A Game Theory Toolkit for Voting Rules

MOSES CHARIKAR
Stanford University
and
PRASANNA RAMAKRISHNAN
Stanford University
and
KANGNING WANG
Rutgers University

Classical voting theory is often viewed as a field of impossibilities, where even the most basic criteria can be unsatisfiable. A growing line of work seeks to sidestep these hurdles by considering voting rules that output randomized distributions over candidates. In this survey, we explore a particular family of randomized voting rules derived from equilibria of simple two-player games, most notably maximal lotteries and stable lotteries. We survey a flurry of recent work using these ideas to prove several positive results in voting theory. These include improved bounds for randomized voting rules in the *distortion* problem, and progress on *committee selection* questions that appear deterministic in nature, such as Condorcet winning sets, approximately stable committees, and approximately dominating sets.

Categories and Subject Descriptors: [**Theory of Computation**]: Algorithmic Game Theory and Mechanism Design

General Terms: Algorithms, Economics, Theory

Additional Key Words and Phrases: Computational Social Choice, Game Theory, Distortion, Committee Selection

1. INTRODUCTION

Voting is a fundamental tool for aggregating individual preferences into a collective decision. While elections are most associated with representative democracy, the same abstraction is at the heart of many other settings: an organization making new hires, judges choosing prize winners, academic conferences choosing papers, a group of friends choosing a restaurant, or human labelers fine-tuning a language model towards better responses. In each case, voters express preferences over a set of candidates, and the goal is to select an outcome that is in some sense fair, representative, or socially efficient.

The persistent challenge in voting theory is that even very modest desiderata are not always satisfiable. This issue is prominently exemplified by the impossibility theorems of Arrow and Gibbard–Satterthwaite, but the earliest example is Condorcet's paradox from 1785, which says that a candidate that beats each other candidate in a head-to-head majority vote (called a *Condorcet winner*) does not always exist. That is, there are elections where no matter which candidate wins, a majority of voters will prefer an alternative.

Authors' addresses: `moses@cs.stanford.edu`, `pras1712@stanford.edu`, `kn.w@rutgers.edu`.

One approach to circumventing these negative results is to use *randomized* voting rules that choose a distribution over candidates rather than a single winning candidate. For example, we can use game theory to prove the existence of a randomized analogue of a Condorcet winner called a *maximal lottery*, which beats each candidate for half the voters in expectation. Although the richer design space of randomized voting rules opens up new possibilities, the idea of deciding elections by chance can still be instinctively unsettling, especially given the close association between voting and high-stakes, politically charged decisions. But randomization can feel much more natural in lower-stakes decisions among friends, in processes that are already inherently noisy, or even in settings involving language models, which are stochastic by design. Even in political contexts, randomized or fractional solutions are not without precedent. For example, ancient Athens used sortition to fill many public offices by lot, and the Roman Republic divided executive authority between two annually elected consuls.

In this survey, we will explore a particular family of randomized voting rules that can be derived from the mixed-strategy Nash equilibria of simple two-player games, like maximal lotteries. We will see how these voting rules have played pivotal roles in new results from seemingly different parts of computational social choice: *committee selection* and *distortion*. A recurring theme is that these randomized solutions are useful not only as an end in themselves, but also as a tool to show the existence of *deterministic* structures with desirable properties.

2. MAXIMAL LOTTERIES AND STABLE LOTTERIES

Imagine that an election is taking place, and two gamblers decide to play the following game. Each gambler independently chooses a candidate to bet on, and then they ask a random voter to compare their chosen candidates. Whoever chose the candidate that the voter prefers is the winner of the game. (If the gamblers happen to choose the same candidate, they flip a coin to decide the winner.)

Intuitively, the gamblers want to choose the "best" candidate — betting on an unpopular candidate is clearly a losing proposition. Conversely, it might be reasonable to think that effective strategies for choosing candidates in the game (given knowledge of how voters' preferences are distributed) correspond to voting rules that choose more popular candidates. Arguably, the optimal strategy would be to *randomize* over the candidates according to the *mixed-strategy Nash equilibrium* of the game. The equilibrium distribution is identical for each gambler since the game is symmetric, and playing according to the equilibrium guarantees a win with probability at least $\frac{1}{2}$ against any fixed choice of the opponent.

This guarantee underpins the concept of a *maximal lottery*. To give a more formal definition, we first introduce some basic notation. An *election* (or *preference profile*) $(V, C, \succ_V)$ consists of a set V of n voters, a set C of m candidates, and a strict linear order $\succ_v$ over the candidates C for each voter $v \in V$. We use $\frac{1}{n}|a \succ b|$ to denote the fraction of voters who prefer candidate a over candidate b (i.e., $\frac{1}{n}|a \succ b| = \Pr_{v\sim V}[a \succ_v b]$). For simplicity, we adopt the convention that $\frac{1}{n}|a \succ a| = 0$, but note that it can also be natural to set $\frac{1}{n}|a \succ a| = \frac{1}{2}$ (akin to the coin-flip tiebreaker between the gamblers). Results will be stated so that they are true with either convention (and are often slightly stronger with the latter).

THEOREM 2.1 ([KREWERAS 1965; FISHBURN 1984]). *In any election, there exists a distribution D_{ML} over candidates such that for all candidates $a \in C$,*

$$\mathop{\mathbb{E}}_{b \sim D_{\mathrm{ML}}}[\tfrac{1}{n}|a \succ b|] \leq \frac{1}{2}.$$

A distribution satisfying this condition is called a maximal lottery.

The proof is a fairly straightforward application of von Neumann's minimax theorem applied to our game between the gamblers. We will defer a proof until Theorem 2.2, which is a more general result.

For some illustrative examples, consider elections with three candidates a, b, c, and voters of three types A, B, C with preferences $a \succ b \succ c$, $b \succ c \succ a$, and $c \succ a \succ b$ respectively. With one voter of each type, this is the classic example of Condorcet's paradox, and the unique maximal lottery is the uniform distribution over the candidates. If we have two voters of types A and B each and one voter of type C, then the maximal lottery would choose (a, b, c) with probabilities $(\frac{3}{5}, \frac{1}{5}, \frac{1}{5})$. If we have three voters of type A, two voters of type B, and one voter of type C, any distribution which chooses (a, b, c) with probabilities $(p, 0, 1-p)$ for $p \in [\frac{2}{3}, 1]$ is a maximal lottery. In elections that have a Condorcet winner (a candidate a such that $\frac{1}{n}|a \succ b| > \frac{1}{2}$ for all candidates $b \neq a$), the maximal lottery deterministically chooses that candidate. If one takes the view that Condorcet's paradox is unlikely[1], then a maximal lottery can be thought of as like any other deterministic Condorcet-consistent voting rule, with a natural randomized tie-breaking scheme if needed.

Maximal lotteries were originally introduced by Kreweras [1965] in a short paper primarily about Condorcet's paradox and Arrow's theorem. In discussing ways of escaping these challenges at the very end, he mentions that using game theory, one can show that the set of convex combinations of candidates no longer suffers from the curse of intransitivity. Fishburn [1984] gave a more in-depth treatment of maximal lotteries, formally proving their existence, analyzing their properties, and also defining the broader class of "maximal lottery methods" which can be derived from variants of the earlier zero-sum game over candidates. For example, if the gamblers decide the winner by using a majority vote instead of a random voter's preference, the resulting equilibrium is a *C1-maximal lottery* (named much later by Brandl et al. [2018]). In surveying randomized voting rules, Brandt [2017] vividly recounts that maximal lotteries have been rediscovered numerous times by researchers across academic disciplines.

Recently, Cheng et al. [2020] generalized maximal lotteries to *stable lotteries*, which are distributions over *sets* of candidates (called *committees*). If we modify our game between the gamblers so that one is allowed to pick k candidates and wins if a random voter prefers *any* of their candidates over the opponent's candidate, then their equilibrium distribution is a stable lottery.

To give a formal description, we introduce some notation. We say that voter v prefers candidate a over committee S, or $a \succ_v S$, if $a \succ_v b$ for each $b \in S$. We define $\frac{1}{n}|a \succ S|$ to be the fraction of voters that prefer candidate a over S (i.e.,

[1] For example, with single-peaked preferences, or in an impartial culture (see Gehrlein [2002] for more).

$\frac{1}{n}|a \succ S| = \Pr_{v \sim V}[a \succ_v S]$). In Section 4.1, we briefly use $|a \succ S|$ to denote the *number* of voters that prefer a over S.

