Smoothed Analysis of Social Choice Revisited

Bailey Flanigan
Carnegie Mellon University
bflaniga@andrew.cmu.edu

Daniel Halpern
Harvard University
dhalpern@g.harvard.edu

Alexandros Psomas
Purdue University
apsomas@cs.purdue.edu

Abstract

A canonical problem in voting theory is: which voting rule should we use to aggregate voters' preferences into a collective decision over alternatives? When applying the axiomatic approach to evaluate and compare voting rules, we are faced with prohibitive impossibilities. However, these impossibilities occur under the assumption that voters' preferences (collectively called a profile) will be worst-case with respect to the desired criterion.

In this paper, we study the axiomatic approach slightly *beyond* the worst-case: we present and apply a "smoothed" model of the voting setting, which assumes that while inputs (profiles) may be worst-case, all inputs will be perturbed by a small amount of noise. In defining and analyzing our noise model, we do not aim to substantially technically innovate on Lirong Xia's recently-proposed smoothed model of social choice; rather, we offer an alternative model and approach to analyzing it that aims to strike a different balance of simplicity and technical generality, and to correspond closely to Spielman and Teng's original work on smoothed analysis [28].

Within our model, we then give simple proofs of smoothed-satisfaction or smoothed-violation of several axioms and paradoxes, including most of those studied by Xia as well as some previously unstudied. Novel results include smoothed analysis of Arrow's theorem and analyses of the axioms Consistency and Independence of Irrelevant Alternatives. In independent work from Xia's recent paper [33], we also show the smoothed-satisfaction of coalition-based notions of Strategy-Proofness, Monotonocity, and Participation. A final, central component of our contributions are the high-level insights and future directions we identify based on this work, which we describe in detail to maximally facilitate additional research in this area.

1 Introduction

The defining question of voting theory is: how do we aggregate people's preferences into a collective decision? This is a broadly impactful question, with elections happening at all scales throughout society.

The task of aggregating people's preferences is performed by *voting rules* — functions that take as input a set of voters' preferences over alternatives, called a *preference profile*, and output a collectively-chosen alternative. To determine which voting rule to use, we must be able to evaluate and compare voting rules. In social choice, the canonical criteria used for such comparisons are *axioms*, and our analysis will focus primarily on this type of criterion. Other criteria considered more recently include distortion [25], time-complexity of computing winner(s) [12], and communication complexity of soliciting preferences [9].

These criteria all have in common that their satisfaction is traditionally determined from a *worst-case* perspective — i.e., a voting rule is said to satisfy a criterion only if it does so on *all* possible profiles. While the stringency of this requirement gives strong guarantees when they exist, it more often leads to prohibitive impossibilities; for example, Arrow's Impossibility Theorem [1] and the Gibbard-Satterthwaite Theorem [16] — some of the most foundational results in social choice — show that *no* voting rule can satisfy certain combinations of fundamental axioms. Such results imply that essentially all voting rules violate at least one standard axiom, and that some natural axioms are violated by almost any reasonable voting rule.[1] In other words, due to these widespread impossibilities, axioms — at least when studied from a worst-case perspective — are not always useful for determining which voting rule to use.

Fortunately, there is hope in that worst-case analysis may, in some cases, be unrealistically pessimistic: while it may be *possible* to find a profile in which a voting rule violates a given criterion, such a profile may be extremely unlikely to occur in practice. To move beyond the worst case in pursuit of more practical, positive results, it is standard practice to add randomness to the input. When doing so, we face a trade-off: while increased randomness may permit more positive guarantees, these guarantees grow weaker as the model's assumptions about the distribution of inputs grow stronger. In some sense, then, we want to know, roughly, the "minimal" amount of randomness required to ensure that a voting rule satisfies a given criterion with sufficiently high probability.

One approach to modeling inputs that is conducive to using minimal randomness is *smoothed analysis*, introduced by Spielman and Teng in their analysis of the Simplex algorithm's runtime [28]. The smoothed model assumes only that all inputs will be perturbed by a small amount of mild, generically-distributed random noise—in other words, that the world is slightly noisy.[2] Beyond this assumption, the smoothed model is worst-case because for a criterion to be satisfied within this model, it must hold after perturbation with sufficiently large probability *on all possible inputs*. When this is true, we say the criterion is *smoothed-satisfied*. For a primer on smoothed analysis, see [27].

To apply smoothed analysis in the voting setting, one needs a model that adds random noise to profiles. Lirong Xia recently proposed such a model [29], and then applied it to study the smoothed satisfaction of various social choice criteria [30, 31, 32, 33, 35] (with Weiqiang Zheng on [35]). Xia's model is highly general; it is, in fact, *more* general than the traditional smoothed analysis framework. Xia's results are technically sophisticated, often using intricate methods to get precise asymptotic bounds.

In this paper, we will present a different smoothed model for social choice, prove its convergence to well-behaved distributions, and then use these convergence bounds to study the smoothed-satisfaction of several axiomatic criteria. In the first two of these steps, we do not aim to significantly technically innovate

[1] https://en.wikipedia.org/wiki/Comparison_of_electoral_systems

[2] In the voting setting, assuming that inputs (profiles) are slightly noise is reasonable: people's preferences have been shown to be susceptible to small perturbations in daily life, and even in the preference solicitation process [10, 18, 19].

on Xia's work — in fact, as we discuss in the Related Work (Section 1.1), our model is a restriction of Xia's model, and some of our results can also be proven via Xia's. Rather, our goal is to propose an approach to smoothed analysis of social choice that strikes a different balance between interpretability and technical power: while Xia's model significantly generalizes the smoothed analysis framework, we propose a model that closely matches the smoothed framework; while Xia's results seek highly precise asymptotic bounds, we seek bounds that are close to asymptotically comparable but can be proven more simply. In the spirit of interpretability, we also present convergence rates in non-asymptotic form, in terms of intuitive constants.

As we will discuss in Section 1.1, we take this simplicity-driven approach because, in our view, the generality it sacrifices is modest, while the gains are fruitful: it allows us to quickly analyze most of the axioms and paradoxes studied across Xia's line of work, plus some yet unstudied in the smoothed setting; it yields high-level insights about where smoothed analysis is likely to circumvent impossibilities; and it illuminates promising new research directions. Our contributions below discuss these points in more detail.

Contributions:

(1) **Present an alternative model of smoothed social choice** that closely matches the core features of smoothed analysis, and establish its convergence to good behavior as the number of voters grows larger (Section 2);

(2) **Provide simple sufficient conditions** for smoothed-satisfaction and smoothed-violation of general criterion, with non-asymptotic convergence bounds (Sections 3 and 4);

(3) **Apply these conditions** to analyze the smoothed-satisfaction of several axiomatic criteria by large classes of voting rules (Sections 3 to 5).

In particular, we find that for a large classes of voting rules, the smoothed model is *not* sufficient to circumvent violations of many axioms, including Condorcet Consistency, Majority, Consistency, and Independence of Irrelevant Alternatives, but it *is* able to circumvent impossibilities in group notions of Strategyproofness, Participation, and Monotonicity. Interestingly, we find that the impossibility specified by Arrow's theorem is resolved with high probability in our smoothed model, but not necessarily so under a slightly stronger notion of Non-Dictatorship.

From these findings, we also distill high-level intuition about *types* of voting rules and axioms for which the smoothed model will (and will not) be sufficient to circumvent worst-case impossibilities.

(4) **Identify substantive new directions** in smoothed analysis of social choice, including an approach to relaxing one of the two major assumptions made in both ours and Xia's work (Section 6).

1.1 Related Work

We now give a detailed comparison of our model, techniques, and results to Xia's, and where relevant, to Spielman and Teng's original smoothed model [28]. To our knowledge, Xia's is the only work so far on smoothed analysis in social choice; however, the broader approach of beyond-worst-case analysis in social choice is far from new. There is such research on axiom satisfaction in average-case models like the Impartial Culture (i.i.d. uniform) model [15, 17, 21], and empirical work examining the frequency of criteria violations on simulated inputs and real elections [24, 14]. Further afield, semi-random models have also been used in other areas of algorithmic economics to escape worst-case constructions in fair division and mechanism design [26, 5, 2], and to circumvent computational hardness, e.g., [6, 7] study the smooth complexity of computing a Nash equilibrium.

1.1.1 Comparison of Models

The smoothed analysis framework was originally proposed by Spielman and Teng to provide theoretical justification for the Simplex algorithm's fast runtime in real-world instances, despite its exponential worst-case complexity [28]. In their analysis, they go beyond the worst-case by adding Gaussian noise independently to each entry of the real-valued constraint matrix that is the input to Simplex algorithm. Then, they bound the expected run time of Simplex in the worst-case over inputs, where these guarantees are parameterized by the variance of the Gaussian noise added. Stated as a more general framework, the idea of smoothed analysis is to fix an arbitrary instance, add noise from a parameterized distribution, and then measure the quality of the expected outcome on the worst input (or, whether a property is satisfied with high probability—an alternative formulation proposed by Spielman and Teng that is closer to ours).

This is precisely how our model works, but rather than its input being a constraint matrix of real numbers, it is a *base profile* of complete rankings over m alternatives. This base profile is perturbed by applying generically-structured noise independently across each of its rankings (we refer to this as the *independence* assumption). When evaluating the probability of a criterion being satisfied post-perturbation, we assume that the base profile is chosen adversarially, i.e., to minimize this probability. The noise distribution we apply to each ranking is parameterized by a value $\phi \in [0, 1]$ to measure the quantity of noise added, analogous to Spielman and Teng's variance parameter. Our main departure from the original smoothed model is, where Spielman and Teng specify Gaussian noise, we do not assume a specific noise distribution, instead allowing *any* ϕ-parameterized distribution that is *neutral* over alternatives.[3] *Neutrality* means that for a given ϕ, if we permute a ranking and add noise, this is equivalent to adding noise and then permuting the output. In that sense, our noise model over profiles can be specified by a single distribution over one ranking, permuted to be applied to ranking permutations. On this distribution over rankings, we also assume *positivity*—that when $\phi > 0$, this distribution assigns positive probability to all rankings. Under these assumptions, our class of noise models generalizes the popular Mallows noise model (e.g., [20]) as well as one-dimensional parameterizations of the Plackett-Luce model (e.g., see [8]).[4]

Like ours, Xia's model [29] assumes independence and positivity (as above), but it otherwise generalizes the smoothed framework much further than we do. In Xia's model, instead of voters having rankings, voters have *types*. Then, whereas in our model, a ranking is perturbed by applying a fixed noise distribution, in Xia's, each type is associated with an arbitrary distribution from which a "noisy" ranking is drawn. The worst-case is then taken over type assignments, rather than base profiles. This generalization departs from our model by not enforcing neutrality and allowing multiple possible noise distributions simultaneously. As such, Xia's model subsumes the concept of the smoothed model, and correspondingly, ours (for fixed ϕ): one could implement our model in Xia's by choosing Xia's types as all possible rankings, where each type's associated distribution is noise added to its ranking by our model.

