)$.

Our results provide interesting contrasts to the results of (Anshelevich, Bhardwaj, and Postl 2015). Under adversarial candidates, all of the above-mentioned voting rules have distortion $\Omega(n)$; the focus on representative candidates allowed us to distinguish the performance of Borda Count and $\Theta(n)$ -approval from that of the other voting systems. Thus, an analysis in terms of representative candidates allows us to draw distinctions between voting systems which in a worst-case setting seem to be equally bad.

²For consistency, we always use male pronouns for candidates and female pronouns for voters.

As a by-product of the proof of our main theorem, in Lemma 3.3, we show that *every* voting system (positional or otherwise) has distortion $O(n)$ with representative candidates. Combined with the lower bound alluded to above, this exactly pins down the distortion of Plurality, Veto, and $O(1)$ -approval with representative candidates to $\Theta(n)$. For Veto, this result also contrasts with the worst-case bound of (Anshelevich, Bhardwaj, and Postl 2015), which showed that the distortion can grow unboundedly even for $n = 3$ candidates.

2 Preliminaries

2.1 Voters, Metric Space, and Preferences

The voters/candidates are embedded in a closed metric space (Ω, d) , where $d_{\omega, \omega'}$ is the distance between points $\omega, \omega' \in \Omega$. The distance captures the dissimilarity in opinions between voters (and candidates) — the closer two voters or candidates are, the more similar they are. The distribution of voters in Ω is denoted by the (measurable) density function q_ω . We allow for q to have point masses.³ Unless there is no risk of confusion, we will be careful to distinguish between a location $\omega \in \Omega$ and a specific voter j or candidate i at that location. We apply d equally to locations/voters/candidates.

We frequently use the standard notion of a *ball* $B(\omega, r) := \{\omega' \mid d_{\omega, \omega'} \leq r\}$ in a metric space. For balls (and other sets) B , we write $q_B := \int_{\omega \in B} q_\omega d\omega$.

An election is run between $n \geq 2$ candidates according to rules defined in Section 2.3. The n candidates are assumed to be representative of the population, in the sense that their locations are drawn i.i.d. from the distribution q of voters.

Each voter ranks the n candidates i by non-decreasing distance from herself in (Ω, d) . Ties are broken arbitrarily, but consistently⁴, meaning that all voters at the same location have the same ranking. We denote the ranking of a voter j or a location ω over candidates i by $\pi_j(i)$ or $\pi_\omega(i)$. The distance-based ranking assumption means that $\pi_\omega(i) < \pi_\omega(i')$ implies that $d_{\omega, i} \leq d_{\omega, i'}$ and $d_{\omega, i} < d_{\omega, i'}$ implies that $\pi_\omega(i) < \pi_\omega(i')$. As mentioned in the introduction, we assume that voters are not strategic; i.e., they express their true ranking of candidates based on proximity in the metric space.

2.2 Social Cost and Distortion

Candidates are “better” if they are closer to voters on average. The *social cost* of a candidate (or location) i is $c_i = \int_\omega d_{\omega, i} q_\omega d\omega$.

The socially optimal candidate among the set C of candidates running is denoted by $o(C) := \operatorname{argmin}_{i \in C} c_i$. The overall optimal location is denoted by $\hat{o} \in \operatorname{argmin}_{\omega \in \Omega} c_\omega$, which is any 1-median of the metric space. (If there are multiple optimal locations, consider one of them fixed arbitrarily.) The argmin always exists, because the metric space is assumed to be closed, and the cost function is continuous and bounded below by 0. Note that it is not necessary that there be any voters located at \hat{o} .

³Since the continuum model allows for point masses, it subsumes finite sets of voters. Changing all our results to finite or countable voter sets is merely cosmetic.

⁴Our results do not depend on specific tie breaking rules.

Based on the votes, a voting system will determine a *winner* $w(C)$ for the set C of candidates, who will often be different from $o(C)$. The *distortion* measures how much worse the winner is than the optimum $D(C) = \frac{c_{w(C)}}{c_{o(C)}}$. We are interested in the *expected distortion* of positional voting systems under i.i.d. random candidates, i.e., $\mathbb{E}_{C \sim \mathcal{Q}} [D(C)]$.

Our distortion bounds are achieved by lower-bounding $c_{o(C)} \geq c_{\hat{o}}$. A particularly useful quantity in this context is the fraction of voters outside a ball of radius r around \hat{o} , which we denote by $H(r) := 1 - q_{B(\hat{o}, r)}$. The following lemma captures some useful simple facts that we use:

Lemma 2.1. 1. For any candidate or location i ,

$$c_i \leq c_{\hat{o}} + d_{i, \hat{o}}. \quad (1)$$

2. The cost of any candidate or location i can be written as

$$c_i = \int_0^\infty (1 - q_{B(i, r)}) dr. \quad (2)$$

3. For all $r \geq 0$, the cost of the optimum location \hat{o} is lower-bounded by

$$c_{\hat{o}} \geq rH(r). \quad (3)$$

Proof. 1. The proof of the first inequality simply applies the triangle inequality under the integral: $c_i = \int_\omega d_{\omega, i} q_\omega d\omega \leq \int_\omega (d_{\omega, \hat{o}} + d_{i, \hat{o}}) q_\omega d\omega = c_{\hat{o}} + d_{i, \hat{o}}$.

2. For the second equation, observe that $c_i = \mathbb{E}_{\omega \sim \mathcal{Q}} [d_{i, \omega}]$, and the expectation of any non-negative random variable X can be rewritten as $\mathbb{E}[X] = \int_0^\infty \Pr[X \geq x] dx$.