If $a \in S$, we use the convention $\frac{1}{n}|a \succ S| = 0$, but there is also a natural alternate convention[2] which breaks ties between duplicate candidates uniformly,

analogous to setting $\frac{1}{n}|a \succ a| = \frac{1}{2}$. Like before, results will be stated to be true with either convention.

THEOREM 2.2 (STABLE LOTTERY [CHENG ET AL. 2020]). *In any election, there exists a distribution $D_{k\text{-SL}}$ over committees of size k such that for all candidates $a \in C$,*

$$\mathop{\mathbb{E}}_{S \sim D_{k\text{-SL}}}[\tfrac{1}{n}|a \succ S|] \leq \frac{1}{k+1}.$$

A distribution satisfying this condition is called a stable k-lottery*.*

PROOF. Consider the following zero-sum game. The *attacker* chooses a distribution D_a over candidates, the *defender* chooses a distribution D over committees of k candidates, then we sample $a \sim D_a, S \sim D$, and the defender pays the attacker $\frac{1}{n}|a \succ S|$ (one can imagine, a penny for each voter that prefers the attacker's candidate over the defender's committee). Since the players can randomize, the minimax theorem [von Neumann 1928] shows that under optimal play, the order of the players does not change the expected value of the game. That is,

$$\max_{D_a} \min_{D} \mathop{\mathbb{E}}_{\substack{a \sim D_a \\ S \sim D}}[\tfrac{1}{n}|a \succ S|] = \min_{D} \max_{D_a} \mathop{\mathbb{E}}_{\substack{a \sim D_a \\ S \sim D}}[\tfrac{1}{n}|a \succ S|].$$

If the defender plays second, then they can simply sample from the attacker's distribution k times (setting $D = D_a^k$). The defender's samples are just as likely to be a voter's favorite as the attacker's sample, so the defender pays at most $\frac{1}{k+1}$ in expectation.

□

Before diving into the various ways of applying maximal lotteries and stable lotteries, we briefly mention some useful structural facts.

The first concerns efficient computation of these distributions. The equilibria of zero-sum games can be computed using linear programming in time polynomial in the size of the payoff matrix. Thus, maximal lotteries can be computed in poly(m) time, and more generally, stable lotteries can be computed in poly(m^k) time given the payoff matrix as input (which can be computed in time linear in n given the preferences). Cheng et al. [2020] show that the exponential dependence on k can be avoided at the cost of approximation; using tools from no-regret learning, they show that ε-approximate stable k-lotteries can be computed in poly($m, \frac{1}{\varepsilon}$) time.

Second, while stable lotteries are defined as distributions over committees, one can show that there always exists a stable lottery that is a *product distribution*. That is, in any election there is a distribution D over candidates such that D^k is a stable k-lottery.

[2] This convention also makes it easier to generalize S to be a *multiset*, which is cleaner when S is sampled from a distribution over candidates with replacement, as in the proof of Theorem 2.2.

This fact is proven implicitly by Charikar et al. [2025], and explicitly by Charikar et al. [2025]. It turns out that if we take $k \to \infty$, then D tends to the distribution which chooses a uniformly random voter's favorite candidate (called a *random dictatorship*).

With this in mind, one can alternatively view stable k-lotteries as interpolating between a maximal lottery ($k = 1$) and a random dictatorship ($k \to \infty$). In some settings where a random dictatorship is useful (like in the metric distortion setting considered by Charikar et al. [2025] and Cai et al. [2026]), we can use a stable k-lottery instead and achieve comparable quantitative guarantees while satisfying additional normative properties.

Finally, it is worth mentioning that some papers do not define stable lotteries in exactly the same way we do here. Cheng et al. [2020] originally used a weaker definition with $< \frac{1}{k}$ in place of $\leq \frac{1}{k+1}$, to align with the existing notion of a "stable committee" when the distribution is deterministic.[3] Cai et al. [2026] instead define a stable lottery directly as the Nash equilibrium of the game described in the proof of Theorem 2.2, following Charikar et al. [2024], which does the same with maximal lotteries.[4] Ultimately, all these choices are reasonable, and their differences have essentially no impact on results.

3. APPLICATION 1: COMMITTEE SELECTION

The committee selection problem addresses social choice settings where the goal is to choose a *committee* of k winners rather than a single winner. For two familiar motivating examples, consider an academic department choosing candidates to hire (or interview), or a conference choosing papers to accept.[5] One commonly studied goal is to choose committees such that no large group of voters would have preferred a candidate that is excluded from the committee. What it means for a group to be "large" and to "prefer an excluded candidate" is open to interpretation, and different choices lead to a variety of definitions which are more or less attainable.

In this section, we will explore how maximal lotteries, stable lotteries, and related randomized voting rules have been used to establish positive results on the existence of three closely related solution concepts in committee selection: Condorcet winning sets, approximately stable committees, and approximately dominating sets.

The fact that these solution concepts are all fundamentally deterministic makes it notable that probabilistic tools have been critical in understanding them. The techniques are applied in similar ways across the three problems, so we will showcase them in one context (Condorcet winning sets), and briefly discuss the others.

Lastly, we mention that committee selection is a rich field of study, and the setting of ranked preferences and the desiderata that we focus on here are only a thin slice. For more in-depth surveys on committee selection, we refer the reader to Faliszewski et al. [2017] and Lackner and Skowron [2023].

[3] See Section 3.2 for more details.

[4] Though, a maximal lottery as defined in Theorem 2.1 is always an equilibrium of the game, but the same is not necessarily true of a stable lottery as defined in Theorem 2.2.

[5] Of course, it would often be unreasonable to elicit full ranked preferences in these settings, but it is still a helpful abstraction for understanding what objectives are attainable with enough effort.

3.1 Condorcet winning sets

Condorcet's paradox points to an unsettling possibility — no matter which candidate wins, a majority of voters may be upset. For example, if a department can only hire one candidate, most of the faculty may think another candidate would have been better. Even if the department can make *multiple* hires, there could be a risk of an even more embarrassing situation: most faculty think some other candidate would have been better than *all* the hires. Of course, the department could avoid this kind of disaster by hiring every applicant, but can it do so without having to make too many hires?

In this example, the department is hoping to find what Elkind et al. [2015] called a *Condorcet winning set*: a committee of candidates S such that no candidate a is preferred over all members of S by a majority of voters. More formally, we have the following definition.

Definition 3.1. A set S of candidates is a *Condorcet winning set* if for all candidates $a \notin S$, $\frac{1}{n}|a \succ S| < \frac{1}{2}$. The *Condorcet dimension* of an election is the size of its smallest Condorcet winning set.

Elkind et al. [2015] asked: how large can the Condorcet dimension of an election be? Condorcet's paradox shows that there are elections with Condorcet dimension at least 2, and Elkind et al. [2015] constructed a simple example of an election with Condorcet dimension at least 3. On the positive side, they showed that every election with m candidates has Condorcet dimension at most $\lceil \log_2 m \rceil + 1$, and conjectured that the Condorcet dimension must grow with the number of candidates in the election.

This conjecture turned out to be false. Charikar et al. [2025] observed that earlier work by Jiang et al. [2020] on approximately stable committees (see Section 3.2) implies that the Condorcet dimension of any election is at most 32, and introduced some new ideas to improve the result to 6. This bound was further sharpened to 5 by Song et al. [2026], whose substantially simpler proof remains the state of the art.

Theorem 3.2 (Song et al. [2026]). *Every election has Condorcet dimension at most* 5.

Charikar et al. [2025] conjectured that these results can be pushed even farther, and that the lower bound of Elkind et al. [2015] is tight.

Conjecture 3.3 (Charikar et al. [2025]). *In every election, there exists a committee of at most k candidates such that for all candidates $a \notin S$,*

$$\tfrac{1}{n}|a \succ S| < \frac{2}{k+1}.$$

In particular, all elections have Condorcet dimension at most 3.

In this section, we will start by motivating the common ideas in the proofs from Jiang et al. [2020], Charikar et al. [2025], and Song et al. [2026]. To demonstrate the proof strategy, we will give a simple proof that the Condorcet dimension of any election is at most 15 (following Jiang et al. [2020] with minor modifications). For the interested reader, we have written a short companion note with a full proof of Theorem 3.2 [Charikar et al. 2026b].