While Xia's model allows more flexible distributions over voters, this added generality comes at a cost. Because the worst-case is no longer taken over base profiles — the inputs to the overall problem — we lose the intuitive interpretation of "the world adds noise to arbitrary problem inputs", and the connection to smoothed analysis is obscured. The generality in the noise permitted by Xia's model also makes it less clear how to implement a noise measurement parameter ϕ — a feature of our model which, as we illustrate in Section 6, lends itself to greater intuition and new, potentially fruitful research directions.

We are further motivated to explore a simplified model because in our view, it is not clear that this

[3]We do not commit to a noise distribution because there is no single well-established distribution that is obvious to apply (unlike in the real-valued setting, where the Gaussian is standard).

[4]Traditionally, the Plackett-Luce model takes a single m real-valued parameters, one per ranking position. Our model generalizes a variation where each parameter is expressed as a function of ϕ.

generality is worth the cost. First, the generality gained by permitting various non-neutral noise distributions at once is modest under conditions where the quantity of noise is very minimal — a core assumption of smoothed analysis. Especially in light of this, our neutrality assumption is weak in comparison to the much stronger assumption, made by both Xia's and our models, that noise is added independently across agents. As we will discuss in the next section, this much more central assumption forms the foundation of both our and Xia's analyses.

1.1.2 Comparison of Techniques and Results

At a high level, Xia's work and ours take a similar technical approach: both show that their noise models become well-behaved as the number of voters n grows large, and then use these convergence results to upper and lower bound how much probability mass is placed on "bad" profiles. Because of this correspondence, many of our results are qualitatively quite similar. They diverge in some special cases where Xia uses more intricate techniques to achieve stronger asymptotic convergence rates, while on the other hand, we aim to take the most straightforward approach using more standard techniques, which we hope offers more intuitive proofs. We now briefly highlight the similarities and differences in our respective analyses.

In both models, the key assumption is independence. This means that as n grows large, the resulting distribution over profiles will converge to distributions we understand, analogous to how sums of i.i.d. random variables converge to Gaussians via the Central Limit Theorem, or concentrate around their mean via Hoeffding's inequality. We prove the convergence of our model in Lemmas 2 and 3, corresponding to Xia's Lemma 1 of [29].[5] The convergence rates across the respective lemmas match asymptotically except when roughly, the set of "bad" profiles for a criterion is extremely small, which tends to only come into effect for very specific kinds of criterion, none of which we formally analyze in this paper. The difference stems from our proof relying on a multi-dimensional version of the well-known Berry-Esseen theorem, while Xia's relies on faster convergence of Poisson Multinomial variables.

In both ours and Xia's work, these convergence rates translate to sufficient conditions to prove either smoothed satisfaction or smoothed violation of various criteria in their respective models. In the latter part of Section 3, we show how our sufficient conditions can be applied to analyze many axioms for many voting rules. These results, plus several corollaries of these results (Section 5), echo the main conclusions of Xia's papers so far, up to a few differences in the set of voting rules we choose to study (our results can easily extend to additional rules).[6] Our results also include the first smoothed analysis of Arrows theorem, plus key axioms thus far unstudied in the smoothed setting including INDEPENDENCE OF IRRELEVANT ALTERNATIVES and CONSISTENCY. We prioritize proving these results via simple arguments with consistent structure — an effort that leads to higher-level insights across axioms, which we discuss in Section 6.

In Section 4, we analyze an axiom we call $o(\sqrt{n})$-GROUP-STABILITY, which roughly states that a coalition of size $o(\sqrt{n})$ cannot influence the outcome of an election. It implies group-based notions of the standard axioms STRATEGY-PROOFNESS, PARTICIPATION, and MONOTONICITY. The key difference here is that $o(\sqrt{n})$-GROUP-STABILITY does not satisfy the sufficient conditions for convergence either by our earlier lemmas, or Xia's aforementioned ones, so new techniques are required to analyze it. We find that

[5] Notice that when comparing formal statements, both models consider the space of histograms (vectors representing the fraction of the profile composed of each ranking) rather than profiles directly, and, while we tend to analyze the set of histograms corresponding to "bad" profiles (e.g., those in which an axiom is not satisfied) directly, Xia considers sets which are solutions to a system of linear equations and inequalities. However, in practice, these approaches end up being quite similar.

[6] These include Condorcet's voting paradox and the ANR impossibility theorem [29]; satisfaction of RESOLVABILITY (up to k-ties) [30]. Although our results don't explicitly discuss k-ties, we show the smoothed-satisfaction of RESOLVABILITY which requires there exist no $\geq$ 2-way ties); Moulin's impossibility theorem on PARTICIPATION and CONDORCET [31] [23]; and the individual satisfaction of the axioms PARTICIPATION and CONDORCET by various voting rules [32]

despite being violated in the worst-case by essentially all voting rules,[7] this axiom is smoothed-satisfied by a broad class of voting rules. This result generalizes a similar result obtained in a broad class of average-case models by Procaccia *et al.* [22]. It is also implied by recent work by Lirong Xia [33], though our results were obtained independently; among other results, Xia derives more precise convergence rates via more intricate techniques than we use here, but arrives at a qualitatively similar outcome.

2 Model

2.1 Preliminaries

As in the standard social choice setting, we have a set of m candidates M and a set of n agents $N = [n]$. Throughout the paper, we treat m as a fixed constant such that $m \geq 3$.

Rankings $\pi \in \mathcal{L}$. The agents express their preferences over the candidates as complete rankings. Formally, a ranking π is a bijection $[m] \to M$, mapping indices to candidates, where $\pi(j)$ represents the candidate in the j'th position of the ranking. We also use the standard notation $a \succ_\pi b$ to denote that candidate a is preferred to b in ranking π, or formally, $\pi^{-1}(a) < \pi^{-1}(b)$. We let $\mathcal{L}(M)$ be the set of $m!$ possible rankings over the candidates, written as $\mathcal{L}$ when M is clear from context. We fix an arbitrary order over rankings in $\mathcal{L}$ so we can talk about the j-th ranking in $\mathcal{L}$, or in some cases, the π-th ranking. With this indexing fixed, we refer to the last ($m!$-th) ranking in $\mathcal{L}$ as π_{-1}, and the set of rankings without this element as $\mathcal{L}_{-1} = \mathcal{L} \setminus \{\pi_{-1}\}$. For convenience, we will often work with $\mathcal{L}_{-1}$ instead of $\mathcal{L}$. We will occasionally use the *Kendall-Tau* distance between two rankings π, π', defined as the total number of swaps required to transform π into π'.

Profiles $\boldsymbol{\pi} \in \Pi$. We express a profile as a vector $\boldsymbol{\pi}$ of n rankings, $\boldsymbol{\pi} = (\pi_i | i \in [n])$, where π_i is agent i's ranking. We define addition over profiles in the natural way, so that the profile $(\boldsymbol{\pi} + \boldsymbol{\pi'}) = (\pi_i | i \in [n]) \| (\pi'_i | i \in [n'])$.[8] We extend this operation to permit positive integer multiples of profiles, such that adding a profile together z times is expressed as $z\boldsymbol{\pi}$. We let $\Pi_n = \prod_{i=1}^{n} \mathcal{L}$ be the set of all profiles on n voters and $\Pi = \bigcup_{n \in \mathbb{Z}^+} \Pi_n$ be the set of all profiles with any number of voters. We use $|\boldsymbol{\pi}|$ to denote the number of voters in profile $\boldsymbol{\pi}$. We say that a *pairwise-dominates* b in $\boldsymbol{\pi}$ when a is ranked ahead of b by over half of agents, i.e., when $\left|(\pi_i | a \succ_{\pi_i} b)\right| > n/2$. We say that a and b *pairwise tie* when $\left|(\pi_i | a \succ_{\pi_i} b)\right| = n/2$.

Histograms $\mathbf{h}$ and the histogram operator $\mathcal{H}(\cdot)$. Instead of working directly with profiles, we will work primarily with *histograms*. A *histogram* $\mathbf{h}$ is an $|\mathcal{L}_{-1}|$-length vector with nonnegative entries summing to at most 1. The π-th entry h_π is the proportion of agents with ranking π.[9] We further define the histogram operator $\mathcal{H}(\cdot)$, which takes as input a profile $\boldsymbol{\pi}$, and returns the associated histogram, that is,

$$\mathcal{H}(\boldsymbol{\pi})_\pi = {}^1\!/_n \, |\{i \,:\, \pi_i = \pi\}|.$$

Since we will use this operation so much, we also use the notation $\mathbf{h}^{\boldsymbol{\pi}} := \mathcal{H}(\boldsymbol{\pi})$. Note that a given histogram $\mathbf{h}^{\boldsymbol{\pi}}$ corresponds to an infinite number of profiles of different sizes and permutations of agents.

[7] The difficulty of satisfying strategy-proofness follows from the Gibbard-Satterthwaite theorem, which says that any onto voting rule choosing a single winner that is non-dictatorial and permits $m > 2$, must not be strategyproof [16].

[8] The indexing of agents in the resulting profile is not relevant to our results, but for formality, we can assume that the agents in $(\boldsymbol{\pi} + \boldsymbol{\pi'})$ will be indexed such that agent i in $\boldsymbol{\pi}$ is at index i in $(\boldsymbol{\pi} + \boldsymbol{\pi'})$, and agent i in $\boldsymbol{\pi'}$ is at index $i + n$ in $(\boldsymbol{\pi} + \boldsymbol{\pi'})$.

[9] We define histogram vectors a indexed by $\mathcal{L}_{-1}$ because the $m!$-th index is redundant, as the histogram over $\mathcal{L}$ must add to 1.

Because the notion of a profile histogram is so natural, we slightly abuse notation and allow the operator $\mathcal{H}(\cdot)$ to translate many kinds of profile-based objects into histogram-based objects. For instance, we will apply this operator to individual rankings $\pi \in \mathcal{L}_{-1}$, where $\mathcal{H}(\pi)$, which is a $|\mathcal{L}_{-1}|$-length basis vector with a 1 at the π-th index (and $\mathcal{H}(\pi_{-1})$ is the 0s vector). Accordingly, $\mathcal{H}(\Pi)$ is the set of all possible profile histograms; $\mathcal{H}(\mathcal{L})$ is the set of all rankings in their basis vector representations. We will also apply this operator to distributions over profiles and rankings in the natural way, first drawing a profile or ranking from the distribution, and then considering the corresponding histogram.

The simplex of profile histograms H. We will study the space of profiles as histograms in $|\mathcal{L}_{-1}|$-dimensional space. In this space, we define the simplex of all possible histograms H as

$$H := \left\{ \mathbf{h} \in [0,1]^{|\mathcal{L}_{-1}|} \ : \ \sum_{\pi \in \mathcal{L}_{-1}} h_\pi \leq 1 \right\}.$$

Note that H includes vectors with irrational entries that could never correspond to a well-defined profile. Nonetheless, it will be useful to consider the completed space. In order to talk about only the histograms that are realizable from well-defined profiles $\mathcal{H}(\Pi)$, we also define $H^{\mathbb{Q}} = H \cap \mathbb{Q}^{|\mathcal{L}_{-1}|}$ to be the subset of H of vectors with rational components, noting that $\mathcal{H}(\Pi) = H^{\mathbb{Q}}$.