3. For the third inequality, we apply the previous part with $i = \hat{o}$, and lower bound $\int_0^\infty \Pr_{\omega \sim \mathcal{Q}} [d_{\hat{o}, \omega} \geq x] dx \geq \int_0^r \Pr_{\omega \sim \mathcal{Q}} [d_{\hat{o}, \omega} \geq r] dx + \int_r^\infty 0 dx = r \cdot H(r)$. ■

2.3 Positional Voting Systems and Scoring Rules

We are interested in *positional voting systems*. Such systems are based on *scoring rules*: voters give a ranking of candidates, and with each position is associated a score.

Definition 2.1 (Scoring Rule). A *scoring rule* for n candidates is a non-increasing function $g_n : \{0, \dots, n-1\} \rightarrow [0, 1]$ with $g_n(0) = 1$ and $g_n(n-1) = 0$.

Definition 2.2 (Positional Voting System). A *positional voting system* is a sequence of scoring rules g_n , one for each number of candidates $n = 1, 2, \dots$

The interpretation of g_n is that if voter j puts a candidate i in position $\pi_j(i)$ on her ballot, then i obtains $g_n(\pi_j(i))$ points from j . The total score of candidate i is then $\sigma(i) = \int_\omega g_n(\pi_\omega(i)) q_\omega d\omega$. The winning candidate is one with highest total score, i.e., for a set C of n candidates, $w(C) \in \operatorname{argmax}_{i \in C} \sigma(i)$; again, ties are broken arbitrarily, and our results do not depend on tie breaking.

The restriction to monotone non-increasing scoring rules is standard when studying positional voting systems. One justification is that in any positional voting system violating this restriction, truth-telling is a dominated strategy, rendering such a system uninteresting for most practical purposes. Given this restriction, the assumption that $g_n(0) = 1$ and

$g_n(n-1) = 0$ is without loss of generality, because a score-based rule is invariant under affine transformations.

Next, we want to capture the notion that a positional voting system is “consistent” as we vary the number of candidates n . Intuitively, we want to exclude contrived voting systems such as “If the number of candidates is even, then use Borda Count; otherwise use Plurality voting.” This is captured by the following definition.

Definition 2.3 (Consistency). Let \mathcal{V} be a positional voting system with scoring rules $\{g_n \mid n \in \mathbb{N}\}$. We say that \mathcal{V} is consistent if there exists a function $g : \mathbb{Q} \cap [0, 1] \rightarrow [0, 1]$ such that for each rational quantile $x \in [0, 1]$ and accuracy parameter $\epsilon > 0$, there exists a threshold n_0 such that $g_n(\lfloor x(n-1) \rfloor) \geq g(x) - \epsilon$ and $g_n(\lceil x(n-1) \rceil) \leq g(x) + \epsilon$ for all $n \geq n_0$. We call g the limit scoring rule of \mathcal{V} .

Intuitively, this definition says that the sequence of scoring rules g_n is consistent with a single scoring rule g in the limit. Using the fact that g_n is monotone non-increasing for each n , it can be shown that g is also monotone non-increasing. We note that g_n converges pointwise to g in a precise and natural sense. Formally, when $x \in [0, 1]$ is rational, there exists an infinite sequence of integers n with $\lfloor x(n-1) \rfloor = \lceil x(n-1) \rceil = x(n-1)$, and consistency implies that $g(x)$ must equal the limit of $g_n(x(n-1))$ for that sequence of values of n . Therefore the limit scoring rule g is uniquely defined if it exists.

All positional voting systems we are aware of are consistent according to Definition 2.3.

Example 2.1. To illustrate the notion of a consistent positional voting system, consider the following examples, encompassing most well-known scoring rules.

- In *Plurality voting* with n candidates, $g_n(0) = 1$ and $g_n(k) = 0$ for all $k > 0$. The limit scoring rule is $g(0) = 1$ and $g(x) = 0$ for all $x > 0$.
- In *Veto voting* with n candidates, $g_n(k) = 1$ for all $k < n-1$ and $g_n(n-1) = 0$. The limit scoring rule is $g(x) = 1$ for all $x < 1$ and $g(1) = 0$.
- In *k-approval voting* with constant k , we have $g_n(k') = 1$ for $k' \leq \min(k-1, n-1)$, and $g_n(k') = 0$ for all other k' . The limit scoring rule is $g(0) = 1$ and $g(x) = 0$ for all $x > 0$, i.e., the same as for *Plurality voting*. (This relies on k being constant, or more generally, $k = o(n)$.)
- In *k-approval voting* with linear k , there exists a constant $\gamma \in (0, 1)$ with $g_n(k) = 1$ for all $k \leq \gamma n$, and $g_n(k) = 0$ for all larger k . The limit scoring rule is $g(x) = 1$ for $x \leq \gamma$ and $g(x) = 0$ for $x > \gamma$.
- The *Borda voting rule* has $g_n(k) = 1 - \frac{k}{n-1}$ (after normalization). The limit scoring rule is $g(x) = 1 - x$.
- The *Dowdall method* used in *Nauru* (Fraenkel and Grofman 2014; Reilly 2002) has $g_n(k) = 1/(k+1)$. After normalization, the rule becomes $g_n(k) = \frac{1}{n-1} \cdot (\frac{n}{k+1} - 1)$. The limit scoring rule is $g(0) = 1$ and $g(x) = 0$ for all $x > 0$, i.e., the same as for *Plurality voting*. This is because for every constant quantile x , the score of the candidate at x is $\frac{1}{n-1} (\frac{1}{x} - 1) \xrightarrow{n \rightarrow \infty} 0$.

3 The Main Characterization Result

In this section, we state and prove our main theorem, characterizing positional voting systems with constant distortion.