At a very high level, the three papers all proceed via the probabilistic method: construct a certain distribution D over committees or candidates, and argue that with positive probability, we can get a Condorcet winning set by sampling from D. One natural starting point would be to use a stable lottery as D, since Theorem 2.2 says that voters tend to prefer samples from $D_{k\text{-SL}}$ over any candidate a (in expectation, $\frac{1}{n}|a \succ S|$ is small when $S \sim D_{k\text{-SL}}$). To get a single committee that has this kind of property against all candidates, a naive approach would be to use a concentration bound to get a high-probability guarantee for each candidate a, and then union bound over all candidates. However, one quickly finds that it is impossible to shake a dependence on the number of candidates, even if the sharpest concentration bounds could apply.

The key idea is this: compare the sampled committee S against D *itself*, instead of each candidate a.

Building on this idea, all three papers use the following general proof strategy.

(1) Construct a distribution D over committees (or candidates) such that voters tend to prefer samples from D over any candidate a.
(2) Show that if we sample the committee S from D, with positive probability, voters tend to prefer S over samples from D.
(3) Argue that if voters tend to prefer S over samples from D, and samples from D over a, then a majority of voters cannot prefer a over S.

The distribution used by Jiang et al. [2020] is precisely a stable lottery. Charikar et al. [2025] use an equilibrium distribution derived from a similar game, but with some modifications to the payoffs that complement the way the distribution is used in the proof and allow for sharper bounds. Song et al. [2026] use a distribution derived from the equilibrium of a carefully designed market rather than a game (in particular, their distribution adapts the well-studied *Lindahl equilibrium*).

To distill the way that voters compare candidates and committees against a distribution D, we use the notion of *rank*, introduced by Charikar et al. [2025].

Given a distribution D over committees, voter v's rank of candidate a with respect to D is the probability that v prefers a over a committee sampled from D.[6] Formally,

$$\operatorname{rank}_v(a; D) \coloneqq \Pr_{S' \sim D}[a \succ_v S'].$$

To define rank for committees, we extend each voter's preference $\succ_v$ over candidates to a preference over committees as follows. For committees $S \neq T$, we say that $S \succ_v T$ if and only if v prefers their favorite candidate in $S \setminus T$ over their favorite candidate in $T \setminus S$. In another view, each voter v can encode a committee S as a binary string $s_v \in \{0,1\}^m$ such that $s_v(i) = 1$ if and only if v's ith favorite candidate is in S. Voter v's preference over committees is precisely the lexicographic order of the corresponding binary strings.

[6] The intention with this terminology is that "higher ranked" candidates are more preferred, while "lower ranked" candidates are less preferred, though we acknowledge the potential for confusion since colloquially, high ranks can correspond to low numerical ranks.

We then define the rank of a committee S as the probability that v *weakly* prefers S over a committee drawn from D. Formally,

$$\mathrm{rank}_v(S; D) \coloneqq \Pr_{S' \sim D}[S \succeq_v S'].$$

The ranks are a convenient primitive to work with because they encapsulate the voter's preferences over candidates and committees in a way that neatly fits the demands of the proof template. In particular, one can check that the ranks satisfy the following two properties.

PROPOSITION 3.4. *The ranks are consistent with preferences: if* $a \succ_v S$ *then* $\mathrm{rank}_v(a; D) \geq \mathrm{rank}_v(S; D)$.

PROPOSITION 3.5. *The random variable* $\mathrm{rank}_v(S; D)$ *with* $S \sim D$ *stochastically dominates* $r \sim \mathrm{Unif}(0, 1)$.

Proposition 3.4 allows us to use the ranks to interface with voters' preferences over candidates and committees. For example, if we can show that $\mathrm{rank}_v(a; D)$ is small for most voters and $\mathrm{rank}_v(S; D)$ is large for most voters, then we can argue that most voters cannot prefer a over S. Proposition 3.5 is critical in the second step of the proof: it gives us a simple way of understanding the distribution of $\mathrm{rank}_v(S; D)$ when S is sampled from D.

The easiest way to understand why the ranks satisfy the two properties above is with a picture.

Imagine creating a block for each committee S, whose width is the probability mass of S in D, and having each voter arrange these blocks along the interval $[0, 1]$ in increasing order of preference (see Figure 1 for an example). Then $\mathrm{rank}_v(S; D)$ is precisely the top (rightmost point) of the block corresponding to S, and $\mathrm{rank}_v(a; D)$ sits just above the blocks of committees that only consist of candidates that v does not prefer to a. This view explains why we define ranks for candidates and committees as we did, making the ranks for candidates as small as possible, and the ranks for committees as large as possible, while ensuring that if $a \succ_v S$ then $\mathrm{rank}_v(a; D) \geq \mathrm{rank}_v(S; D)$ (precisely Proposition 3.4).

Fig. 1. Depiction of the ranks for a voter v with preference $a \succ b \succ c \succ d$ with respect to a distribution D which chooses the committees $\{c, d\}, \{a, b\}, \{b, c\}, \{a, c\}$ with probabilities $0.1, 0.2, 0.3, 0.4$ respectively. The yellow dot represents $r \sim \mathrm{Unif}(0, 1)$, which can be used to sample from D.

The pictorial representation also gives us a simple way of interpreting samples from D. From each voter's perspective, $S \sim D$ is equivalent to sampling a real number $r \sim \mathrm{Unif}(0, 1)$ and choosing the committee whose block the number r lands in. (Formally, choose the committee S whose $\mathrm{rank}_v(S; D)$ is minimal but at least r.) Proposition 3.5 is immediate from this perspective.

We are now ready to give a constant upper bound on the Condorcet dimension of any election.

Theorem 3.6. *Every election has Condorcet dimension at most* 15.

Proof. First, we claim that the average rank of any candidate with respect to a stable k-lottery is at most $\frac{1}{k+1}$. That is, if $D_{k\text{-SL}}$ is the distribution of a stable k-lottery, then for all candidates a,

$$\frac{1}{n}\sum_{v\in V}\mathrm{rank}_v(a; D_{k\text{-SL}}) \le \frac{1}{k+1}. \tag{1}$$

This claim follows by a straightforward translation of Theorem 2.2. In particular for any distribution D over committees, we have

$$\mathop{\mathbb{E}}_{S\sim D}[\tfrac{1}{n}|a\succ S|] = \mathop{\mathbb{E}}_{S\sim D}\left[\frac{1}{n}\sum_{v\in V}\mathbf{1}[a\succ_v S]\right] = \frac{1}{n}\sum_{v\in V}\Pr_{S\sim D}[a\succ_v S] = \frac{1}{n}\sum_{v\in V}\mathrm{rank}_v(a;D).$$

With these equalities, (1) is equivalent to the statement of Theorem 2.2.

Next, we claim that some committee S in the support of the stable lottery has high ranks for a large fraction of voters. In fact, for any distribution over committees D, a random committee $S\sim D$ has $\mathrm{rank}_v(S;D) > \frac{1}{4}$ for at least a $\frac{3}{4}$ fraction of voters v in expectation. This claim follows immediately from the fact that $\mathrm{rank}_v(S;D)$ with $S\sim D$ stochastically dominates $r\sim \mathrm{Unif}(0,1)$, since r is greater than $\frac{1}{4}$ with probability $\frac{3}{4}$. Figure 2 may be helpful for visual intuition.

Finally, we claim that if $k\ge 15$, then this S (with $D_{k\text{-SL}}$ in place of D) is a Condorcet winning set. Suppose towards a contradiction that for some candidate a, we have that $\frac{1}{n}|a\succ S|\ge\frac{1}{2}$. Since $\mathrm{rank}_v(S;D_{k\text{-SL}}) > \frac{1}{4}$ for at least a $\frac{3}{4}$ fraction of voters v, it means that for at least a $\frac{1}{4}$ fraction of voters v, we have that $\mathrm{rank}_v(S;D_{k\text{-SL}}) > \frac{1}{4}$ *and* $a\succ_v S$. But then, $\mathrm{rank}_v(a;D_{k\text{-SL}}) > \frac{1}{4}$ for this $\frac{1}{4}$ fraction of voters v. It follows that

$$\frac{1}{n}\sum_{v\in V}\mathrm{rank}_v(a;D_{k\text{-SL}}) > \frac{1}{4}\cdot\frac{1}{4} = \frac{1}{16},$$

which contradicts (1) for $k\ge 15$. □

3.2 Approximately stable committees

While Condorcet winning sets focus on finding the best possible *small* committees, one can hope to get strong asymptotic results as a function of the committee size as well. The gold standard is for a committee to be *stable* (following the classic notion of *core stability* from cooperative game theory [Moulin 1988]). The normative basis for stable committees is that if we are to choose a committee S of size k, then any coalition of $\frac{1}{k}$ fraction of voters deserves to pick one candidate on the committee. As such, if there is a candidate a that a $1/k$ fraction of voters prefers over each candidate in S, then the committee S is considered *unstable*. (Perhaps these voters would want to break away from the larger group, and make their decisions separately, like an underrepresented academic community forming a new conference.) If there is no such candidate a, then S is stable.