Criteria C. We will establish conditions for the smoothed-satisfaction of general *criteria*, where a criterion $C : \Pi \mapsto \{\mathsf{True}, \mathsf{False}\}$ is a predicate that, conceptually, specifies whether a profile has a certain property. For example, a criterion could represent the statement "$\boldsymbol{\pi}$ contains 1 voter", so that $C(\boldsymbol{\pi}) = \mathsf{True}$ when $\boldsymbol{\pi} \in \Pi_1$ and False otherwise. We say C is *satisfied* if $C(\boldsymbol{\pi}) = \mathsf{True}$ for all $\boldsymbol{\pi} \in \Pi$. We say that C is *violated* if there exists a *counterexample* to C, i.e., a $\boldsymbol{\pi}' \in \Pi$ such that $C(\boldsymbol{\pi}') = \mathsf{False}$. Π^C and $\Pi^{\neg C}$ are the sets of profiles in which C is True and False, respectively.

Voting Rules R. A *voting rule* is a function $R : \Pi \mapsto 2^M$ mapping a given profile to a set of winning candidates. Then, $R(\boldsymbol{\pi})$ is the set of winners chosen by the voting rule on $\boldsymbol{\pi}$. If a voting rule by its standard definition results in a tie, rather than specifying a tie-breaking rule, we assume it returns all such winners; this assumption, however, is for ease of exposition only, and is not necessary for any of our results. Let $\mathcal{R}$ be the set of all voting rules. We will study several specific voting rules, defined colloquially below and formally in Appendix A.1.

- Positional Scoring Rules (PSRs) are represented by m-length vectors of weakly decreasing scores $s_1 \geq s_2 \geq \cdots \geq s_m$, where without loss of generality, $s_1 = 1$ and $s_m = 0$. Candidate a receives s_i points for each voter that ranks it i-th, and the candidate with the most points wins. We will specifically mention three standard PSRs, defined by their score vectors: Plurality: $(1, 0, \ldots, 0, 0)$, Borda: $(1, 1 - \frac{1}{m-1}, \ldots, \frac{1}{m-1}, 0)$, and Veto: $(1, 1, \ldots, 1, 0)$.
- Minimax selects the candidate whose maximum pairwise domination by any other candidate is the smallest.
- Kemeny-Young selects the candidate ranked first in the ranking with the minimal sum of Kendall-Tau distances from voters' rankings.
- Copeland selects the candidate with the highest Copeland score, computed as follows: candidate a receives 1 point for each candidate b it pairwise dominates, and 1/2 point for each candidate b with which it pairwise ties.

Axioms A. Let an axiom be a function $A : \mathcal{R} \to (\Pi \to \{\mathsf{True}, \mathsf{False}\})$, i.e., a mapping from a voting rule $R \in \mathcal{R}$ to a mapping describing whether A is satisfied by R on that profile. We can then think of $A(R)$ as a criterion representing the statement "R is consistent with A on $\boldsymbol{\pi}$". We will mainly discuss five standard axioms in this paper, described below and defined formally in Appendix A.2.

- R satisfies Resolvability on $\boldsymbol{\pi}$ if it selects a single winner.
- R satisfies Condorcet Consistency (abbreviated as Condorcet) on $\boldsymbol{\pi}$ if it selects the *Condorcet winner* (the candidate that pairwise dominates all other candidates), or by default if $\boldsymbol{\pi}$ has no Condorcet winner.
- R satisfies Majority on $\boldsymbol{\pi}$ if it selects the *majority winner* (the candidate ranked first by a majority of voters), or by default if $\boldsymbol{\pi}$ does not have a majority winner.
- R satisfies Consistency on $\boldsymbol{\pi}$ if there is no partition of $\boldsymbol{\pi}$ into subprofiles such that a unique candidate b, which is *not* the winner in $\boldsymbol{\pi}$, is chosen as the winner on all subprofiles in that partition.
- R satisfies Independence of Irrelevant Alternatives (IIA) on $\boldsymbol{\pi}$ if the winner, a, cannot be made to lose to b by adjusting votes in a way that does not change the relative positions of a and b.

We will also study group-level notions of the standard axioms Strategyproofness, Participation, and Monotonicity, and we also briefly consider the axioms Unanimity and Non-Dictatorship in our discussion of Arrow's theorem. We will describe all these axioms where we use them.

2.2 Noise distribution

We think of perturbing a ranking as selecting a random permutation of indices by which to permute it. Formally, for a permutation $\sigma : [m] \to [m]$, if $\sigma(i) = j$, we mean that the i'th ranked candidate post-noise will be the original j'th candidate pre-noise.[10] Then, we write $\pi + \sigma$ to denote the ranking achieved by permuting the candidates in π by permutation σ (so $+$ here represents composition).

A *noise distribution* $\mathcal{S}_\phi$, parameterized by a value $\phi \in [0, 1]$, is simply a distribution over permutations—the distribution from which we will draw a random permutation by which to permute π. Hence, for a ranking π, $\pi + \mathcal{S}_\phi$ will be the distribution over rankings resulting from applying $\mathcal{S}_\phi$ to π. Since we will so often work with this distribution $\pi + \mathcal{S}_\phi$, to streamline this notation, we use $\mathcal{S}_\phi(\pi)$ to denote this addition, which should be read as "the noise distribution $\mathcal{S}_\phi$ applied to ranking π". Since profiles are the true inputs to the social choice setting, we will slightly abuse notation to extend this perturbation model to entire profiles: for a profile $\boldsymbol{\pi} \in \Pi$, let $\mathcal{S}_\phi(\boldsymbol{\pi})$ denote the distribution over profiles achieved by adding $\mathcal{S}_\phi$ in an i.i.d. fashion to each voter's ranking—that is, $\mathcal{S}_\phi(\boldsymbol{\pi}) = \prod_{i=1}^n \mathcal{S}_\phi(\pi_i)$, where each $\mathcal{S}_\phi(\pi_i)$ is independent. Note that we are treating $\mathcal{S}_\phi$, $\mathcal{S}_\phi(\pi)$, and $\mathcal{S}_\phi(\boldsymbol{\pi})$ as distributions and random variables interchangeably.

A noise distribution $\mathcal{S}_\phi$ is parameterized by a dispersion parameter $\phi \in [0, 1]$. This parameter is analogous to the variance of a Gaussian distribution, as applied to perturb real values in past prominent applications of smoothed analysis (e.g., Spielman and Teng [28]). Extending the analogy between our work and theirs, our $\mathcal{S}_\phi$ corresponds to a 0-mean Gaussian distribution, and $\pi + \mathcal{S}_\phi$ as analogous to adding Gaussian noise to a real-valued input corresponding to π.

[10] We note that here, what we are using as σ is technically σ^{-1} per the formal definition of σ. We do this to make the notation clearer. This will not affect any of the later exposition, as the only specific permutation we reason about is the identity, whose inverse is itself.

Our class of noise models $\boldsymbol{\mathcal{S}}$. $\boldsymbol{\mathcal{S}}$ is the class of all noise models we consider. In particular, we permit a noise model within this class $\mathcal{S}_\phi \in \boldsymbol{\mathcal{S}}$ to be *any distribution over permutations*, so long as all $\mathcal{S}_\phi$ satisfy the minimal assumptions and regularity conditions below. These assumptions ensure that ϕ reasonably measures the amount of noise, and that the noise distribution is practical to work with. Before beginning, we define $\textsc{min-prob}(\mathcal{S}_\phi)$ to be the smallest probability $\mathcal{S}_\phi$ assigns to any permutation. That is,

$$\textsc{min-prob}(\mathcal{S}_\phi) := \min_\sigma \Pr[\mathcal{S}_\phi = \sigma]$$

This term will appear in many of our bounds, as well as multiple of our assumptions.

Assumption 1 (Extremal values). The distribution $\mathcal{S}_0$ is the point mass on the identity (so that $\phi = 0$ corresponds to no noise added). The distribution $\mathcal{S}_1$ is uniform over all permutations (so that $\phi = 1$ corresponds maximum noise). Note that this implies that for all profiles $\boldsymbol{\pi} \in \Pi$, $\mathcal{S}_1(\boldsymbol{\pi})$ is equivalent to the impartial culture model (see, e.g., [11]).

Assumption 2 (Positivity). For all $\phi \in (0, 1]$, $\textsc{min-prob}(\mathcal{S}_\phi) > 0$. That is, for any nonzero amount of noise, the resulting distribution assigns positive probability to all permutations.

Assumption 3 (Weak Monotonicity). The value $\textsc{min-prob}(S_\phi)$ is non-decreasing in ϕ.[11]

Assumption 4 (Continuity). For each permutation σ, the probability $\mathcal{S}_\phi$ places on σ is continuous in ϕ.

Proportions version of a noise distribution. We will primarily work in the space of histograms, so we will usually reason about the projection of distributions over rankings and profiles into histograms space. We express the histograms version of $\mathcal{S}_\phi(\pi)$, a distribution over rankings, as $\mathcal{H}(\mathcal{S}_\phi(\pi))$, a distribution over basis vectors. Then, the noise distribution over profile histograms is defined in terms of that over ranking histograms as

$$\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi})) = {}^1\!/\!{}_n \sum_{i=1}^{n} \mathcal{H}(\mathcal{S}_\phi(\pi_i)). \tag{1}$$

2.3 Convergence properties of noise distribution

Here, we will prove that despite the fact that the distribution over *rankings* induced by our noise model is essentially unstructured, the distribution over *proportions*, $\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi}))$, has useful convergence properties as n grows large. Intuitively this holds because, per Equation (1), $\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi}))$ is the sum (scaled by $1/n$) of independent distributions over rankings, i.e., the sum of (scaled) independent random basis vectors.

We first prove the following lemma (proof in Appendix B.1). One takeaway from this lemma it that $\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi}))$'s covariance matrix is quite easy to work with: its inverse not only exists but has a simple closed-form, and its eigenvalues are lower-bounded by a constant.

Lemma 1. For all noise models $\mathcal{S}$, parameters $\phi \in (0, 1]$, and rankings $\boldsymbol{\pi} \in \Pi_n$, the covariance matrix $\mathrm{Cov}[\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi}))]$ is invertible and has all positive real eigenvalues lower bounded by $\textsc{min-prob}(S_\phi)/(m!n)$.

[11] We note that our results don't centrally depend on Assumption 3 — Assumptions 2 and 4 are sufficient to give the same high-level results. We include Assumption 3 because it is not prohibitive, it simplifies the exposition, and allows us to give more usefully parameterized bounds.