Theorem 3.1. *Let \mathcal{V} be a positional voting system with a sequence g_n of scoring rules for $n = 1, 2, \dots$. Then, \mathcal{V} has constant expected distortion if and only if there exist constants n_0 and $y \in (0, 1)$ such that for all $n \geq n_0$,*

$$\begin{aligned} y \cdot \sum_{k=0}^{\lceil y(n-1) \rceil - 1} (g_n(k) - g_n(\lceil y(n-1) \rceil)) \\ > (1-y) \cdot \sum_{k=n-\lceil y(n-1) \rceil}^{n-1} (1 - g_n(k)). \end{aligned} \quad (4)$$

We prove Theorem 3.1 in Sections 3.1 (sufficiency) and 3.2 (necessity). Condition (4) is quite unwieldy. In most cases of practical interest, we can use Corollary 3.2.

Corollary 3.2. *Let \mathcal{V} be a consistent positional voting system with limit scoring rule g .*

1. *If g is not constant on the open interval $(0, 1)$, then \mathcal{V} has constant expected distortion.*
2. *If g is equal to a constant other than 1 on the open interval $(0, 1)$, then \mathcal{V} does not have constant expected distortion.*

The constant in Theorem 3.1 and Corollary 3.2 depends on \mathcal{V} , but not on the metric space or the number of candidates. Corollary 3.2 has the advantage of determining constant expected distortion only based on the limit scoring rule g . The only case when it does not apply is when $g(x) = 1$ for all $x \in [0, 1)$. In that case, the higher complexity of Theorem 3.1 is indeed necessary to determine whether \mathcal{V} has constant distortion. Fortunately, Veto voting is the only rule of practical importance for which $g(x) \equiv 1$ on $[0, 1)$, and it is easily analyzed.

Before presenting the proofs, we apply the characterization to the positional voting systems from Example 2.1. Using the limit scoring rules derived in Example 2.1, Corollary 3.2 implies constant expected distortion for Borda Count and k -approval with linear $k = \Theta(n)$, and super-constant expected distortion for Plurality, k -approval with $k = o(n)$, and the Dowdall method.

This leaves Veto voting, for which it is easy to apply Theorem 3.1 directly. Because $g_n(k) = 1$ for all $k < n - 1$, for any constant $y < 1$ and large enough n , the left-hand side of (4) is 0, while the right-hand side is positive. Hence, (4) can never be satisfied for sufficiently large n , implying super-constant expected distortion. The proof easily generalizes to show that when voters can veto $o(n)$ candidates, the distortion is super-constant.

3.1 Sufficiency

In this section, we prove that condition (4) suffices for constant distortion. First, because of the monotonicity of g_n , if (4) holds for $y \in (0, 1)$, then it also holds for all $y' \in [y, 1)$. Now, the high-level idea of the proof is the following: we define a radius \hat{r} large enough so that the ball $B(\hat{d}, \hat{r})$ around the socially optimal location \hat{d} contains a very large (but still

constant) fraction y of all voters, such that y satisfies (4). If the number of candidates n is large enough (a large constant), standard Chernoff bounds ensure that as $r \geq \hat{r}$ grows large, most candidates who are running will be from inside $B(\hat{d}, r)$. In turn, if many candidates inside $B(\hat{d}, r)$ are running, all candidates outside $B(\hat{d}, 3r)$ are very far down on almost everyone's ballot, and therefore cannot win. In particular, Inequality (4) implies that the total score of an average candidate in $B(\hat{d}, r)$ exceeds the maximum possible total score of a candidate outside $B(\hat{d}, 3r)$. This allows us to bound the expected distortion in terms of the cost of \hat{d} .

The case of small n is much easier, since we can treat n as a constant. In that case, the following lemma is sufficient.

Lemma 3.3. *If n candidates are drawn i.i.d. at random from q , the expected distortion is at most $n + 1$.*

Proof. The proof illustrates some of the key ideas that will be used later in the more technical proof for a large number of candidates. We want to bound

$$\begin{aligned} \mathbb{E}_C [c_{w(C)}] &\stackrel{(1)}{\leq} c_{\hat{d}} + \mathbb{E}_C [d_{\hat{d}, w(C)}] \\ &= c_{\hat{d}} + \int_0^\infty \Pr_C [d_{\hat{d}, w(C)} \geq r] dr. \end{aligned}$$

In order for a candidate at distance at least r from \hat{d} to win, it is necessary that at least one such candidate be running. By a union bound over the n candidates, the probability of this event is at most $\Pr_C [d_{\hat{d}, w(C)} \geq r] \leq nH(r)$, so

$$\mathbb{E}_C [c_{w(C)}] \leq c_{\hat{d}} + n \int_0^\infty H(r) dr \stackrel{(2)}{=} (n+1)c_{\hat{d}}.$$

Lower-bounding the cost of the optimum candidate from C in terms of the overall best location \hat{d} , the expected distortion is

$$\begin{aligned} \mathbb{E}_C \left[\frac{c_{w(C)}}{c_{o(C)}} \right] &\leq \mathbb{E}_C \left[\frac{c_{w(C)}}{c_{\hat{d}}} \right] = \frac{1}{c_{\hat{d}}} \mathbb{E}_C [c_{w(C)}] \\ &\leq \frac{1}{c_{\hat{d}}} \cdot (n+1)c_{\hat{d}} = n+1. \quad \blacksquare \end{aligned}$$

In preparation for the case of large n , we begin with the following technical lemma, which shows that whenever (4) holds, it will also hold when the terms on the left-hand side are "shifted," and the right-hand side can be increased by a factor of 2 (or, for that matter, any constant factor).