Definition 3.7. A committee S of size k is *stable* if for all candidates $a\notin S$, $\frac{1}{n}|a\succ S| < \frac{1}{k}$.

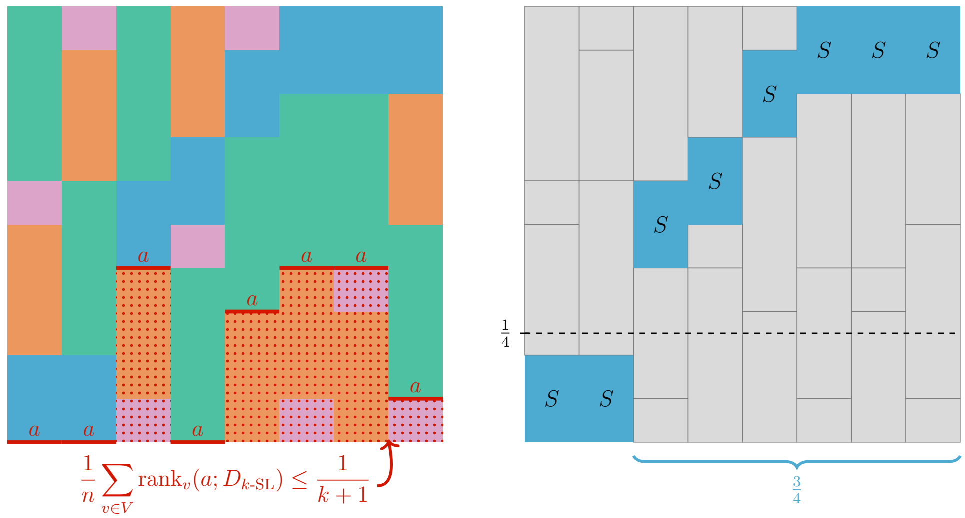

Fig. 2. A visual depiction of the first two steps of the proof. Within the square, each column represents a voter. Each committee is a block in the column, with height corresponding to their probability mass in $D_{k\text{-SL}}$, ordered by the preference of the voter (higher is more preferred). Each candidate a sits at height $\mathrm{rank}_v(a; D_{k\text{-SL}})$ in the column for voter v. The left diagram shows how each candidate is ranked low: the area under each a is at most $\frac{1}{k+1}$. The right shows how some committee S must be high: the top of block S reaches height above a quarter for three quarters of the voters.

Note that the definition is identical to that of a Condorcet winning set (Definition 3.1), but with the majority $\frac{1}{2}$ threshold replaced with $\frac{1}{k}$. The aforementioned counterexample of Elkind et al. [2015] against Condorcet winning sets of size 2 also shows that for each $k > 1$, stable committees of size k need not exist.

Instead, Jiang et al. [2020] proposed the following notion of *approximately* stable committees, and showed that they always exist.

THEOREM 3.8 (JIANG ET AL. [2020]). *A committee S of k candidates is a c-*stable committee *if for all candidates a,*

$$\tfrac{1}{n}|a \succ S| < \frac{c}{k}.$$

There exists a universal constant c such that for all positive integers k, every election has a c-stable committee of size at most k.

Jiang et al. [2020] proved this theorem with the particular constant $c = 16$. It is not hard to see that a c-stable committee of size $2c$ is a Condorcet winning set, and so this result is already enough to establish that all elections have Condorcet dimension at most 32. With some additional ideas[7] the correspondence can be made to work both ways. Charikar et al. [2025] improve the constant c to 9.82 and the techniques of Song et al. [2026] can be used to improve it to 4.91.[8]

[7] The key idea is to sample part of the committee, and then recurse on the subset of voters that are unhappy with the sampled committee (have low ranks).

[8] Song et al. [2026] cite a bound of 5.03 from prior work of Song and Nguyen [2026], but this bound can be improved marginally; see Charikar et al. [2026b] for more details.

Finally, we briefly mention that the interest in (approximately) stable committees extends far beyond the setting of ranked preferences that we focus on in this survey. With any monotone preference structure over committees (meaning $S \succeq_v S'$ if $S' \subseteq S$), the notion of approximately stable committees naturally extends as follows.

Definition 3.9. A committee S of k candidates is a *c-stable committee* if for all committees S' of k' candidates, $\frac{1}{n}|S' \succ S| < c \cdot \frac{k'}{k}$.

Intuitively, a coalition of $\frac{k'}{k}$ fraction of voters "deserves" k' candidates on the committee, so if they prefer some committee S' of size k' over S, then S is unstable. Jiang et al. [2020] proved a general result that for *any* monotone preferences, 32-stable committees of size at most k always exist.

Along with ranked preferences, *approval* preferences are a special case of significant interest. In this setting, each voter v has a subset A_v of candidates that they approve. Voters prefer committees that have more candidates that they approve (i.e., $S \succeq_v S'$ if $|S \cap A_v| \geq |S' \cap A_v|$.) In contrast to the ranking setting, it remains a major open question whether an exactly stable committee always exists (see, e.g., Peters [2025] and Becker et al. [2026]). It is known that 3.65-stable committees exist, due to Gao et al. [2026].

3.3 Approximately dominating sets

A common feature of Condorcet winning sets and approximately stable committees is that they compare a committee *collectively* against each candidate, but in some applications the resulting conditions are far too weak. For example, it would not be reasonable to justify the set of accepted papers for a conference by arguing that less than half of the program committee would think any rejected paper would be a best paper candidate if it were accepted. Instead we could aim for a stronger condition which asks voters to compare excluded candidates *head-to-head* against a single candidate on the committee. One natural proposal would be to choose a committee that is a *dominating set*.

Definition 3.10. A committee S is a *dominating set* if for each candidate $a \notin S$, there exists $b \in S$ such that $\frac{1}{n}|a \succ b| \leq \frac{1}{2}$.

The terminology comes from graph theory, where a dominating set is a subset of vertices S such that every vertex has a neighbor in S. If one considers the *majority graph* of an election, which is the directed graph whose vertices are candidates and where we have a directed edge $a \to b$ if a majority of voters prefers a over b, then the two notions of dominating sets are aligned.

Unfortunately, unlike Condorcet winning sets and approximately stable committees, dominating sets may need to be arbitrarily large. This fact is a consequence of two classic results. The first is McGarvey's theorem [McGarvey 1953], which shows that *any* tournament graph[9]

is the majority graph of some election. The second is a result due to Erdős [1963] that in random tournament graphs on m vertices, the smallest dominating sets have size $\Theta(\log m)$.

[9]A directed graph obtained by assigning directions to the edges of a complete graph.

Instead of relaxing the condition by allowing collective domination, perhaps we can relax the $\frac{1}{2}$ threshold in the definition of a dominating set.

Definition 3.11 *(Approximately dominating sets).* A committee S is a $(\frac{1}{2}+\varepsilon)$-*dominating set* if for each candidate $a \notin S$, there exists $b \in S$ such that $\frac{1}{n}|a \succ b| \leq \frac{1}{2}+\varepsilon$.

It turns out that for any constant ε, $(\frac{1}{2}+\varepsilon)$-dominating sets of constant size exist. In particular, the following theorem was proven independently by Bourneuf et al. [2025] and Charikar et al. [2026a].

THEOREM 3.12. *Every election has a $(\frac{1}{2}+\varepsilon)$-dominating set of size $O(1/\varepsilon^2)$.*

Intriguingly, the techniques in the two papers are substantially different. Bourneuf et al. [2025] prove the result as one of many applications of a new tool they call the *dense neighborhood lemma*, which bounds a variant of the VC-dimension in several combinatorial structures. We will focus on the proof from Charikar et al. [2026a], which constructs *approximate* maximal lotteries with small support size.