Using Lemma 1, we show that $\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi}))$ uniformly (over profiles) converges at a rate of $O(1/\sqrt{n})$ to a multi-dimensional Gaussian distribution (Lemma 2). The uniform nature of this convergence means that for all profiles $\boldsymbol{\pi}$ of any fixed n, $\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi}))$ is at most $O(1/\sqrt{n})$ "distance away" from the the Gaussian distribution with expectation and variance corresponding to that of $\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi}))$. As written formally below, we measure the distance between our distribution and the Gaussian by comparing the probability mass the respective distributions place on arbitrary convex sets. The proof of this lemma relies on a relatively general form of the Berry Esseen bound, as stated in [3]. We restate the bound we apply in the proof of the lemma, located in Appendix B.2.

Lemma 2. Let $\mathcal{S}$ be a noise models, $\phi \in [0,1]$ a parameter, and $\boldsymbol{\pi} \in \Pi_n$ a profile on n agents. Then, for all convex sets $X \subseteq \mathbb{R}^{m!-1}$,

$$\left| \Pr\left[\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi})) \in X\right] - \Pr\left[\mathcal{N}\left(\mathbb{E}[\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi}))],\, \mathrm{Cov}[\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi}))]\right) \in X\right] \right| \leq \frac{O((m!)^{7/4})}{\sqrt{n} \cdot \textsc{min-prob}(\mathcal{S}_\phi)^{3/2}}.$$

Next, we show that $\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi}))$ concentrates at an exponential rate around its expectation (Lemma 3). The proof of this lemma is a straightforward application of Hoeffding's inequality, and can be found in Appendix B.3.

Lemma 3. Let $\mathcal{S}$ be a noise model, $\phi \in [0,1]$ a parameter, and $\boldsymbol{\pi} \in \Pi_n$ a profile on n agents. Then,

$$\Pr\left[\left\|\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi})) - \mathbb{E}[\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi}))]\right\|_1 < \varepsilon\right] > 1 - 2m!\exp\left(-\frac{2\varepsilon^2 n}{m!}\right).$$

We also prove another concentration bound: for sufficiently small ϕ, we have that $\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi}))$ concentrates an an exponential rate around the starting histogram, $\mathbf{h}^{\boldsymbol{\pi}}$. The proof is found in Appendix B.4.

Lemma 4. Let $\mathcal{S}$ be a noise model. For all $\varepsilon > 0$, there exists a $\phi \in (0,1]$ such that for all $\phi' \in [0,\phi]$ and profiles $\boldsymbol{\pi} \in \Pi_n$ on n agents,

$$\Pr\left[\left\|\mathcal{H}(\mathcal{S}_{\phi'}(\boldsymbol{\pi})) - \mathbf{h}^{\boldsymbol{\pi}}\right\|_1 < \varepsilon\right] > 1 - \exp\left(\frac{\varepsilon^2 n}{2}\right).$$

2.4 *Smoothed satisfaction* and *smoothed violation* of C

We have already defined above what it means for a criterion C to be *satisfied* and *violated*, in the worst-case sense. Now, we define the smoothed analogs to these terms, *smoothed-satisfied* and *smoothed-violated*. Note that satisfied implies smoothed-satisfied, and smoothed-violated implies violated. Recall that $\Pi^{\neg C}$ refers to the set of profiles that are counterexamples to C.

Definition 5 (smoothed-satisfied). A criterion C is $\mathcal{S}$-*smoothed-satisfied* at a rate $f(n,\phi)$ such that for all $n \in \mathbb{Z}^+$ and $\phi \in (0,1]$,

$$\sup_{\phi' \in [\phi,1]} \sup_{\boldsymbol{\pi} \in \Pi_n} \Pr\left[\mathcal{S}_{\phi'}(\boldsymbol{\pi}) \in \Pi^{\neg C}\right] \leq f(n,\phi).$$

Definition 6 (smoothed-violated). A criterion C is $\mathcal{S}$-*smoothed-violated* at a rate of $f(n)$ if there exists $\phi \in (0,1]$ and a profile $\boldsymbol{\pi}$ of size n such that for all $z \in \mathbb{Z}^+$,

$$\inf_{\phi' \in [0,\phi]} \Pr\left[\mathcal{S}_{\phi'}(z\boldsymbol{\pi}) \in \Pi^{\neg C}\right] \geq 1 - f(zn).$$

In words, when C is smoothed-satisfied, the probability of C's satisfaction after applying our noise distribution converges to 1 as n grows large; when C is smoothed-violated, the probability of C's satisfied converges to 0. We emphasize that, per the above definition, convergence to smoothed-satisfaction occurs eventually for *all* $\phi \in (0, 1]$, although the rate depends on ϕ—specifically, on MIN-PROB_ϕ. This is a strong definition in the sense that, when something is smoothed-satisfied, even a trivial amount of noise ϕ (and any larger amount) is enough to eventually achieve satisfaction at a rate depending on ϕ. When we show smoothed-violation, in contrast, we are saying there exists a constant amount of noise ϕ such that this amount, or any less, will not be enough to ensure satisfaction of C as n grows large. Note that there is a gap between these two definitions, i.e., smoothed-violated is not the negation of smoothed-satisfied, and a claim could therefore be "in-between" these definitions, satisfying neither. This will not end up being the case for any of the voting rules or criteria we study.

3 Sufficient conditions for smoothed satisfaction and violation of C

This section will give sufficient conditions for the smoothed satisfaction and smoothed violation of a general criterion C. For both, we provide convergence rates with explicit constants. At a high level, our sufficient condition for smoothed-satisfaction (Theorem 7) is that $\mathcal{H}\left(\Pi^{\neg C}\right)$, the set of counterexample histograms, is coverable by a finite collection of measure-zero convex sets. This will ensure the set is sufficiently "small", and hence has little probability mass placed on it. Our sufficient condition for smoothed-violation (Theorem 8) is that there is a positive-radius ball contained in the set counterexamples. We essentially show that there are instances in which after adding noise, profiles will concentrate in this ball.

Theorem 7. Fix a noise model $\mathcal{S}$ and criterion C. If there exists some set X such that (1) $X = \bigcup_{j=1}^{\ell} X_j$ where each $X_j \subseteq H$ is convex, (2) $\mathcal{H}\left(\Pi^{\neg C}\right) \subseteq X$, and (3) X is measure zero (noting that X is necessarily measurable), then C is **smoothed-satisfied** at a rate of

$$f(n, \phi) = \frac{\ell \cdot O((m!)^{7/4})}{\sqrt{n} \cdot \text{MIN-PROB}(\mathcal{S}_\phi)^{3/2}}.$$

Proof. Fix $\mathcal{S}$, C, and X satisfying the preconditions of the theorem: $X = \bigcup_{j=1}^{\ell} X_j$ where each X_j is convex, $\mathcal{H}(\Pi^{\neg C}) \subseteq X$, and X is measure 0. Since X is measure 0, this implies each X_j has measure 0. Fix $\phi \in (0, 1]$, a profile $\boldsymbol{\pi} \in \Pi_n$, and $\phi' \in [\phi, 1]$.

Fix an arbitrary X_j. Note that since X_j has measure zero, the probability mass placed on X_j by a Gaussian with an invertible covariance matrix is 0. Hence, Lemma 2 immediately implies that

$$\Pr[\mathcal{H}(\mathcal{S}_{\phi'}(\boldsymbol{\pi})) \in X_j] - 0 \leq \frac{O((m!)^{7/4})}{\sqrt{n} \cdot \text{MIN-PROB}(\mathcal{S}_{\phi'})^{3/2}}.$$

Using the monotonicity of MIN-PROB (Assumption 3), we get that this is at most $\frac{O((m!)^{7/4})}{\sqrt{n} \cdot \text{MIN-PROB}(\mathcal{S}_\phi)^{3/2}}$ (with ϕ' replaced with ϕ). Union bounding over all ℓ sets X_j tells us the following, as needed:

$$\Pr[\mathcal{H}(\mathcal{S}_{\phi'}(\boldsymbol{\pi})) \in X] \leq \frac{\ell \cdot O((m!)^{7/4})}{\sqrt{n} \cdot \text{MIN-PROB}(\mathcal{S}_\phi)^{3/2}}. \quad \square$$

Theorem 8. Fix a noise model $\mathcal{S}$ and a criterion C. Suppose there exists a profile $\boldsymbol{\pi}$ and radius $r > 0$ such that for all profiles $\boldsymbol{\pi}' \in \Pi$ satisfying (1) $|\boldsymbol{\pi}'| = z|\boldsymbol{\pi}|$ for some $z \in \mathbb{Z}^+$ and (2) $\|\mathbf{h}^{\boldsymbol{\pi}'} - \mathbf{h}^{\boldsymbol{\pi}}\|_1 < r$, it is the case that $\boldsymbol{\pi}' \in \Pi^{\neg C}$. Then, C is **smoothed-violated** at a rate of

$$f(n) = \exp\left(\frac{-r^2 n}{2}\right).$$

Proof. Fix a noise model $\mathcal{S}$, criterion C. Fix a profile $\boldsymbol{\pi}$ and radius $r > 0$ satisfying the preconditions of the theorem. Choose $\varepsilon = r$ and let ϕ be the one from Lemma 4 corresponding to ε and $\mathcal{S}$. Fix an arbitrary $z \in \mathbb{Z}^+$ and $\phi' \in [0, \phi]$. Notice that $\mathbf{h}^{\boldsymbol{\pi}} = \mathbf{h}^{z\boldsymbol{\pi}}$. Hence, Lemma 4 guarantees that the post-noise histogram lies in an r-radius ball around the original histogram $\mathbf{h}^{\boldsymbol{\pi}}$ with high probability — that is, $\mathcal{H}(\mathcal{S}_{\phi'}(z\boldsymbol{\pi})) \in B_r^{L_1}(\mathbf{h}^{\boldsymbol{\pi}})$ with probability at least $1 - f(z|\boldsymbol{\pi}|) = 1 - \exp\left(\frac{-r^2 z|\boldsymbol{\pi}|}{2}\right)$. Note also that every profile in the support of $\mathcal{S}_{\phi'}(z\boldsymbol{\pi})$ is guaranteed to have $z|\boldsymbol{\pi}|$ voters. These two facts, taken with both preconditions of the theorem, imply that the post-noise profile histogram is a counterexample with high probability — that is, $\mathcal{S}_{\phi'}(z\boldsymbol{\pi}) \in \Pi^{\neg C}$ with probability at least $1 - f(z|\boldsymbol{\pi}|)$. It then follows, by the definition of smoothed violation (Definition 6), that C is smoothed-violated at a rate of $f(z|\boldsymbol{\pi}|)$, as needed. □

Theorem 8 is straightforward to apply to voting rules and axioms, as we will illustrate in Section 3.1, because it allows us to reason about the profiles in the immediate vicinity of an initial counterexample. However, taking this approach requires the amount of noise applied to be quite small to ensure that concentration of the noise occurs around a point in the vicinity of the original profile histogram. To establish smoothed violation under *larger* quantities of noise (i.e., all the way up to the impartial culture model at $\phi = 1$), one must consider the existence of a positive-measure region around the expected histogram post-noise, $\mathbb{E}\left(\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi}))\right)$, rather than around the original profile histogram $\mathbf{h}^{\boldsymbol{\pi}}$. We give the concentration bound necessary to apply this more general approach in Lemma 3, and we illustrate its use in the proof of Proposition 23 in Section 6.3.