Lemma 3.4. *Assume that there exist $y \in (0, 1)$ and n_0 such that (4) holds. Then, there exists $z_0 \in (\frac{1}{2}, 1)$ such that for all $n \geq n_0$, all $z \geq z_0$, and all integers $0 \leq m \leq (1-z) \cdot n$,*

$$\begin{aligned} z \cdot \sum_{k=0}^{\lceil z \cdot (n-1) \rceil - 1} (g_n(m+k) - g_n(m + \lceil z \cdot (n-1) \rceil)) \\ > 2(1-z) \cdot \sum_{k=n-\lceil z \cdot (n-1) \rceil}^{n-1} (1 - g_n(k)). \end{aligned} \quad (5)$$

We now flesh out the details of the construction. By Lemma 3.4, there exists $z_0 \in (\frac{1}{2}, 1)$ and n_0 such that (5) holds for all $z \geq z_0$, all $n \geq n_0$, and all integers $0 \leq m \leq (1-z) \cdot n$. For simplicity of notation, write $z := z_0$, and let $\mu := (1 - \frac{1}{e}) + \frac{1}{e} \cdot z \in (z, 1)$ and $\hat{n} := \max(n_0, \frac{1}{1-z})$.

Notice that μ, z, \hat{n} only depend on \mathcal{V} , but not on the metric space or number of voters.

Let $\hat{r} := \inf\{r \mid q_{B(\hat{o}, r)} \geq \mu\}$, so that $q_{B(\hat{o}, \hat{r})} \geq \mu$, and $q_{\{w \mid d_{\hat{o}, w} \geq \hat{r}\}} \geq 1 - \mu$. (Both inequalities hold with equality unless there is a discrete point mass at distance \hat{r} from \hat{o} .)

Consider any $r \geq \hat{r}$ and write $T := B(\hat{o}, 3r)$ and $S := B(\hat{o}, r)$. When n candidates are drawn i.i.d. from \mathcal{q} , the expected fraction of candidates drawn from outside of S is exactly $H(r) \leq 1 - \mu$. Let \mathcal{E}_r be the event that more than $(1 - z)n$ candidates are from outside S . Lemma 3.5 uses Chernoff bounds and the definitions of the parameters to show that \mathcal{E}_r happens with sufficiently small probability; Lemma 3.6 then shows that unless \mathcal{E}_r happens, the distortion is constant.

Lemma 3.5. $\Pr[\mathcal{E}_r] \leq \frac{e}{1-z} \cdot H(r)$.

Lemma 3.6. *Whenever \mathcal{E}_r does not happen, the winner of the election is from $B(\hat{o}, 3r)$.*

Proof. Let $Z := \lceil z \cdot (n - 1) \rceil$. Assume that exactly $s \geq Z$ out of the n candidates are drawn from S . Consider a candidate $i \notin T$. We will compare the average number of points of candidates in S with the maximum possible number of points of candidate i , and show that the former exceeds the latter.

- Each voter $j \notin S$ gives at most one point to i . On the other hand, even if j ranks all of S in the last s positions, the total number of points assigned by j to S is at least $\sum_{k=n-s}^{n-1} g_n(k)$. The difference between the number of votes to i and the average number of votes to candidates in S is thus at most $\frac{1}{s} \cdot \sum_{k=n-s}^{n-1} (1 - g_n(k))$.

Because no more than a $1 - \mu$ fraction of voters are strictly outside S , the total advantage of i over an average candidate in S resulting from such voters is at most $\Delta_i := \frac{1}{s} \cdot (1 - \mu) \cdot \sum_{k=n-s}^{n-1} (1 - g_n(k))$.

- Each voter $j \in S$ will rank all candidates in S (who are at distance at most $2r$ from her) ahead of all candidates outside T (who are at distance strictly more than $3r - r = 2r$ from her).

Let $m \geq 0$ be such that j ranks i in position $s + m$. Then, i gets $g_n(s + m)$ points from j . Because j ranks all of S ahead of i , she gives at least $\sum_{k=0}^{s-1} g_n(k + m)$ points in total to S . Hence, the difference in the number of points that j gives to an average candidate in S and the number of votes that j gives to i is at least $\frac{1}{s} \cdot \sum_{k=0}^{s-1} (g_n(k + m) - g_n(s + m))$.

Because at least a μ fraction of voters are in S , the total advantage of an average candidate in S resulting from voters in B is at least $\Delta_S := \frac{1}{s} \cdot \mu \cdot \sum_{k=0}^{s-1} (g_n(k + m) - g_n(s + m))$.

We show that $\Delta_S > \Delta_i$, using condition (5). Because g_n is monotone non-increasing, and because $s \geq Z$, we get that

$$\begin{aligned} & \sum_{k=0}^{s-1} (g_n(k + m) - g_n(s + m)) \\ & \geq \sum_{k=0}^{Z-1} (g_n(k + m) - g_n(Z + m)) \\ & \stackrel{(5)}{>} \frac{2(1-z)}{z} \cdot \sum_{k=n-Z}^{n-1} (1 - g_n(k)). \end{aligned}$$

Because $z > \frac{1}{2}$ and g_n is monotone, we get that $\sum_{k=n-Z}^{n-1} (1 - g_n(k)) \geq \frac{1}{2} \sum_{k=n-s}^{n-1} (1 - g_n(k))$. Hence,

$$\begin{aligned} \Delta_S & > \frac{1}{s} \cdot \mu \cdot \frac{1-z}{z} \cdot \sum_{k=n-s}^{n-1} (1 - g_n(k)) \\ & \stackrel{\mu \geq z}{\geq} \frac{1}{s} \cdot (1 - \mu) \cdot \sum_{k=n-s}^{n-1} (1 - g_n(k)) = \Delta_i. \quad \blacksquare \end{aligned}$$