THEOREM 3.13 (APPROXIMATE MAXIMAL LOTTERIES). *In any election, there exists a distribution D over candidates with support size at most $(1+o(1))\frac{\pi}{8\varepsilon^2}$ such that for all candidates $a \in C$,*

$$\mathop{\mathbb{E}}_{b\sim D}[\tfrac{1}{n}|a \succ b|] \leq \frac{1}{2}+\varepsilon.$$

*A distribution satisfying this condition is called an ε-*approximate maximal lottery.

It is not hard to see that this theorem generalizes Theorem 3.12, since the support of D must be a $(\frac{1}{2}+\varepsilon)$-dominating set. On their own, approximate maximal lotteries can be an attractive solution concept even in comparison to dominating sets. In choosing papers for a conference, for example, it may still be unreasonable to justify rejecting a paper by arguing that some accepted paper is preferred by most program committee members. But if many members would think the rejected paper is worse than the *average* accepted paper (or even the average accepted *borderline* paper, since the other decisions are straightforward), rejection seems much more defensible.

Approximate maximal lotteries are also interesting for their connection to approximate Nash equilibria. In fact, an approximate maximal lottery exactly corresponds to an approximate Nash equilibrium in the gambler game at the start of Section 2. However, while ε-approximate Nash equilibria have support size $\Theta(\frac{\log m}{\varepsilon^2})$ in the worst case (where m is the number of pure strategies), the particular structure of elections allows for approximate maximal lotteries to avoid the $\log m$ factor.

At a high level, the proof works by arguing that if we draw k samples from a maximal lottery, then with positive probability, the empirical distribution over the samples is an $O(1/\sqrt{k})$-approximate maximal lottery. Like with Condorcet winning sets in Section 3.1, a naive argument would use concentration inequalities and a union bound over candidates, but a logarithmic dependence on m appears unavoidable. Once again, one can get a better handle on the structure of voters' preferences over candidates using the ranks, and get a bound independent of the number of candidates with a simple proof.

Finally, we mention that the right dependence on ε for the support size of approximate maximal lotteries (and $(\frac{1}{2}+\varepsilon)$-dominating sets) remains open. Charikar et al. [2026a] explain that a construction of Alon et al. [2006] implies a lower bound of $\Omega(\frac{1}{\varepsilon})$, and recently Lin et al. [2026] sharpened this lower bound to $\frac{1}{2\varepsilon}$. Charikar et al. [2026a] conjectured that the upper bound is tight.

Conjecture 3.14 (Charikar et al. [2026a]). *For each $\varepsilon > 0$, there is an election in which all $(\frac{1}{2}+\varepsilon)$-dominating sets have size $\Omega(1/\varepsilon^2)$.*

4. APPLICATION 2: DISTORTION

One takeaway from the history of social choice theory is that no voting rule is perfect, and the best we can hope is to find voting rules that are reasonably good. In other words, if we cannot always make optimal choices, can we at least make *approximately* optimal choices?

The *distortion* framework, introduced by Procaccia and Rosenschein [2006], gives a way of understanding what it means for a voting rule to be approximately optimal. In this setting, we imagine that the voters have some cardinal utilities (or costs) for the candidates. If we knew these cardinal values, the optimal choice would be the candidate that maximizes total social welfare (or minimizes the total social cost). However, the voting rule only has access to the *ordinal* preferences of the voters. With only this limited information, can a voting rule always choose an approximately optimal candidate? What is the best approximation factor (called the *distortion*) that can be guaranteed? (See Figures 3 and 4 for illustrative examples.)

Under completely arbitrary utilities or costs, the distortion can be unbounded, but with mild structural assumptions, it is possible to prove meaningful positive results. The literature has focused on two structural models that have each received substantial attention over the last two decades: *utilitarian distortion* and *metric distortion*. At a high level, utilitarian distortion imposes very mild normalization assumptions that voters' utilities are on a similar scale, while metric distortion imagines that the preferences are derived from the distances in a metric. The tradeoff is that in utilitarian distortion, the structural assumption is weaker, but the distortion guarantees still blow up as the number of candidates grows, while metric distortion imposes stronger structural assumptions and gets distortion guarantees that stay constant with any number of candidates.

Instead of covering the vast literature (for which we defer to a survey of Anshelevich et al. [2021]), we will focus on the slice of results that crucially rely on maximal lotteries and stable lotteries. As it happens, these results are also the state of the art for some of the most central questions in both utilitarian distortion and metric distortion.

4.1 Utilitarian distortion

In *utilitarian distortion* [Procaccia and Rosenschein 2006], each voter v has an underlying utility function $u_v : C \to \mathbb{R}_{\geq 0}$ satisfying the condition that if $a \succ_v b$ then $u_v(a) \geq u_v(b)$. The utility functions are also assumed to be normalized,[10] typically so that each voter's utilities sum to one ($\sum_{a \in C} u_v(a) = 1$ for each voter

[10] We refer the reader to Aziz [2019] for compelling normative arguments in favor of normalization.

v). The *social welfare* of a candidate a is $\mathsf{SW}(a) := \sum_{v \in V} u_v(a)$, and a voting rule has distortion α if it chooses a candidate b such that $\mathsf{SW}(b) \geq \frac{1}{\alpha} \max_{a \in C} \mathsf{SW}(a)$ in any election, and with any underlying utility functions that are consistent with the preferences. For a randomized voting rule that chooses a distribution D over candidates, we replace $\mathsf{SW}(b)$ with $\mathbb{E}_{b \sim D}[\mathsf{SW}(b)]$.

Preferences

$v_1 : a \succ b \succ c$

$v_2 : b \succ c \succ a$

$v_3 : c \succ a \succ b$

Possible Utilities

	a	b	c
v_1	1/3	1/3	1/3
v_2	0	1/2	1/2
v_3	0	0	1

Social Welfare

$a :$ 1/3

$b :$ 5/6

$c :$ 11/6

Fig. 3. An example of the utilitarian distortion in an election with three candidates and three voters with cyclical preferences. The choice of utilities in the middle results in candidate a having distortion 11/2.

Procaccia and Rosenschein [2006] showed that for deterministic voting rules, the plurality rule (which chooses the most common first choice of the voters) achieves the best possible utilitarian distortion of $\Theta(m^2)$.

As such, the problem is most interesting for *randomized* voting rules. Trivially, choosing a uniformly random candidate has distortion $O(m)$, but this is far from optimal. Boutilier et al. [2015] designed an intricate rule with distortion $O(\sqrt{m} \log^* m)$ and showed that any randomized voting rule has distortion at least $\Omega(\sqrt{m})$. This $\log^* m$ gap persisted for several years, until Ebadian et al. [2024] showed that a simple rule which mixes between a random candidate and a stable lottery gets distortion $O(\sqrt{m})$.

Theorem 4.1 (Ebadian et al. [2024]). *Consider the following voting rule: with probability $\frac{1}{2}$, choose a uniformly random candidate, and with probability $\frac{1}{2}$, sample a committee S from a stable $\sqrt{m}$-lottery, and choose a uniformly random candidate from S. This rule has distortion $O(\sqrt{m})$ in any election with unit-sum utilities.*

Note that Ebadian et al. [2024] prove the same result for a much broader class of utility functions, but we focus on the unit-sum case for simplicity.

Here is some rough intuition for the argument. Imagine that we just sample a committee S of size k from a stable k-lottery, and then choose a uniformly random candidate from S. Since we expect that nearly all voters ($\frac{k}{k+1}$ fraction) will prefer some candidate in S over the optimal candidate $a^\star$, we can argue that a uniformly random candidate from S captures nearly a $1/k$ fraction of the welfare of $a^\star$. There is an additive loss of at most $\frac{n}{k(k+1)}$, coming from the voters who prefer $a^\star$ over S. (The factor $1/k$ in both these terms comes from the fact that we choose one of the k candidates of S at random.) If we set $k \approx \sqrt{m}$ then the roughly n/m additive loss can be made up for by mixing with a uniform lottery, which alone always captures utility at least n/m (since the total utility of all candidates is n). For the curious reader, the full proof is below.