3.1 Application of Theorems 7 and 8 to smoothed-analyze $C = A(R)$

Now, we will demonstrate how to apply the conditions given by Theorems 7 and 8 to smoothed analyze criteria of the form $A(R)$, for several voting rules R and axioms A. Colloquially, we are particularly interested in cases where R violates A but smoothed-satisfies it. We summarize our findings in Table 1.

Our results illustrate the broad applicability of the conditions given by Theorems 7 and 8: while $C = A(R)$ may not satisfy the preconditions of either theorem for *all possible* A and R, for all mainstream rules and many axioms we study, we find that one of the two theorems applies. This is because popular rules and axioms tend to be simple, translating to regularity in the regions of histogram space where they are satisfied or violated. Here, we will show that not only can we apply Theorems 7 and 8 to many voting rule-axiom pairs via simple arguments; we can often exploit structural commonalities across many rules or axioms, giving intuition about what "types" of rule-axiom pairs smoothed analysis will be fruitful.

Voting Rules. Here, we will study *hyperplane rules* (Definition 9), a class of voting rules which is known to be equivalent to generalized scoring rules [34] and encompasses essentially every popular voting rule considered in the social choice literature.

Definition 9 (Hyperplane rules [22]). Note that given a set of ℓ affine-hyperplanes, these hyperplanes partition the space of histograms into at most 3^ℓ regions, as every point is either on a hyperplane or on one of two sides. We say that R is a *hyperplane rule* if there exists a finite set of affine hyperplanes $H_1, \ldots, H_\ell$ such that R is constant on each such region.

In our analysis, we will sometimes subdivide this class of voting rules into *decisive* and *non-decisive* hyperplane rules. Decisive hyperplane rules are those which output a single winner on profiles that are not on any hyperplane; non-decisive hyperplane rules are all others. Sometimes, we will study specific rules within these classes: we study Plurality, all non-Plurality Positional Scoring Rules (PSRs) (which include the popular rules Borda Count and Veto), Minimax, Kemeny-Young (all decisive hyperplane rules) and Copeland (a non-decisive hyperplane rule). These rules are formally defined in Appendix A.1.

Axioms. We apply our conditions to study the smoothed-satisfaction of the axioms Resolvability, Condorcet, Majority, Consistency and Independent of Irrelevant Alternatives (IIA). See Appendix A for definitions of axioms.

Results. Table 1 summarizes the classifications we make for the voting rules and axioms listed above.

First, we analyze Resolvability, finding that whether a hyperplane rule smoothed-satisfies this axiom depends on whether it is decisive.

Proposition 10. All decisive hyperplane rules smoothed-satisfy Resolvability. All non-decisive hyperplane rules smoothed-violate Resolvability.

We prove smoothed-satisfaction by decisive hyperplane rules via the measure-zero nature of hyperplanes on which ties occur via Theorem 7. The proof of smoothed-violation by non-decisive hyperplane rules is by the existence of regions beyond just the hyperplanes where it is not resolvable. The formal proof is located in Appendix C.1.

The remaining results are stated in Proposition 11 and appear in the gray shaded region of the table. For every axiom A and rule R in that region of the table, we find that given essentially any counterexample to $A(R)$ (even small, canonical ones), all histograms in the immediately surrounding ball must also be counterexamples, immediately implying via Theorem 8 that $A(R)$ is smoothed-violated. The arguments that show this are simple and homogeneous across $A(R)$, and are provided in Appendix C.2.

Proposition 11. For all $R \in$ PSRs $\cup$ {Minimax, Kemeny-Young, Copeland} and $A \in$ {Condorcet, Majority, Consistency, IIA}, if R violates A, then R smoothed-violates A.

The common motifs we find across rules and axioms in these negative results suggest that Proposition 11 could be generalized to all rules R and axioms A in large classes. Such a characterization would illuminate swaths of the social choice space where smoothed analysis fundamentally cannot circumvent worst-case impossibilities. We discuss more concrete research questions in this direction in Section 6.

	Axioms				
Voting Rules	Condorcet	Majority	Consistency	IIA	Resolvability
Plurality	**s-violated**	satisfied	satisfied	**s-violated**	s-satisfied
(non-Plurality) PSRs	**s-violated**	**s-violated**	satisfied	**s-violated**	s-satisfied
Minimax	satisfied	satisfied	**s-violated**	**s-violated**	s-satisfied
Kemeny-Young	satisfied	satisfied	**s-violated**	**s-violated**	s-satisfied
Copeland	satisfied	satisfied	**s-violated**	**s-violated**	**s-violated**

Table 1: Classification of rule-axiom pairs as satisfied, smoothed(s)-satisfied, or smoothed(s)-violated. Gray shading denotes the set of claims for which we use common "robustness" properties across voting rules and axioms to apply Theorem 8.

4 Smoothed satisfaction of $\rho(n)$-*Group-Stability*

In this section, we introduce a new axiom, $\rho(n)$-Group-Stability, which colloquially requires that the outcome of voting rules be *stable* to a change in the behavior of up to $\rho(n)$ voters, where $\rho(n)$ is some function increasing in n.

Definition 12 ($\rho(n)$-Group-Stability). For a given rule R, $\rho(n)$-Group-Stability(R) is satisfied if, for every pair of profiles $\boldsymbol{\pi}, \boldsymbol{\pi}'$ that differ on at most $\rho(n)$ of agents, we have that $R(\boldsymbol{\pi}) = R(\boldsymbol{\pi}')$.

We will prove the smoothed-satisfaction of $\rho(n)$-Group-Stability(R) via the following technical theorem, which states a fundamental property of our noise model: on any $\boldsymbol{\pi}$, the histogram drawn from our noise model $\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi}))$ is unlikely to be within a certain distance $\delta(n)$ of any specific hyperplane as n grows large, where $\delta(n)$ is decreasing sufficiently quickly in n. We think of this region of interest as a "thick" hyperplane whose width is shrinking as n grows large.

Theorem 13. Let $\mathcal{G}$ be the set of all hyperplanes in $\mathbb{R}^{m!-1}$. For all noise models $\mathcal{S}$, parameters $\phi \in [0, 1]$, and $\delta(n) \in o(1/\sqrt{n})$, we have the following, where d is the L_1 distance.

$$\sup_{G\in\mathcal{G}} \sup_{\phi'\in[\phi,1]} \sup_{\boldsymbol{\pi}\in\Pi_n} \Pr\left[d(\mathcal{H}(\mathcal{S}_{\phi'}(\boldsymbol{\pi})), G) \le \delta(n)\right] \in o(1).$$

We prove this theorem and give a non-asymptotic version of the upper-bound in Appendix D.1. This is an anti-concentration result: we show that even if the expected histogram post-noise falls within the thick hyperplane (in which case the outcome *can* be influenced by $\rho(n)$ voters), the width of this thick hyperplane is shrinking faster than the distribution over histograms concentrates as n grows large. Intuitively, the condition that $\delta(n)$ need be in $o(1/\sqrt{n})$ arises from the fact that our bound on the probability of the post-perturbation profile being in the thick hyperplane depends on $\delta(n) \cdot \sqrt{n}$, and we want this bound to approach 0 as n gets large.

One might ask why this theorem does not follow from Theorem 7; the issue is finding a finite number of measure-zero sets that cover the entire space of possible counterexamples. When trying to identify a finite number of measure zero convex sets covering all profiles whose histograms are within $\delta(n)$ of a hyperplane G, the pigeonhole principle implies that for any finite n, at least one of these sets must cover points whose convex hull is positive measure, implying that this set must have positive measure.

4.1 Application of Theorem 13 to smoothed-analyze $C = A(R)$

Voting Rules. Here we will prove results about all hyperplane rules (Definition 9).

Axioms. Here we will analyze the axiom $\rho(n)$-Group-Stability (Definition 12). This axiom implies strong, $\rho(n)$-parameterized versions of three axioms of common interest, which we define formally in **??** and intuitively here:

- $\rho(n)$-Group-Strategyproofness: No group of up to $\rho(n)$ agents can strategically misreport their preferences in concert and cause R to output an alternative they all weakly prefer, and at least one of them strongly prefers.
- $\rho(n)$-Group-Participation: No group of up to $\rho(n)$ agents can leave the election and cause R to output an alternative they all weakly prefer, and at least one of them strongly prefers.

- $\rho(n)$-Group-Monotonicity: No group of up to $\rho(n)$ agents can weakly decrease the position of an alternative c, which is not currently a winner by R, in their rankings, and cause c to become a winner by R.

The fact that $\rho(n)$-Group-Stability implies the first and third of these is obvious; the second is technically not implied, but does hold for a slight generalization of $\rho(n)$-Group-Stability. This generalization does not affect any results asymptotically but makes a few bounds marginally messier, so for ease of exposition, we discuss the simpler version of the axiom here and discuss the more precise bounds in the proof. When analyzing these axioms, we will often set $\rho(n) \in o(\sqrt{n})$, as the convergence rate will depend on $\rho(n)/\sqrt{n}$. This bound arises from Theorem 13, which requires $\delta(n)$ (closely related to $\rho(n)$) to be $o(1/\sqrt{n})$.

Results. We conclude Proposition 14, simply by invoking Theorem 13 and union bounding over the finite number of hyperplanes referred to in Definition 9. This conclusion is striking as, by the Gibbard-Satterthwaite theorem [16], even a single-voter version of this axiom is worst-case violated by essentially all reasonable voting rules. Theorem 13.

Proposition 14. $\rho(n)$-group-stability(R) with $\rho(n) \in o(\sqrt{n})$ is smoothed-satisfied by all hyperplane rules.

Of course, this result directly implies that all hyperplane rules also either satisfy or smoothed-satisfy the axioms $o(\sqrt{n})$-Group-Strategyproofness, $o(\sqrt{n})$-Group-Participation, and $o(\sqrt{n})$-Group- Monotonicity. The group sizes in these axioms are stated asymptotically; to understand precise group sizes for which these axioms are smoothed-satisfied, see the non-asymptotic bounds in the proof of Theorem 13.

5 Extending our model to other axiom-based criteria

So far, we have applied our sufficient conditions to study the smoothed-satisfaction of criteria of the form $C = A(R)$ for many rules R and axioms A. Naturally, our approach can be used to apply our model to study additional criteria of this form. Now, we show how to also extend our model to analyze axiom-based criteria of other forms.