We now wrap up the sufficiency portion of the proof of Theorem 3.1. We distinguish two cases, based on the number of candidates n . If $n < \hat{n}$, then Lemma 3.3 implies an upper bound of $n + 1 \leq \hat{n} \leq \max(n_0, \frac{1}{1-z}) = O(1)$ on the expected distortion. Now assume that $n \geq \hat{n}$. Recall that $\hat{r} := \inf\{r \mid q_{B(\hat{o}, r)} \geq \mu\}$. By Lemmas 3.5 and 3.6, for any $r \geq \hat{r}$, the probability that the election's winner is outside $B(\hat{o}, 3r)$ is at most $\frac{e}{1-z} \cdot H(r)$. The rest of the proof is similar to that of Lemma 3.3. We again use that $\mathbb{E}_C [c_{w(C)}] \leq c_{\hat{o}} + \int_0^\infty \Pr_C [d_{\hat{o}, w(C)} \geq r] dr$, and bound

$$\begin{aligned} & \int_0^\infty \Pr_C [d_{\hat{o}, w(C)} \geq r] dr \\ & = \int_0^{3\hat{r}} \Pr_C [d_{\hat{o}, w(C)} \geq r] dr + \int_{3\hat{r}}^\infty \Pr_C [d_{\hat{o}, w(C)} \geq r] dr \\ & \stackrel{\text{Lemmas 3.5, 3.6}}{\leq} \int_0^{3\hat{r}} 1 dr + \int_{3\hat{r}}^\infty \frac{e}{1-z} \cdot H(r) dr \\ & \leq 3\hat{r} + \frac{e}{1-z} \cdot \int_0^\infty H(r) dr \stackrel{(2)}{=} 3\hat{r} + \frac{e}{1-z} \cdot c_{\hat{o}}. \end{aligned}$$

To upper-bound \hat{r} , recall that at least a $1 - \mu$ fraction of voters are outside of $B(\hat{o}, \hat{r})$ or on the boundary. Therefore, by Inequality (3), $c_{\hat{o}} \geq \hat{r} \cdot (1 - \mu)$. Substituting this bound, the expected cost of the winning candidate is at most

$$\left(1 + \frac{3}{1-\mu} + \frac{e}{1-z}\right) \cdot c_{\hat{o}} = \left(1 + \frac{4}{1-\mu}\right) \cdot c_{\hat{o}} = O(c_{\hat{o}}),$$

as y depends only on the voting system \mathcal{V} , but not on the metric space or the number of candidates. This completes the proof of sufficiency.

3.2 Necessity

Next, we prove that the condition in Theorem 3.1 is also necessary for constant distortion. We assume that the condition (4) does not hold, i.e., for every $y \in (0, 1)$ and n_0 , there exists an $n \geq n_0$ such that

$$\begin{aligned} & y \cdot \sum_{k=0}^{\lceil y \cdot (n-1) \rceil - 1} (g_n(k) - g_n(\lceil y \cdot (n-1) \rceil)) \\ & \leq (1 - y) \cdot \sum_{k=n - \lceil y \cdot (n-1) \rceil}^{n-1} (1 - g_n(k)). \end{aligned} \quad (6)$$

We will show that the distortion of \mathcal{V} is not bounded by any constant.

The high-level idea of the construction is as follows: we define two tightly knit clusters A and B that are far away from each other. A contains a large α fraction of the population, and thus should in an optimal solution be the one that the winner is chosen from. We will ensure that with probability at least $\frac{1}{2}$, the winner instead comes from B . Because

B is far from A , most of the population then is far from the chosen candidate, giving much worse cost than optimal.

The metrics underlying A and B are as follows: B will essentially provide an “ordering,” meaning that whichever set of candidates is drawn from B , all voters in B (and essentially all in A) agree on their ordering of the candidates. This will ensure that one candidate from B will get a sufficiently large fraction of first-place votes, and will be ranked highly enough by voters from A , too. A will be based on a large number M of discrete locations ω . Their pairwise distances are chosen i.i.d.: as a result, the rankings of voters are uniformly random, and there is no consensus among voters in A on which of their candidates they prefer. Because the vote is thus split, the best candidate from B will win instead.

The following parameters (whose values are chosen with foresight) will be used to define the metric space.

- Let $c > 1$ be any constant; we will construct a metric space and number of candidates for which the distortion is at least c .
- Let $\beta \in (0, \frac{1}{2})$ solve the quadratic equation $\frac{2\beta+1}{3\beta} \cdot (1 - \beta) = 2c - 1$. A solution exists because at $\beta = \frac{1}{2}$, the left-hand side is $\frac{2}{3} < 2c - 1$; it goes to infinity as $\beta \rightarrow 0$, while the right-hand side is a positive constant. β is the fraction of voters in the small cluster B .
- Let $\alpha = 1 - \beta$ denote the fraction of voters in the large cluster A .
- Let $s = \frac{1+\beta}{\beta}$ be the distance between the clusters B and A . (Each cluster will have diameter at most 2.)
- Let $\hat{\alpha} \geq \frac{1}{2} + \frac{\alpha}{2} > \alpha$ satisfy $4\hat{\alpha} \cdot (1 - \hat{\alpha}) < \alpha \cdot (1 - \alpha)$; such an $\hat{\alpha}$ exists because the left-hand side goes to 0 as $\hat{\alpha} \rightarrow 1$. $\hat{\alpha} < 1$ is a high-probability upper bound on the fraction of candidates that will be drawn from A .
- Let $n_0 = \frac{4}{\beta^2} > 16$; this is a lower bound on the number of candidates that ensures that the actual fraction of candidates drawn from A is at most $\hat{\alpha}$ with sufficiently high probability.
- Let $n \geq n_0$ be the n whose existence is guaranteed by the assumption (6) (for $y = \hat{\alpha}$ and n_0).
- Let $M = n^3$; this is the number of discrete locations ω we construct within the larger cluster A .