PROOF. Suppose that $a^\star$ is the candidate with the maximum social welfare. Let D_S be the distribution that chooses a uniformly random candidate from the committee S of size k. The key claim is that

$$\mathsf{SW}(D_S) \geq \frac{\mathsf{SW}(a^\star) - |a^\star \succ S|}{k}.$$

If a voter v prefers some candidate $b \in S$ over $a^\star$ (i.e., $a^\star \not\succ_v S$), then we have $u_v(a^\star) \leq u_v(b)$. Even more loosely, these voters satisfy

$$u_v(a^\star) \leq \sum_{b \in S} u_v(b).$$

For the remaining voters v satisfying $a^\star \succ_v S$, we can default to the weak bound $u_v(a^\star) \leq 1$. Putting the two together, we have

$$\begin{aligned}
\mathsf{SW}(a^\star) &= \sum_{v : a^\star \succ_v S} u_v(a^\star) + \sum_{v : a^\star \not\succ_v S} u_v(a^\star) \\
&\leq |a^\star \succ S| + \sum_{v \in V} \sum_{b \in S} u_v(b) \\
&= |a^\star \succ S| + \sum_{b \in S} \mathsf{SW}(b).
\end{aligned}$$

Since $\mathsf{SW}(D_S) = \frac{1}{k} \sum_{b \in S} \mathsf{SW}(b)$, the claim follows after rearranging.

Let $D_{k\text{-SL}}$ be the distribution of a stable k-lottery, and let D_k be the distribution which samples $S \sim D_{k\text{-SL}}$ and then chooses a uniformly random candidate from S. Using the guarantee of Theorem 2.2, we get

$$\mathsf{SW}(D_k) = \mathop{\mathbb{E}}_{S \sim D_{k\text{-SL}}}[\mathsf{SW}(D_S)] \geq \frac{\mathsf{SW}(a^\star) - \mathbb{E}_{S \sim D_{k\text{-SL}}}[|a^\star \succ S|]}{k} \geq \frac{\mathsf{SW}(a^\star)}{k} - \frac{n}{k(k+1)}.$$

In particular, $\mathsf{SW}(D_{\sqrt{m}}) \geq \frac{1}{\sqrt{m}} \mathsf{SW}(a^\star) - \frac{n}{m}$. Finally, the social welfare of the uniform distribution over candidates D_{Unif} is exactly n/m, since

$$\mathsf{SW}(D_{\text{Unif}}) = \frac{1}{m} \sum_{a \in C} \sum_{v \in V} u_v(a) = \frac{1}{m} \sum_{v \in V} \sum_{a \in C} u_v(a) = \frac{n}{m}$$

by the unit-sum assumption. Averaging over the two distributions, we get

$$\mathsf{SW}(\tfrac{1}{2} D_{\sqrt{m}} + \tfrac{1}{2} D_{\text{Unif}}) \geq \frac{1}{2} \left(\frac{1}{\sqrt{m}} \cdot \mathsf{SW}(a^\star) - \frac{n}{m} \right) + \frac{1}{2} \cdot \frac{n}{m} = \frac{1}{2\sqrt{m}} \cdot \mathsf{SW}(a^\star)$$

which gives a distortion bound of $O(\sqrt{m})$ as claimed. □

4.2 Metric distortion

In *metric distortion* [Anshelevich et al. 2015], we imagine that the voters and candidates lie in a metric space with a distance metric $d(\cdot, \cdot)$. Voters prefer candidates that are close to them, meaning that if $a \succ_v b$ then $d(v, a) \leq d(v, b)$. This idea is inspired by a vast literature on proximity spatial models of voting [Enelow and Hinich 1984; Enelow and Hinich 1990; Merrill and Grofman 1999; Armstrong et al. 2020], which includes concepts like the "political spectrum." The *social cost* of a candidate a is $\mathsf{SC}(a) := \frac{1}{n}\sum_{v\in V} d(a,v)$, and a voting rule has distortion α if it chooses a candidate b such that $\mathsf{SC}(b) \le \alpha \min_{a\in C} \mathsf{SC}(a)$ in any election, and with any underlying metric space that is consistent with the preferences. For a randomized voting rule that chooses a distribution D over candidates, we replace $\mathsf{SC}(b)$ with $\mathbb{E}_{b\sim D}[\mathsf{SC}(b)]$.

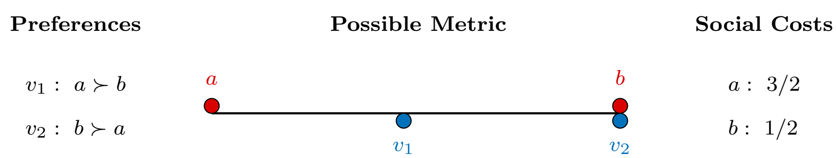

Fig. 4. An example of the metric distortion in an election with two candidates and two disagreeing voters. Candidates and voters lie on a line with a and b at -1 and 1, and one voter is impartial (at 0) while the other is collocated with their preferred candidate. In this case, one candidate has 3 times the cost of the other.

In contrast to the utilitarian distortion model, two differences stand out in the literature on metric distortion. First, despite the broad array of possible metric spaces, their structure is enough for several voting rules to have *constant* distortion, independent of the number of candidates. On the other hand, previously studied voting rules all fall short of the *optimal* distortion, and successive improvements in the best known distortion have come from designing several natural new voting rules.

Anshelevich et al. [2015] initially showed that all deterministic voting rules have distortion at least 3 (due to the simple example in Figure 4), and a long line of work culminated in proving a matching upper bound with simple voting rules [Gkatzelis et al. 2020; Kizilkaya and Kempe 2022; Kizilkaya and Kempe 2023]. Beyond these results, 3 turns out to be a critical threshold for the metric distortion problem as a whole. In particular, a distortion 3 candidate is in a way analogous to a Condorcet winner,[11] and other important classes of voting rules also have optimal distortion 3.[12]

With the distortion of deterministic voting rules resolved, attention turned to *randomized* voting rules. It turns out that getting distortion 3 in this setting is easy: Anshelevich and Postl [2017] and Feldman et al. [2016] independently showed that Random Dictator (which chooses a uniformly random voter's favorite candidate) achieves this. Both papers also noted that no randomized voting rule can have distortion less than 2 (by the example in Figure 4), which was conjectured to be optimal [Goel et al. 2017]. This conjecture was refuted independently by Charikar and Ramakrishnan [2022] and Pulyassary and Swamy [2021], with the former giving the best known lower bound of 2.112. Finally, Charikar et al. [2024] designed a new randomized voting rule with distortion 2.753, separating the best possible distortion of deterministic and randomized voting rules and breaking the barrier of 3.

[11]If a majority of voters prefers b over a then $\mathsf{SC}(b) \le 3\mathsf{SC}(a)$.

[12]Rules that exclusively choose from top-ranked candidates, and those that only need the aggregate preferences between pairs of candidates (called *weighted tournament rules*).

The voting rule used by Charikar et al. [2024] works by randomizing between a maximal lottery and a new voting rule called RaDiUS, which chooses a random voter's favorite candidate in the *Weighted Uncovered Set* (introduced by Munagala and Wang [2019]). Rather than detail the technical proof, we will highlight the way in which maximal lotteries are uniquely well-suited to this application.

Charikar et al. [2024] start by pinning down the distortion of a maximal lottery in isolation.

THEOREM 4.2 (CHARIKAR ET AL. [2024]). *In any election, a maximal lottery has distortion 3.*

Theorem 4.2 is not surprising if one thinks of a maximal lottery as a randomized analogue of a Condorcet winner, but the proof is not as straightforward as this correspondence would suggest. Instead of the full proof, we give an explanation of the additional structure the proof implies that is not present in proofs that other voting rules have distortion 3.

Suppose that our goal is to show that $\mathsf{SC}(b) \leq 3\mathsf{SC}(a^\star)$, or equivalently that $\mathsf{SC}(b) - \mathsf{SC}(a^\star) \leq 2\mathsf{SC}(a^\star)$ (for technical reasons, this rearrangement is cleaner to work with).