5.1 Condorcet's Voting Paradox

To begin with a simple example, consider *Condorcet's Voting Paradox*, which is said to occur in profiles that contain a *Condorcet cycle*, i.e., have no Condorcet winner. This paradox has already garnered interest from the smoothed analysis perspective [29]. Placing this within our framework, the criterion C represents the statement "There exist no Condorcet cycles in $\boldsymbol{\pi}$".

Proposition 15. Let C be the criterion that $\boldsymbol{\pi}$ contains no Condorcet Cycle. C is smoothed-violated.

This proposition is proven via Theorem 8 using much the same procedure as for the rules and axioms before: We find a profile where all nearby histograms have a Condorcet Cycle, implying the preconditions of Theorem 8.

5.2 Simultaneous satisfaction of multiple axioms

Some of the most influential worst-case impossibilities in social choice are of the form "there exists no voting rule that satisfies the intersection of all axioms in $A \in \mathcal{A}$", where $\mathcal{A}$ is some collection of axioms. Stated in terms of criteria, there exists no R that satisfies all criteria in the set $A(R)|A \in \mathcal{A}$. This type of axiomatic impossibility encompasses, e.g., Arrow's theorem [1], Gibbard-Satterthwaite [16], the ANR impossibility theorem (e.g., see [29]), and the impossibility of simultaneously satisfying Condorcet and Participation as identified by Moulin [23], referred to henceforth as the Condorcet-Participation impossibility. We address all but Arrow's theorem in this section, and then address Arrow's theorem in Section 6.1.

Positive direction (e.g. ANR, Gibbard-Satterthwaite, Condorcet-Participation). A positive smoothed result pertaining to such an impossibility would show that there exists a voting rule R such that, for all $A \in \mathcal{A}$, $A(R)$ is either satisfied or smoothed-satisfied. This type of result follows naturally from the approaches we used in previous sections. For instance, our results from Table 1 and Proposition 14 immediately give the following corollaries:

Corollary 16 (Smoothed resolution of ANR impossibility theorem). Per Table 1, there exist several voting rules that simultaneously satisfy Anonymity and Neutrality and smoothed-satisfy Resolvability.

Corollary 17 (Smoothed resolution of Gibbard-Satterthwaite). Per Proposition 14, there exist several voting rules that simultaneously are non-dictatorial, permit consideration of more than two alternatives, and smoothed-satisfy Strategyproofness ($o(\sqrt{n})$-Group-Strategyproofness $\implies$ 1-Strategyproofness, the relevant definition for the this impossibility).

Corollary 18 (Smoothed resolution of Condorcet-Participation impossibility). Per Proposition 14, there exist several voting rules that simultaneously satisfy Condorcet and smoothed-satisfy Participation ($o(\sqrt{n})$-Group-Participation $\implies$ 1-Participation, the relevant definition for this impossibility).

Negative direction. A negative result pertaining to such an impossibility would show that, for all possible voting rules R, there exists an $A \in \mathcal{A}$ such that $A(R)$ is smoothed-violated. Pursuing this type of result is harder, because doing so requires reasoning about *all possible* voting rules, including those with no geometric regularity. This is in contrast to the voting rules to which we can apply our conditions, upon which we impose geometric restrictions. The good news is, the restrictions we place on voting rules are very weak, and our model would highly amenable to showing impossibilities across all, e.g., hyperplane rules (Definition 9). Given the weakness of these restrictions, such a result would already be practically relevant. One opportunity to pursue such a result is offered by the open question related to Arrow's theorem we identify in Section 6.1.

6 Open Directions

This section is dedicated to outlining some higher-level insights that arise from this paper, and discussing how they lead to impactful open questions that future work could study within our model. We note that in the introduction, we also identified additional, non-axiom-based criteria that might also be worth studying in the smoothed model: the distortion, and the complexity of computing a voting rule's winners.[12] We do not discuss these additional potential directions here.

[12] We note that there already exists work by Lirong Xia and Weiqiang Zheng along these lines, examining the smoothed time-complexity of computing the winners of Kemeny and Slater voting rules [35].

6.1 Smoothed-Analysis of Arrow's Theorem

In light of our findings on the smoothed-violation of Independence of Irrelevant Alternatives (IIA), one might guess that in the smoothed model, Arrow's Impossibility still holds. However, due to how the axiom Non-dictatorship is defined by Arrow [1], this is actually not the case:

Observation 19. There exists a voting rule which simultaneously satisfies Non-Dictatorship and smoothed-satisfies IIA and Unanimity.

A voting rule is a dictatorship if there exists some voter such that the outcome of the voting rule is always that voter's first choice. Now define the voting rule, *almost-dictatorship*, that chooses voter 1's favorite alternative on every profile except one arbitrary profile. The existence of this single exceptional profile means that our rule satisfies Non-dictatorship. However, since the rule only differs from a dictatorship on this one profile, it can easily be checked that it smoothed-satisfies IIA and Unanimity.

Conceptually, this positive result is not very satisfying, as it does not get at the heart of what makes Arrow's axioms inconsistent. In some sense, this is because satisfying non-dictatorship is too easy; to say something more meaningful, we propose a strengthening of Non-Dictatorship, called Local Non-Dictatorship, that may be more interesting, particularly in the smoothed context.

Definition 20 (Local Non-Dictatorship). Fix a profile $\boldsymbol{\pi}$ and define voter i's neighborhood around $\boldsymbol{\pi}$, $N_i(\boldsymbol{\pi}) \subseteq \Pi$, to be the set of all profiles reachable by switching i's ranking for another ranking. Then, we say rule R satisfies Local Non-Dictatorship on profile $\boldsymbol{\pi}$, if for all voters i, there is a $\boldsymbol{\pi}' \in N_i(\boldsymbol{\pi})$ such that i's first choice in $\boldsymbol{\pi}'$ doesn't win. In a sense, then, satisfying this axiom disallows a voter from being a dictator in any *local* area of profile space.

Observe that the satisfaction of Local Non-Dictatorship implies the satisfaction of Non-Dictatorship, making the former stronger; thus, Arrow's impossibility also implies inconsistency between IIA, Unanimity, and Local Non-Dictatorship. As we hoped, this version of the axiom is sufficiently stronger such that our aforementioned rule, almost-dictatorship, no longer smoothed-satisfies it, since voter 1 is the dictator across almost every neighborhood. In fact, it is so strong that none of the rules we consider actually satisfy it: for example, consider Plurality when among all voters other than i, there is an m-way tie. In this case, whoever i ranks first wins, so they are a local dictator. Nonetheless, we still believe Local Non-Dictatorship is an interesting axiom, as all reasonable voting rules at least smoothed-satisfy it (due to pivotality being a result of multi-way ties).

Future work: smoothed analysis of a strengthening of Arrow's theorem. The aforementioned rule, almost-dictatorship, does not smoothed-satisfy Local Non-Dictatorship. Then, the question remains open: does there exist a voting rule that simultaneously smoothed-satisfies Local Non-Dictatorship, IIA, and Unanimity?

6.2 Defining where smoothed analysis cannot overcome worst-case impossibilities

In Section 3, we show that several canonical axioms — Condorcet, Majority, Consistency, and IIA — are smoothed violated by *every* voting rule we study that worst-case violates them. In other words, for all rules and axioms in Section 3 for which we prove smoothed-violation, it is implied almost directly by violation. Moreover, the proofs of these results across across rule-axiom pairs are essentially the same, all built from small, simple counterexamples using one common property of rules and axioms: their tendency to have similar behavior on similar profiles. This property is inextricably linked to voting rules and axioms being "simple" and "natural"—central properties of standard rules and axioms — suggesting that there may be large classes of attractive rules and axioms in which smoothed analysis is fundamentally insufficient to circumvent worst-case impossibilities.

Future work: a sufficient condition for smoothed impossibility. Motivated by the potential to formally characterize where smoothed analysis fundamentally cannot help us circumvent worst-case impossibilities, we propose the pursuit of a sufficient condition on axioms. Such a condition would imply that an axiom must be either satisfied or smoothed-violated by any rule in a large class. Here, we lay out some intermediate steps toward such a condition that build on our results so far:

1. *Generalize Proposition 11 across rules.* For each axiom $A \in \{\textsc{Condorcet}, \textsc{Majority}, \textsc{Consistency}, \text{IIA}\}$ individually, show that all voting rules in a natural class either satisfy or smoothed-violate A. Proving this for, e.g., a large sub-class of hyperplane rules, would be a slight generalization of our results so far — it would require showing only that, for every rule in the class considered, the existence of any counterexample implies the existence of a strict counterexample, by the relevant definition of "strict" given in Appendix C.2.

2. *Generalize across absolute axioms.* For a large class of voting rules, pursue a sufficient condition across exclusively *absolute axioms* – axioms whose satisfaction is determined on a single profile (e.g., Condorcet is an absolute axiom, because the Condorcet winner is determined on a single profile; Strategy-proofness is not an absolute axiom, because its satisfaction depends on the behavior of the voting rule across pairs of profiles). This is a natural intermediate step because absolute axioms are easier to characterize than axioms over multiple profiles, called *relative axioms*, for the reasons discussed next.

3. *Generalize across relative axioms.* For any of the classes of rules, above, pursue a sufficient condition across exclusively relative axioms. Such a general condition is perhaps the trickiest part because, as illustrated in our proofs for Consistency and IIA, the arguments differ more across relative axioms. This is because positive measure sets must exist around *multiple* of the profiles that go into constructing the counterexample in order to prove smoothed-violation, and the relationships between these profiles vary across relative axioms.

6.3 More meaningfully using the measurement of noise.

In our results so far, the ϕ-dependent feature of the distribution by which we have parameterized our results is $\textsc{min-prob}_\phi$. Given that $\textsc{min-prob}_\phi$ is such a simple static of the entire distribution, we are in some sense only weakly using ϕ as a measurement of the quantity of noise. This is in contrast to, for example, Spielman and Teng's results, which are parameterized in terms of σ, the variance of their noise distribution—a much richer statistic. A natural question, then, is how we can get bounds in terms of ϕ, where ϕ represents a richer measure of the noise distribution.

Our use of such a simple statistic is a direct result of the fact that we study such a general class of noise distributions; for these distributions, we do not fully specify relationship between the noise distribution and ϕ beyond weak regularity conditions. However, one *can* get bounds in terms of ϕ by defining a specific ϕ-parameterized noise model within our class. Then, our work can be applied to directly yield bounds in the positive direction, and open the door for interesting future work in the negative direction. We illustrate both through an example in which we choose our noise model to be the popular Mallows noise model, defined below.

Definition 21 (Mallows noise model [20]). Let $d : \mathcal{L} \times \mathcal{L} \to \mathbb{N}$ be the Kendall-Tau distance. Then, the Mallows model $\mathcal{S}_\phi^{Mallows}$, is defined as

$$\Pr\left[\mathcal{S}_\phi^{Mallows}(\pi) = \pi'\right] = \frac{1}{Z}\phi^{d(\pi,\pi')},$$

where $Z = \sum_{\pi' \in \mathcal{L}} \phi^{d(\pi',\pi)}$, a normalizing term.