We now formally define the metric space consisting of two clusters:

Definition 3.1. *The metric space consists of two clusters A and B . A has M discrete locations, and \mathbf{q} has a point mass of $\frac{\alpha}{M}$ on each such location. The total probability mass on B is $q_B = 1 - \alpha$, distributed uniformly over the interval $[1, 2]$. Locations in B are identified by $x \in [1, 2]$. The distances are defined as follows:*

1. For each distinct pair $\omega, \omega' \in A$, the distance $d_{\omega, \omega'}$ is drawn independently uniformly at random from $[1, 2]$.
2. For each distinct pair $x, x' \in B$ of locations, the distance is defined to be $d_{x, x'} := \min(x, x')$.
3. Partition $B = [1, 2]$ into $M!$ disjoint intervals I_π of length $1/M!$ each, one for each permutation of the M locations

in A . For $\omega \in A$ and $x \in I_\pi$, let $\pi^{-1}(\omega)$ be the position of ω in π , and define the distance between ω and x to be $d_{\omega, x} = s + \frac{x}{4} + \frac{\pi^{-1}(\omega)}{M!}$.

Proposition 3.7. *Definition 3.1 defines a metric.*

Proof. Non-negativity, symmetry, and indiscernibles hold by definition. Because all distances within clusters are in $[1, 2]$, and distances across clusters are more than 2, the triangle inequality holds for all pairs $\omega, \omega' \in A$ and all pairs $x, x' \in B$.

Because $d_{\omega, x} \in [s, s + 1]$ for all $\omega \in A$ and $x \in B$, and distances within A or B are at least 1, there can be no shorter path than the direct one between any $\omega \in A$ and $x \in B$. Therefore, the triangle inequality is satisfied. \square

Now consider a (random) set C of n candidates, drawn i.i.d. from \mathbf{q} . We are interested in the event that the resulting slate of candidates is highly representative of the voters, in the following sense.

Definition 3.2. *Let C be the (random) set of n candidates drawn from \mathbf{q} . Let \mathcal{E} be defined as the conjunction of the following:*

1. For each location $\omega \in A$, the set C contains at most one candidate from ω .
2. At least a $\frac{\beta}{2}$ fraction of candidates in C is drawn from B (and thus at most an $\hat{\alpha}$ fraction of candidates are from A).
3. At least an $\frac{\alpha}{2}$ fraction of candidates in C is drawn from A .
4. No pair $x, x' \in B \cap C$ has $|x - x'| < \frac{1}{(M-1)!}$.

Lemma 3.8 uses standard tail bounds to show that \mathcal{E} happens with probability at least $\frac{1}{2}$; then, Lemma 3.9 shows that whenever \mathcal{E} happens, the winner is from B .

Lemma 3.8. *\mathcal{E} happens with probability at least $\frac{1}{2}$.*

Lemma 3.9. *Whenever \mathcal{E} happens, the winning candidate is from B .*

Proof. Let b be the actual number of candidates drawn from B , and $a = n - b$ the number of candidates drawn from A . Because we assumed that \mathcal{E} happened, $b \geq \frac{\beta}{2} \cdot n$ and $a \leq \hat{\alpha} \cdot n$. Let C_A be the set of candidates drawn from A . Under \mathcal{E} , C_A contains at most one candidate from each location $\omega \in A$. As a result, because the random distances within A are distinct with probability 1, there will be no ties in the rankings of any voters.

Let \hat{v} be the candidate from B with smallest value \hat{x} . With probability 1, the x value of \hat{v} is unique. Consider some arbitrary candidate $i \in C_A$ from location ω' . We calculate the contributions to \hat{v} and i from voters in B and in A separately, and show that \hat{v} beats i . Because this holds for arbitrary i , the candidate \hat{v} or another candidate from B wins.

1. We begin with points given out by voters in B . By definition of the distances within B , \hat{v} is ranked first by all voters in B .

Voters in I_π rank the candidates from A according to their order in π . For each ordering of C_A , exactly a $\frac{1}{a!}$

fraction of permutations induces that ordering. In particular, for each $k \in 1, \dots, a$, exactly a $1/a$ fraction of voters places i in position $k + b$. Thus, i obtains a total of $(1 - \alpha) \cdot \sum_{k=n-a}^{n-1} \frac{1}{a} \cdot g_n(k)$ points from voters in B . Overall, \hat{i} obtains an advantage of at least $\Delta_B = (1 - \alpha) \cdot \frac{1}{a} \cdot \sum_{k=n-a}^{n-1} (1 - g_n(k))$.

2. Next, we analyze the number of points given out by voters in A . The distance from any voter location $\omega \in A$ to \hat{i} is at most $s + \frac{\hat{x}}{4} + \frac{n}{M!}$. Under \mathcal{E} , no other candidate from B can be at a location $x \leq \hat{x} + \frac{1}{(M-1)!}$; therefore, the distance from any voter location $\omega \in A$ to any other candidate $x \in B$ is at least

$$d_{\omega, x} \geq s + \frac{x}{4} \geq s + \frac{\hat{x}}{4} + \frac{1}{4(M-1)!} > s + \frac{\hat{x}}{4} + \frac{n}{M!} \geq d_{\omega, \hat{i}},$$

so all voters in A prefer \hat{i} over any other candidate from B . Hence, \hat{i} obtains at least $\alpha \cdot g_n(a)$ points combined from voters in A .