We can use the trick of expressing a mean as an integral over tails to write

$$\mathsf{SC}(b) - \mathsf{SC}(a^\star) = \int_0^\infty \Pr_{v\sim V}[d(v,b) - d(v,a^\star) > t]\,\mathrm{d}t$$

and

$$2\mathsf{SC}(a^\star) = \int_0^\infty \Pr_{v\sim V}[2d(v,a^\star) > t]\,\mathrm{d}t.$$

For the sake of clarity, we point out that in the first equation, it would be more correct to have the integral range from $-\infty$ to ∞, or for the $=$ to be replaced by $\leq$. It turns out that in the worst-case metrics, equality actually holds ($d(v,b) \geq d(v,a^\star)$ for all v). Also, for a distribution D in place of b, we would instead write

$$\mathop{\mathbb{E}}_{b\sim D}[\mathsf{SC}(b)] - \mathsf{SC}(a^\star) = \int_0^\infty \mathop{\mathbb{E}}_{b\sim D}\left[\Pr_{v\sim V}[d(v,b) - d(v,a^\star) > t]\right]\mathrm{d}t.$$

The proofs that earlier voting rules[13] have distortion 3 can all be interpreted as proving something stronger than $\mathsf{SC}(b) - \mathsf{SC}(a^\star) \leq 2\mathsf{SC}(a^\star)$. They actually show that for all $t \geq 0$,

$$\Pr_{v\sim V}[d(v,b) - d(v,a^\star) > t] \leq \Pr_{v\sim V}[2d(v,a^\star) > t]. \tag{2}$$

For readers familiar with this line of work, we take a beat to note that these proofs are not typically written in terms of (2). The usual objective is to construct a perfect matching M between two copies of the set of voters V such that for each pair $(v,u) \in M$, v weakly prefers b over u's favorite candidate. It is not hard to use the triangle inequality to show that such a pair must satisfy $d(v,b) - d(v,a^\star) \leq 2d(u,a^\star)$,

[13] Random Dictator [Anshelevich and Postl 2017; Feldman et al. 2016], Plurality Matching [Gkatzelis et al. 2020], Plurality Veto [Kizilkaya and Kempe 2022], Simultaneous Plurality Veto [Kizilkaya and Kempe 2023].

and with a little more effort, one can see that constructing the matching M is equivalent to showing (2).

The *pièce de résistance* is that the proof that maximal lotteries have distortion 3 shows that in the *worst-case* metrics (termed *biased* metrics by Charikar and Ramakrishnan [2022]), something even stronger than (2) holds. For all $t \geq 0$

$$\mathop{\mathbb{E}}_{b \sim D_{\mathrm{ML}}} \left[\Pr_{v \sim V} [d(v, b) - d(v, a^\star) > t] \right] \leq \min \left(\frac{1}{2}, \Pr_{v \sim V} [2d(v, a^\star) > t] \right). \tag{3}$$

To give a window into the proof, the $\frac{1}{2}$ in this expression is no coincidence. It comes directly from Theorem 2.1.

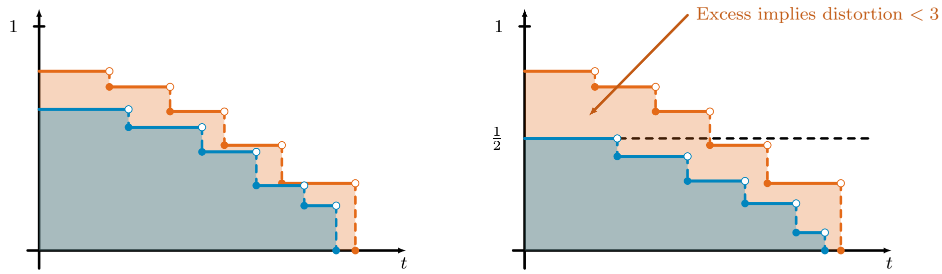

Fig. 5. The blue regions show $\mathsf{SC}(b) - \mathsf{SC}(a^\star) = \int_0^\infty \Pr_{v \sim V}[d(v, b) - d(v, a^\star) > t]\, \mathrm{d}t$, and the orange regions show $2\mathsf{SC}(a^\star) = \int_0^\infty \Pr_{v \sim V}[2d(v, a^\star) > t]\, \mathrm{d}t$. The diagram on the left illustrates the condition (2) usually satisfied when b is chosen by a voting rule with distortion at most 3, while the right illustrates the condition (3) satisfied when b is sampled from a maximal lottery.

The power of the stronger condition (3) is that it implies a precise structure in the metrics where maximal lotteries have high distortion (close to 3). Figure 5 helps show why: if the distortion of maximal lotteries is close to 3, then the function $\Pr_{v \sim V}[2d(v, a^\star) > t]$ (in orange) must be close to $\frac{1}{2}$ for small values of t, meaning that nearly half the voters are tightly clustered around the optimal candidate $a^\star$. The RaDiUS rule complements maximal lotteries by having particularly small distortion in these structured metrics, and so mixing the two voting rules together covers all the bases to beat 3.

Finally, we conclude by mentioning a couple of follow-up results that use stable lotteries in addition to maximal lotteries.

Charikar et al. [2025] study the metric distortion of k-*tournament* rules, whose input is the aggregate preferences of the voters for each set of k candidates (the quantities $\frac{1}{n}|a_1 \succ \cdots \succ a_k|$ for each $(a_1, \ldots, a_k) \in C^k$). They show that a k-tournament rule can have distortion $3 + \tilde{O}(\frac{1}{k^{1/4}})$ without randomness, and strictly less than 3 with randomness even for $k = 3$. Both results crucially make use of stable lotteries, leveraging the intuition that they are similar to random dictatorships for large k.

Cai et al. [2026] study the minimal amount of randomness needed to beat distortion 3, and show that it is possible while randomizing over $O(1)$ candidates. Their voting rule is again built on top of maximal lotteries and stable lotteries, using the fact that they can be approximated with small support (generalizing Theorem 3.13). An interesting open question raised by this work is whether a voting rule which randomizes between at most two candidates can have distortion less than 3.

Conjecture 4.3 (Cai et al. [2026]). *There exists a voting rule which randomizes between at most two candidates and has distortion* $3 - \varepsilon$ *for some fixed* $\varepsilon > 0$.

An exciting aspect of this conjecture is that resolving it plausibly requires completely new techniques.

5. CONCLUSION

Throughout this survey, we have seen myriad ways of using maximal lotteries and stable lotteries to prove positive results. To illustrate the point once again, even though Condorcet's paradox is one of the oldest and most fundamental challenges in voting theory, we discussed at least *four* different relaxations of a Condorcet winner that are universally achievable (maximal lotteries, Condorcet winning sets, approximately dominating sets, distortion 3 and below). The broad applicability of these tools is also incredibly exciting from the perspective of a theoretician. It suggests that they are not overly sensitive to the idiosyncrasies of a particular setting, and instead may capture some part of the more fundamental principles that the field broadly aims for, such as representation, efficiency, and fairness. For this reason, maximal lotteries, stable lotteries, and their relatives deserve a prominent place in any voting theorist's toolkit.

What is especially promising is that so many of these results are recent, and there is no indication that we have found all or even most of the domains of social choice where game-theoretic rules could be impactful. One possible domain on the horizon is pluralistic AI alignment. Recently, researchers have proposed alignment methods based on maximal lotteries, and several papers have made a compelling case for these approaches, both theoretically and experimentally [Munos et al. 2024; Maura-Rivero et al. 2025; Khalaf et al. 2026; Gölz et al. 2026].[14]

Finally, we hope that these new results and directions also motivate the community to step back and develop a deep and comprehensive theory of probabilistic tools in social choice. In doing so, we can find ideas that are flexible enough to apply across the diverse range of scenarios that voting can model, and benefit from the compounding effects of rich theory and positive results.

Acknowledgments. We thank editors-in-chief Nick Arnosti and Sam Taggart for helpful feedback that greatly improved the exposition of the survey.

REFERENCES

Alon, N., Brightwell, G., Kierstead, H. A., Kostochka, A. V., and Winkler, P. 2006. Dominating sets in k-majority tournaments. *Journal of Combinatorial Theory, Series B 96,* 3, 374–387.

[14] See also [Conitzer et al. 2024; Ge et al. 2024; Fish et al. 2026] for some broader agenda-setting work at the intersection of social choice and AI.

Anshelevich, E., Bhardwaj, O., and Postl, J. 2015. Approximating optimal social choice under metric preferences. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*. 777–783. 91

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.

Anshelevich, E. and Postl, J. 2017. Randomized social choice functions under metric preferences. *J. Artif. Intell. Res. 58*, 797–827. 93

Armstrong, D. A., Bakker, R., Carroll, R., Hare, C., Poole, K. T., and Rosenthal, H. 2020. *Analyzing spatial models of choice and judgment*. Chapman and Hall/CRC. 91

Aziz, H. 2019. Justifications of welfare guarantees under normalized utilities. *SIGecom Exch. 17,* 2, 71–75.

Becker, P., Greger, M., and Peters, D. 2026. Core existence in approval-based committee elections with up to five voter types.