In order to get specific bounds in terms of ϕ from our positive results (Theorems 7 and 13)—both which are parameterized by MIN-PROB_ϕ—one just needs to compute MIN-PROB_ϕ as a function of ϕ, and then plug in the result to Theorems 7 and 13. By definition of Mallows, $\text{MIN-PROB}(\mathcal{S}^{Mallows}) \geq \phi^{\binom{m}{2}}/m!$. Then, the bound we get from Theorem 7 is the following (one can get comparable rates for Theorem 13, similarly).

Corollary 22. Fix a rule R, an axiom A, and an $m \geq 3$. Then, the convergence rate given by Theorem 8 is

$$\Pr[\mathcal{H}(\mathcal{S}_{\phi'}^{Mallows}(\boldsymbol{\pi})) \in X] \leq \frac{\ell \cdot O(m!^{13/4})}{\sqrt{n} \cdot \phi^{3\binom{m}{2}/2}}.$$

Future work: convergence to smoothed-*violation* in terms of ϕ. In contrast to the positive direction above, ϕ does not appear in the convergence rate for smoothed violation. Instead, it appears implicitly among the quantifiers: Theorem 8 states that there is an $r > 0$ such that the claim holds; implicit in this r is a quantity of noise ϕ such that for any $\phi' \in [0, \phi]$, ϕ' is insufficient to escape convergence to smoothed violation. While we know such a ϕ exists, the *magnitude* of this ϕ is unknown — it depends on the specific noise model, axiom, and voting rule being analyzed. For specific claims, we might like to know how large this ϕ can be before our sufficient condition for smoothed violation given by Theorem 8 no longer holds—i.e., are certain claims that smoothed-violated even after the addition of a large amount of noise?

To illustrate how such a result might be pursued, we continue our Mallows noise model example and prove Proposition 23 via Theorem 8. This proposition shows that violations of the axiom Condorcet by any positional scoring rule is extraordinarily robust to Mallows noise: we find that there exists a profile to which we can add *arbitrarily close to $\phi = 1$ noise*, and all positional scoring rules will still smoothed-violate Condorcet (for the proof, see Appendix E).

Proposition 23. Let A = Condorcet. Then, there exists a $\boldsymbol{\pi}$ such that for all $R \in PSR$ and for all $\phi \in [0, 1)$,

$$\lim_{z \to \infty} \Pr\left[\mathcal{H}(\mathcal{S}_\phi^{Mallows}(z\boldsymbol{\pi})) \in \mathcal{H}(\Pi^{\neg(A(R))})\right] = 1$$

This type of analysis could be replicated in other noise models, with other voting rules, and with other criteria (even beyond axioms). It also seems promising to try to generalize such a finding across voting rules, axioms, and other criteria, to carve out spaces where impossibilities are maximally robust: e.g., where smoothed violation by some ϕ always implies smoothed violation for all $\phi \in [0, 1)$.

6.4 Going beyond the positivity assumption

A central assumption relied upon by both our work and Xia's is that any noise distribution must assign positive probability to *all rankings*. This assumption is fundamental to both our analyses because it ensures that our respective noise distributions converge to familiar, tractable distributions. This assumption, in our case, manifests in the fact that our upper-bounds are parameterized by MIN-PROB$_\phi$ in a way that makes them undefined if MIN-PROB$_\phi$ is zero, and conversely, stronger as MIN-PROB$_\phi$ increases. This means that in order to get good convergence rates, we need MIN-PROB$_\phi$ to be nontrivial. This is an unattractive scenario because it means that a voter will have some nontrivial probability of drastically changing — even *reversing* their ranking — as a result of our noise. In the voting setting, this is conceptually contrary to what we would think of as a "minimal" perturbation, as it produces a conceptually different ranking.

The good news is, it intuitively seems that MIN-PROB$_\phi$ being large is not a necessary condition for achieving smoothed-satisfaction of some axioms. For instance, suppose the noise distribution is uniform over all rankings except for one, for which it has zero probability. In this example, MIN-PROB$_\phi = 0$, and yet the likelihood of a well-behaved voting rule violating RESOLVABILITY after applying this noise distribution is intuitively extremely low, since all counterexamples lie on hyperplanes away from which this near-uniform distribution must be dispersed. Neither our work so far nor Xia's captures this intuition.

Future work: an intuitive noise model restriction that enables convergence without positivity. The intuition that positivity does not seem to be a necessary condition for smoothed-satisfaction of some axioms encourages us to pursue upper bounds parameterized by something other than MIN-PROB$_\phi$. Then, the question becomes: what parameter of the noise distribution should we use to capture its "noisiness"? If we want to specify a single noise distribution, we can measure its variance; however, if we want to study classes of distributions, we need a simpler, more universal parameter. The intuition about RESOLVABILITY offers an idea: for smoothed-satisfaction, we just need some sufficient dispersion of probability mass in space, so that a nontrivial amount exists outside the hyperplanes at which ties occur. A natural guess for a parameterization, then, might be MAXPROB$_\phi$, the *maximum* probability assigned to any alternative: then, as MAXPROB$_\phi$ decreases, more probability must be dispersed away from the original ranking.

Unfortunately, MAXPROB$_\phi$ being low is not sufficient to guarantee smoothed-satisfaction in general noise models, because the probability mass may be dispersed but oriented poorly, so that it falls entirely in the space of counterexamples (e.g., for RESOLVABILITY, this would be exactly the hyperplanes on which ties occur). The problem of the noise having an extremely specific, pathological orientation likely cannot be excluded for general noise models under many parameterizations, making this a more fundamental problem that must be circumvented in order to get results that do not depend on MIN-PROB$_\phi$.

Here, we offer a simple, natural modeling restriction that both seems like a promising potential solution to this problem, and also leads to noise models capturing the idea that rankings should undergo less *conceptual* change when noise is lower. Suppose we restricted our noise models to those in which the probability of drawing ranking π' from a noise distribution centered at base ranking π was parameterized in terms of the Kendall-Tau (KT) distance between π, π'. This would mean that every ranking with the same KT-distance from π would be chosen with the same probability, meaning that any noise distribution in this class has significant symmetry over ranking space — hopefully enough to avoid distributions that place all their mass on extremely contrived subspaces of profiles. This restriction is also nice because it more firmly anchors the noise model to be "centered" around the base ranking from the standpoint of KT-distance. Under simple types of monotonicity within this class of models, this centering corresponds to the intuition that "noisier" distributions should be more likely to result in more dramatic conceptual changes to rankings, i.e., result in rankings with greater KT-distance from the base ranking.

References

[1] Kenneth Arrow. Individual values and social choice. *Nueva York: Wiley*, 24, 1951.

[2] Yushi Bai, Uriel Feige, Paul Gölz, and Ariel D Procaccia. Fair allocations for smoothed utilities. 2022.

[3] Vidmantas Bentkus. A lyapunov-type bound in rd. *Theory of Probability & Its Applications*, 49(2):311–323, 2005.

[4] Andrew C. Berry. The accuracy of the gaussian approximation to the sum of independent variates. *Transactions of the american mathematical society*, 49(1):122–136, 1941.

[5] Avrim Blum and Paul Gölz. Incentive-compatible kidney exchange in a slightly semi-random model. In *ACM Conference on Economics and Computation*, volume 21, 2021.

[6] Shant Boodaghians, Joshua Brakensiek, Samuel B Hopkins, and Aviad Rubinstein. Smoothed complexity of 2-player nash equilibria. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 271–282. IEEE, 2020.

[7] Shant Boodaghians, Rucha Kulkarni, and Ruta Mehta. Smoothed efficient algorithms and reductions for network coordination games. In *11th Innovations in Theoretical Computer Science Conference, ITCS 2020*, page 73. Schloss Dagstuhl-Leibniz-Zentrum fur Informatik GmbH, Dagstuhl Publishing, 2020.

[8] Weiwei Cheng, Krzysztof Dembczynski, and Eyke Hüllermeier. Label ranking methods based on the plackett-luce model. In *Proceedings of the 27th International Conference on Machine Learning (ICML)*, pages 215–222, 2010.

[9] Vincent Conitzer and Tuomas Sandholm. Communication complexity of common voting rules. In *Proceedings of the 6th ACM Conference on Economics and Computation (EC)*, pages 78–87, 2005.

[10] Leah Melani Dillman, Don A .and Christian. Survey mode as a source of instability in responses across surveys. *Field methods*, 17(1):30–52, 2005.

[11] Ömer Eğecioğlu and Ayça E. Giritligil. The impartial, anonymous, and neutral culture model: a probability model for sampling public preference structures. *The Journal of mathematical sociology*, 37(4):203–222, 2013.

[12] Ulle Endriss and Ronald De Haan. Complexity of the winner determination problem in judgment aggregation: Kemeny, slater, tideman, young. In *Proceedings of the 14th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 117–125, 2015.

[13] Carl-Gustav Esseen. On the liapunov limit error in the theory of probability. *Arkiv för Matematik, Astronomi och Fysik*, 28:1–19, 1942.

[14] Dan S Felsenthal, Zeev Maoz, and Amnon Rapoport. An empirical evaluation of six voting procedures: do they really make any difference? *British Journal of Political Science*, 23(1):1–27, 1993.

[15] William V Gehrlein. Condorcet's paradox and the likelihood of its occurrence: different perspectives on balanced preferences. *Theory and decision*, 52(2):171–199, 2002.

[16] Allan Gibbard. Manipulation of voting schemes. *Econometrica*, 41:587–602, 1973.

[17] James Green-Armytage, T Nicolaus Tideman, and Rafael Cosman. Statistical evaluation of voting rules. *Social Choice and Welfare*, 46(1):183–212, 2016.

[18] Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. In *Handbook of the fundamentals of financial decision making: Part I*, pages 99–127. World Scientific, 2013.

[19] Leonard Lee, On Amir, and Dan Ariely. In search of homo economicus: Cognitive noise and the role of emotion in preference consistency. *Journal of consumer research*, 36(2):173–187, 2009.

[20] Tyler Lu and Craig Boutilier. Learning mallows models with pairwise preferences. In *Proceedings of the 28th International Conference on Machine Learning (ICML)*, pages 145–152, 2011.

[21] Elchanan Mossel. Probabilistic view of voting, paradoxes, and manipulation. *AMERICAN MATHEMATICAL SOCIETY*, 59(3):297–330, 2022.

[22] Elchanan Mossel, Ariel D. Procaccia, and Miklós Z Rácz. A smooth transition from powerlessness to absolute power. *Journal of Artificial Intelligence Research*, 48:923–951, 2013.

[23] Hervé Moulin. Condorcet's principle implies the no show paradox. *Journal of Economic Theory*, 45(1):53–64, 1988.

[24] Florenz Plassmann and T Nicolaus Tideman. How frequently do different voting rules encounter voting paradoxes in three-candidate elections? *Social Choice and Welfare*, 42(1):31–75, 2014.