To analyze the votes from voters in A for candidates from A , we first notice that \mathcal{E} and the draw of candidates are independent of the distances within A . Hence, even conditioned on \mathcal{E} , the distances $d_{\omega, \omega'}$ between locations in A are i.i.d. uniform from $[1, 2]$. In particular, each location $\omega \in A$ ranks the candidates in C_A in uniformly random order. Furthermore, for two locations $\omega \neq \omega'$, the rankings of C_A are independent; the reason is that they are based on disjoint vectors of distances $(d_{\omega, i})_{i \in C_A}, (d_{\omega', i})_{i \in C_A}$. We use this independence to apply tail bounds. Let ω' be the location of i . Voters rank i as follows:

- Among locations ω without a candidate of their own, in expectation, a $1/a$ fraction of voters will rank i in position k , for each $k = 0, \dots, a - 1$.
- Among the $a - 1$ locations $\omega \neq \omega'$ with a candidate of their own, in expectation, a $1/(a - 1)$ fraction of voters will rank i in position k , for each $k = 1, \dots, a - 1$.
- Voters at ω' will rank i in position 0.

For each k , let the random variable X_k be the number of locations that rank i in position k . By the preceding arguments, $\mathbb{E}[X_k] = \frac{M}{a}$, and X_k is a sum of M independent (not i.i.d.) Bernoulli random variables. Hence, by the Hoeffding bound, the probability that more than a $\frac{2}{a}$ fraction of voters rank i in position k is at most $2 \exp(-2 \cdot \frac{1}{a^2} \cdot M) \leq 2 \exp(-n)$. By a union bound over all candidates $i \in C_A$ and all values $k = 0, \dots, a - 1$, with high probability, for all i and k , the fraction of voters (in A) ranking i in position k is at most $\alpha \cdot \frac{2}{a}$. Because the total fraction of voters in A is α , any excess votes for some (early) positions k must be compensated by fewer votes for other (late) positions k' . Relaxing the constraint that the number of votes for each position k must be non-negative, we can upper-bound the total points for i by assuming that each of the positions $k = 0, \dots, a - 2$ receives twice the expected number of votes, while position $k = a - 1$ receives a negative number of votes that compensates for the excess votes. Then, the advantage for i

over \hat{i} from votes from A is at most

$$\begin{aligned} \Delta_A &:= \alpha \left(\sum_{k=0}^{a-2} \frac{2}{a} g_n(k) + \frac{2-a}{a} g_n(a-1) - g_n(a) \right) \\ &= \frac{\alpha}{a} \left(\sum_{k=0}^{a-2} (2g_n(k) - g_n(a-1) - g_n(a)) \right. \\ &\quad \left. + g_n(a-1) - g_n(a) \right) \\ &\stackrel{g_n \text{ monotone}}{\leq} \frac{2\alpha}{a} \cdot \sum_{k=0}^{a-1} (g_n(k) - g_n(a)). \end{aligned}$$

Finally, we can bound

$$\begin{aligned} \Delta_A \cdot a &\leq 2\alpha \cdot \sum_{k=0}^{a-1} (g_n(k) - g_n(a)) \\ &\stackrel{\alpha \leq \hat{\alpha}, g_n \text{ mon.}}{\leq} 2\hat{\alpha} \sum_{k=0}^{\lceil \hat{\alpha}(n-1) \rceil - 1} (g_n(k) - g_n(\lceil \hat{\alpha}(n-1) \rceil)) \\ &\stackrel{(6), \text{ Def. of } n}{\leq} 2(1 - \hat{\alpha}) \cdot \sum_{k=n - \lceil \hat{\alpha}(n-1) \rceil}^{n-1} (1 - g_n(k)) \\ &\stackrel{g_n \text{ mon.}}{\leq} 2(1 - \hat{\alpha}) \cdot \frac{\hat{\alpha} \cdot (n-1)}{a} \sum_{k=n-a}^{n-1} (1 - g_n(k)) \\ &\stackrel{a \geq \alpha \cdot n/2}{\leq} 2(1 - \hat{\alpha}) \cdot \frac{2\hat{\alpha}}{\alpha} \sum_{k=n-a}^{n-1} (1 - g_n(k)) \\ &\stackrel{\text{Def. of } \hat{\alpha}}{<} (1 - \alpha) \cdot \sum_{k=n-a}^{n-1} (1 - g_n(k)) = \Delta_B \cdot a. \end{aligned}$$

Thus, \hat{i} beats all candidates drawn from A , and the winner will be from B . \square

Using the preceding lemmas, the proof of necessity is almost complete. Consider the metric space with all the parameters as defined above. By Lemmas 3.8 and 3.9, with probability at least $\frac{1}{2}$, the winner is from B . The social cost of any candidate from B is at least $\beta \cdot 0 + (1 - \beta) \cdot (s + 1)$. On the other hand, the social cost of any candidate from A is at most $(1 - \beta) \cdot 2 + \beta \cdot (s + 1) = 3$. The distortion in this case is thus at least $\frac{(1-\beta) \cdot (s+1)}{3} = \frac{(2\beta+1) \cdot (1-\beta)}{3\beta} = 2c - 1$.

In the other case (when \mathcal{E} does not occur — this happens with probability at most $\frac{1}{2}$), the distortion is at least 1, so that the expected distortion is at least $\frac{1}{2}(2c - 1) + \frac{1}{2} \cdot 1 = c$.

4 Conclusions

When candidates are drawn i.i.d. from the voter distribution, we showed that whether a positional voting system \mathcal{V} has expected constant distortion can be almost fully characterized by its limiting behavior. In particular, if the limiting scoring rule is not constant on $(0, 1)$, then \mathcal{V} has constant expected distortion; if the limiting scoring rule is a constant other than 1 on $(0, 1)$, then \mathcal{V} has super-constant expected distortion. A more subtle condition depending on the “rate of convergence” to the limit rule completes the characterization.

Our Theorem 3.1 currently does not characterize the order of growth of the distortion. With some effort, the proof could likely be adapted to the case where the y in the theorem is a function $y(n)$, which would allow us to characterize the rate at which the distortion grows with n .