Bourneuf, R., Charbit, P., and Thomassé, S. 2025. A dense neighborhood lemma: Applications of partial concept classes to domination and chromatic number. In *Proceedings of the 66th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 1–37.

Boutilier, C., Caragiannis, I., Haber, S., Lu, T., Procaccia, A. D., and Sheffet, O. 2015. Optimal social choice functions: A utilitarian view. *Artif. Intell. 227*, 190–213.

Brandl, F., Brandt, F., and Stricker, C. 2018. An analytical and experimental comparison of maximal lottery schemes. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence (IJCAI)*. ijcai.org, 114–120.

Brandt, F. 2017. Rolling the dice: Recent results in probabilistic social choice. In *Trends in computational social choice*. AI Access, Lulu, 3–26.

Cai, Z., Gao, D. D., Ramakrishnan, P., and Wang, K. 2026. Distortion of metric voting with bounded randomness. In *Proceedings of the 27th ACM Conference on Economics and Computation (EC)*.

Charikar, M., Lassota, A., Ramakrishnan, P., Vetta, A., and Wang, K. 2025. Six candidates suffice to win a voter majority. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing (STOC)*. ACM, 1590–1601.

Charikar, M. and Ramakrishnan, P. 2022. Metric distortion bounds for randomized social choice. In *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 2986–3004.

Charikar, M., Ramakrishnan, P., Tan, Z., and Wang, K. 2025. Metric distortion for tournament voting and beyond. In *Proceedings of the 26th ACM Conference on Economics and Computation (EC)*. ACM, 790–818.

Charikar, M., Ramakrishnan, P., and Wang, K. 2026a. Approximately dominating sets in elections. In *Proceedings of the 2026 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 1747–1760.

Charikar, M., Ramakrishnan, P., and Wang, K. 2026b. An exposition of five candidates suffice for a majority. 83

Charikar, M., Ramakrishnan, P., Wang, K., and Wu, H. 2024. Breaking the metric voting distortion barrier. *Journal of the ACM 71,* 6, 1–33.

Cheng, Y., Jiang, Z., Munagala, K., and Wang, K. 2020. Group fairness in committee selection. *ACM Transactions on Economics and Computation (TEAC) 8,* 4, 23:1–23:18. 80

Conitzer, V., Freedman, R., Heitzig, J., Holliday, W. H., Jacobs, B. M., Lambert, N., Mossé, M., Pacuit, E., Russell, S., Schoelkopf, H., Tewolde, E., and Zwicker, W. S. 2024. Position: social choice should guide AI alignment in dealing with diverse human feedback. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*. 9346–9360. 95

Ebadian, S., Kahng, A., Peters, D., and Shah, N. 2024. Optimized distortion and proportional fairness in voting. *ACM Transactions on Economics and Computation (TEAC) 12,* 1, 3:1–3:39.

Elkind, E., Lang, J., and Saffidine, A. 2015. Condorcet winning sets. *Social Choice and Welfare 44,* 3, 493–517.

Enelow, J. M. and Hinich, M. J. 1984. *The spatial theory of voting: An introduction*. CUP Archive. 91

Enelow, J. M. and Hinich, M. J. 1990. *Advances in the spatial theory of voting*. Cambridge University Press. 91

Erdős, P. 1963. On a problem in graph theory. *The Mathematical Gazette 47,* 361, 220–223.

Faliszewski, P., Skowron, P., Slinko, A., and Talmon, N. 2017. Multiwinner voting: A new challenge for social choice theory. *Trends in computational social choice 74,* 2017, 27–47.

Feldman, M., Fiat, A., and Golomb, I. 2016. On voting and facility location. In *Proceedings of the 2016 ACM Conference on Economics and Computation (EC)*. 269–286. 93

Fish, S., Gölz, P., Parkes, D., Procaccia, A., Rusak, G., Shapira, I., and Wuthrich, M. 2026. Generative social choice. *Journal of the ACM 73,* 2, 1–52. 95

Fishburn, P. 1984. Probabilistic social choice based on simple voting comparisons. *The Review of Economic Studies 51,* 4, 683–692. 79

Gao, D. D., Sun, Y., and Vondrák, J. 2026. Computation of approximately stable committees in approval-based elections. In *Proceedings of the 27th ACM Conference on Economics and Computation (EC)*.

Ge, L., Halpern, D., Micha, E., Procaccia, A. D., Shapira, I., Vorobeychik, Y., and Wu, J. 2024. Axioms for AI alignment from human feedback. In *Conference on Neural Information Processing Systems (NeurIPS)*. 95

Gehrlein, W. V. 2002. Condorcet's paradox and the likelihood of its occurrence: Different perspectives on balanced preferences. *Theory and Decision 52,* 2, 171–199.

Gkatzelis, V., Halpern, D., and Shah, N. 2020. Resolving the optimal metric distortion conjecture. In *Proceedings of the 61st IEEE Annual Symposium on Foundations of Computer Science (FOCS)*. 1427–1438. 92, 93

Goel, A., Krishnaswamy, A. K., and Munagala, K. 2017. Metric distortion of social choice rules: Lower bounds and fairness properties. In *Proceedings of the 2017 ACM Conference on Economics and Computation (EC)*. 287–304. 92

Gölz, P., Haghtalab, N., and Yang, K. 2026. Distortion of AI alignment: Does preference optimization optimize for preferences? *Advances in Neural Information Processing Systems (NeurIPS) 38*, 25969–26007. 95

Jiang, Z., Munagala, K., and Wang, K. 2020. Approximately stable committee selection. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing (STOC)*. ACM, New York, 463–472.

Khalaf, H., Wang, S. L., Halpern, D., Shapira, I., du Pin Calmon, F., and Procaccia, A. D. 2026. Robust AI evaluation through maximal lotteries. In *Proceedings of the 43rd International Conference on Machine Learning (ICML)*. 95

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. 92, 93

Kizilkaya, F. E. and Kempe, D. 2023. Generalized veto core and a practical voting rule with optimal metric distortion. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC)*. ACM, 913–936. 92, 93

Kreweras, G. 1965. Aggregation of preference orderings. In *Mathematics and Social Sciences I: Proceedings of the seminars of Menthon-Saint-Bernard, France (1–27 July 1960) and of Gösing, Austria (3–27 July 1962)*. 73–79. 79

Lackner, M. and Skowron, P. 2023. *Multi-winner voting with approval preferences*. Springer Nature.

Lin, Y., Qin, S., Wang, K., and Xia, L. 2026. Winning in the limit: Average-case committee selection with many candidates. In *Proceedings of the 27th ACM Conference on Economics and Computation (EC)*.

Maura-Rivero, R.-R., Lanctot, M., Visin, F., and Larson, K. 2025. Jackpot! Alignment as a maximal lottery. *arXiv preprint arXiv:2501.19266*. 95

McGarvey, D. C. 1953. A theorem on the construction of voting paradoxes. *Econometrica: Journal of the Econometric Society*, 608–610. 87

Merrill, S. and Grofman, B. 1999. *A unified theory of voting: Directional and proximity spatial models*. Cambridge University Press. 91

Moulin, H. 1988. *Axioms of Cooperative Decision Making*. Econometric Society Monographs, vol. 15. Cambridge University Press. 85

Munagala, K. and Wang, K. 2019. Improved metric distortion for deterministic social choice rules. In *Proceedings of the 2019 ACM Conference on Economics and Computation (EC)*. 245–262.

Munos, R., Valko, M., Calandriello, D., Azar, M. G., Rowland, M., Guo, Z. D., Tang, Y., Geist, M., Mesnard, T., Fiegel, C., Michi, A., Selvi, M., Girgin, S., Momchev, N., Bachem, O., Mankowitz, D. J., Precup, D., and Piot, B. 2024. Nash learning from human feedback. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*. 95

Peters, D. 2025. The core of approval-based committee elections with few seats. In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence (IJCAI)*. ijcai.org, 4014–4022.

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. 89

Pulyassary, H. and Swamy, C. 2021. On the randomized metric distortion conjecture. *arXiv preprint arXiv:2111.08698*.

Song, H. and Nguyen, T. 2026. Ordinal Lindahl equilibrium for voting. In *Proceedings of the 27th ACM Conference on Economics and Computation (EC)*.

Song, H., Nguyen, T., and Lin, Y. 2026. A few good choices. In *Proceedings of the 2026 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 4861–4874.

von Neumann, J. 1928. Zur theorie der gesellschaftsspiele. *Mathematische annalen 100,* 1, 295–320. 80