For specific voting systems, the proof of Theorem 3.1 can often be adapted to give tighter bounds. For example, straightforward modifications of the proof can be used to show that the distortion of k -approval or k -veto (where each voter can veto k candidates) for constant k grow as $\Omega(n)$. This matches the $O(n)$ upper bound from Lemma 3.3, giving a tight analysis of the distortion of these voting systems.

Similarly, the sufficiency proof can be adapted to show that the distortion of Borda Count is at most 16, for all metric spaces and all n . When the number of candidates grows large enough, the expected distortion is in fact bounded by 10.

Our results indicate that if one is concerned about systematic, and possibly adversarial, bias in which candidates run for office, randomizing the slate of candidates may be part of a solution approach. Such an approach can be considered as a step in the direction of *lottocracy* and *sortition* (Dowlen 2008; Guerrero 2014; Landmore 2013), in which office holders are directly chosen at random from the population. Pure lottocracy does well in terms of *representativeness* of office holders, but one of its main drawbacks is the potential lack of competency. As a broader direction for future research, our work here suggests devising models that capture the tension between these two objectives, and would allow for the design of hybrid mechanisms that navigate the tradeoff successfully.

Acknowledgments Part of this work was done while Yu Cheng was a student at University of Southern California. Yu Cheng was supported in part by Shang-Hua Teng’s Simons Investigator Award. Shaddin Dughmi was supported in part by NSF CAREER Award CCF-1350900 and NSF grant CCF-1423618. David Kempe was supported in part by NSF grants CCF-1423618 and IIS-1619458. We would like to thank anonymous reviewers for useful feedback.

References

- Anshelevich, E., and Postl, J. 2016. Randomized social choice functions under metric preferences. In *Proc. 25th Intl. Joint Conf. on Artificial Intelligence*, 46–59.
- Anshelevich, E.; Bhardwaj, O.; and Postl, J. 2015. Approximating optimal social choice under metric preferences. In *Proc. 29th AAAI Conf. on Artificial Intelligence*, 777–783.
- Anshelevich, E. 2016. Ordinal approximation in matching and social choice. *ACM SIGecom Exchanges* 15(1):60–64.
- Arrow, K. 1951. *Social Choice and Individual Values*. Wiley.
- Barberà, S.; Gul, F.; and Stacchetti, E. 1993. Generalized median voter schemes and committees. *Journal of Economic Theory* 61:262–289.
- Barberà, S. 2001. An introduction to strategy-proof social choice functions. *Social Choice and Welfare* 18:619–653.
- Black, D. 1948. On the rationale of group decision making. *J. Political Economy* 56:23–34.
- Black, D. 1958. *The Theory of Committees and Elections*. Cambridge University Press.
- Boutilier, C., and Rosenschein, J. S. 2016. Incomplete information and communication in voting. In Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D., eds., *Handbook of Computational Social Choice*. Cambridge University Press. chapter 10, 223–257.
- 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., and Procaccia, A. D. 2011. Voting almost maximizes social welfare despite limited communication. *Artificial Intelligence* 175(9):1655–1671.
- Cheng, Y.; Dughmi, S.; and Kempe, D. 2017. Of the people: Voting is more effective with representative candidates. In *Proc. 18th ACM Conf. on Economics and Computation*, 305–322.
- de Borda, J.-C. 1784. Mémoire sur les élections au scrutin. *Histoire de l’Académie Royale des Sciences, Paris* 657–665.
- Dowlen, O. 2008. *The political potential of sortition: A study of the random selection of citizens for public office*. Imprint Academic.
- Downs, A. 1957. An economic theory of political action in a democracy. *The Journal of Political Economy* 65(2):135–150.
- Feldman, M.; Fiat, A.; and Golomb, I. 2015. On voting and facility location. arXiv:1512.05868.
- Fraenkel, J., and Grofman, B. 2014. The Borda count and its real-world alternatives: Comparing scoring rules in Nauru and Slovenia. *Australian Journal of Political Science* 49(2):186–205.
- Gibbard, A. F. 1973. Manipulation of voting schemes: a general result. *Econometrica* 41(4):587–601.
- Goel, A.; Krishnaswamy, A. K.; and Munagala, K. 2017. Metric distortion of social choice rules: Lower bounds and fairness properties. In *Proc. 18th ACM Conf. on Economics and Computation*, 287–304.
- Guerrero, A. A. 2014. Against elections: The lottocratic alternative. *Philosophy & Public Affairs* 42(2):135–178.
- Landmore, H. 2013. Deliberation, cognitive diversity, and democratic inclusiveness: An epistemic argument for the random selection of representatives. *Synthese* 190(7):1209–1231.
- Merrill, S., and Grofman, B. 1999. *A unified theory of voting: Directional and proximity spatial models*. Cambridge University Press.
- Moulin, H. 1980. On strategy-proofness and single peakedness. *Public Choice* 35:437–455.
- Procaccia, A. D., and Rosenschein, J. S. 2006. The distortion of cardinal preferences in voting. In *Proc. 10th Intl. Workshop on Cooperative Inform. Agents X*, 317–331.
- Procaccia, A. D. 2010. Can approximation circumvent Gibbard-Satterthwaite? In *Proc. 24th AAAI Conf. on Artificial Intelligence*, 836–841.
- Reilly, B. 2002. Social choice in the South Seas: Electoral innovation and the Borda count in the Pacific Island countries. *International Political Science Review* 23(4):355–372.
- Richards, D.; Richards, W. A.; and McKay, B. 1998. Collective choice and mutual knowledge structures. *Advances in Complex Systems* 1:221–236.
- Satterthwaite, M. A. 1975. Strategy-proofness and Arrows conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory* 10:187–217.