[25] Ariel D. Procaccia and Jeffrey S Rosenschein. The distortion of cardinal preferences in voting. In *Proceedings of the 10th International Workshop on Cooperative Information Agents (CIA)*, pages 317–331. Springer, 2006.

[26] Alexandros Psomas, Ariel Schvartzman, and S Matthew Weinberg. Smoothed analysis of multi-item auctions with correlated values. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 417–418, 2019.

[27] Tim Roughgarden. Beyond worst-case analysis. *Communications of the ACM*, 62(3):88–96, 2019.

[28] Daniel A. Spielman and Shang-Hua Teng. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. *Journal of the ACM*, 51(3):385–463, 2004.

[29] Lirong Xia. The smoothed possibility of social choice. arXiv:2006.06875, 2020.

[30] Lirong Xia. How likely are large elections tied? In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*, pages 884–885, 2021.

[31] Lirong Xia. A smoothed impossibility theorem on condorcet criterion and participation. arXiv:2107.06435, 2021.

[32] Lirong Xia. The smoothed satisfaction of voting axioms. arXiv:2106.01947, 2021.

[33] Lirong Xia. How likely a coalition of voters can influence a large election? arXiv:2202.06411, 2022.

[34] Lirong Xia and Vincent Conitzer. Generalized scoring rules and the frequency of coalitional manipulability. In *Proceedings of the 9th ACM Conference on Economics and Computation (EC)*, pages 109–118, 2008.

[35] Lirong Xia and Weiqiang Zheng. The smoothed complexity of computing kemeny and slater rankings. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, pages 5742–5750, 2021.

A Definitions of Voting Rules and Axioms

A.1 Voting Rules

We define the voting rules we study as functions of profile histograms $\mathbf{h}$ (rather than a profile $\boldsymbol{\pi}$). Further, we define them in the same form: first, we express how they assign candidates' scores, and then express via an $\arg\max_c$ that the winner or set of winners constitutes the candidate(s) with the highest (or in one case, lowest) score.

Positional Scoring Rules (PSRs). For fixed m, a positional scoring rule is characterized by a vector of weights of length m, $\boldsymbol{\alpha} = (\alpha_c | c \in [m])$, where $\alpha_1 \geq \alpha_2 \geq \cdots \geq \alpha_m$. Without loss of generality, we let these weights be translated and scaled such that $\alpha_1 = 1$ and $\alpha_m = 0$. The winner(s) by a PSR R, characterized by $\boldsymbol{\alpha}$,

$$R(\mathbf{h}) = \arg\max_c \sum_{j \in [m]} \alpha_j \sum_{\substack{\pi \in \mathcal{L} \\ \pi(j) = c}} h_\pi$$

Minimax. The Minimax winner is the candidate whose greatest pairwise defeat is the smallest:

$$\textsc{Minimax}(\mathbf{h}) = \arg\min_c \max_{c' \neq c} \sum_{\substack{\pi \in \mathcal{L} \\ c' \succ_\pi c}} h_\pi$$

Kemeny-Young. We define a candidate c's Kemeny-Young score to be, at a high level, the level of agreement with voters' rankings of the most-agreeing ranking that ranks c first. Then, the set of winners contains the candidate(s) with the highest Kemeny-Young score.

$$\textsc{Kemeny-Young}(\mathbf{h}) = \arg\max_c \max_{\pi : \pi(1) = c} \sum_{\substack{c', c'' \in C \\ c' \succ_\pi c''}} \sum_{\substack{\pi' \in \mathcal{L} \\ c' \succ_{\pi'} c''}} h_{\pi'}$$

Copeland.

$$\textsc{Copeland}(\mathbf{h}) = \arg\max_c \sum_{c' \neq c} \mathbb{I}[c \succ_{\mathbf{h}} c'] + 1/2 \cdot \mathbb{I}[c \sim_{\mathbf{h}} c']$$

A.2 Axioms

We use the following notation. For a ranking $\pi \in \mathcal{L}$, we let $\pi(j)$ for an index $j \in [m]$ be the candidate in the j'th position of π. For a ranking $\pi \in \mathcal{L}$ and distinct candidate $c, c' \in M$, we use $c \succ_\pi c'$ to denote that c is ranked higher than c' in π. $c \succ_{\boldsymbol{\pi}} c'$ when $|\{i \in [n] \ : \ c \succ_{\pi_i} c'\}| > |\{i \in [n] \ : \ c' \succ_{\pi_i} c\}|$.

The axioms we study in Section 3 are defined as follows:

Resolvability. A voting rule R satisfies Resolvability on $\boldsymbol{\pi}$ iff $|R(\boldsymbol{\pi})| = 1$ (i.e., there are no ties).

Condorcet. A *Condorcet winner* is a candidate that would win in a pairwise election against every other candidate. That is, c is a Condorcet winner in $\boldsymbol{\pi}$ if $c \succ_{\boldsymbol{\pi}} c'$ for all $c' \neq c$. A voting rule R satisfies Condorcet Consistency on a given profile $\boldsymbol{\pi}$ if one of two conditions hold: (1) there is no Condorcet winner in $\boldsymbol{\pi}$, or (2) there is a Condorcet winner c, and $R(\boldsymbol{\pi}) = \{c\}$.

Majority. A *majority winner* is a candidate that is ranked first by a majority of agents. That is, c is a majority winner in $\boldsymbol{\pi}$ if

$$|\{i \in [n] \, : \, \pi_i(c) = 1\}| > n/2$$

A voting rule R satisfies Majority on a profile $\boldsymbol{\pi}$ if it satisfies one of two conditions: (1) there is no majority winner in $\boldsymbol{\pi}$, or (2) if there is a majority winner a, then $R(\boldsymbol{\pi}) = \{a\}$.

Consistency. A voting rule R satisfies Consistency on a profile $\boldsymbol{\pi}$ if the following holds: for all partitions of $\boldsymbol{\pi}$ into sub-profiles, $(\boldsymbol{\pi}^1, \ldots, \boldsymbol{\pi}^t)$, if $R(\boldsymbol{\pi}^j)$ for all $j \in [t]$ is the same set of winners W, then $R(\boldsymbol{\pi}) = W$.

Independence of Irrelevant Alternatives (IIA). A voting rule R satisfies Independence of irrelevent alternatives (IIA) on profile $\boldsymbol{\pi}$ if the following holds. Suppose $R(\boldsymbol{\pi}) = a$. Then, for all candidates $b \neq a$, if $\boldsymbol{\pi}'$ is such that $a \succ_{\pi_i} b$ if and only if $a \succ_{\pi_i'} b$ for all voters i, then $R(\boldsymbol{\pi}') \neq b$.

The axioms we study in Section 4 are defined formally as follows:

Definition 24 (**$\rho(n)$-group-strategyproofness**). A voting rule R satisfies $\rho(n)$-group-strategyproofness on a profile $\boldsymbol{\pi}$ if there exists no group of agents of size at most $\rho(n)$ such that if they change their votes, resulting in some profile $\boldsymbol{\pi}'$ with outcome $R(\boldsymbol{\pi}')$, they are all at least as well off and at least one is strictly better off. An agent is at least as well (resp. strictly better) off if their favorite candidate in the set $R(\boldsymbol{\pi}')$ is weakly (resp. strictly) preferred to their favorite candidate in $R(\boldsymbol{\pi})$.

Definition 25 (**$\rho(n)$-group-monotonicity**). A voting rule R satisfies $\rho(n)$-group-monotonicity on a profile $\boldsymbol{\pi}$ if there exists no candidate a and no group of agents of size at most $\rho(n)$ such that if they change their votes *without decreasing the position of a in any of their rankings*, producing a new profile $\boldsymbol{\pi}'$, then it cannot be that $a \in R(\boldsymbol{\pi})$ and $a \notin R(\boldsymbol{\pi}')$.

Definition 26 (**$\rho(n)$-group-participation**). A voting rule R satisfies $\rho(n)$-group-participation on a profile $\boldsymbol{\pi}$ if there exists no group of agents of size at most $\rho(n)$ such that if they collectively leave the election, producing a new profile $\boldsymbol{\pi}'$, they are all at least as well off and at least one is strictly better off with the outcome $R(\boldsymbol{\pi}')$ than $R(\boldsymbol{\pi})$.

B Proofs and Definitions from Section 2

B.1 Proof of Lemma 1

We first express the expectation and variance of $\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi}))$ in terms of the analogous values for the $\mathcal{H}(\mathcal{S}_\phi(\pi_i))$. The relationships between these quantities are shown below, derived by applying simple properties of the expectation and variance in conjunction with Equation (1):

$$\mathbb{E}[\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi}))] = {}^1\!/\!_n \sum_{i=1}^{n} \mathbb{E}[\mathcal{H}(\mathcal{S}_\phi(\pi_i))], \qquad \mathrm{Cov}[\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi}))] = {}^1\!/\!_{n^2} \sum_{i=1}^{n} \mathrm{Cov}[\mathcal{H}(\mathcal{S}_\phi(\pi_i))]$$

We now use these relationships to find closed forms for each of these objects. Note that $\mathbb{E}[\mathcal{H}(\mathcal{S}_\phi(\pi_i))]$ is an $|\mathcal{L}_{-1}|$-length vector whose π-th component is simply $\Pr[\mathcal{S}_\phi(\pi_i) = \pi]$. $\mathrm{Cov}[\mathcal{H}(\mathcal{S}_\phi(\pi_i))]$ is a $|\mathcal{L}_{-1}| \times |\mathcal{L}_{-1}|$ matrix whose entries each correspond to a pair of rankings π, π', such that the π, π'-th entry is equal to the covariance between the random variables $\mathcal{H}(\mathcal{S}_\phi(\pi_i))_\pi$ and $\mathcal{H}(\mathcal{S}_\phi(\pi_i))_{\pi'}$.

The covariance matrix $\mathrm{Cov}[\mathcal{H}(\mathcal{S}_\phi(\pi_i))]$ is a $|\mathcal{L}_{-1}| \times |\mathcal{L}_{-1}|$ matrix whose entries each correspond to a pair of rankings π, π', such that the π, π'-th entry is equal to the covariance between the random variables $\mathcal{H}(\mathcal{S}_\phi(\pi_i))_\pi$ and $\mathcal{H}(\mathcal{S}_\phi(\pi_i))_{\pi'}$. Given that $\mathcal{H}(\mathcal{S}_\phi(\pi_i))$ can take on the values of only basis vectors, the values of these random variables are either 0 or 1. To characterize these entries, we will use the fact that for general X and Y, $\mathrm{Cov}(X, Y) = \mathbb{E}[XY] - \mathbb{E}[X]\mathbb{E}[Y]$.

For distinct rankings $\pi \neq \pi'$, at most one of $\mathcal{H}(\mathcal{S}_\phi(\pi_i))_\pi$ and $\mathcal{H}(\mathcal{S}_\phi(\pi_i))_{\pi'}$ can be nonzero, so the expectation their product must be 0. Then,

$$\mathrm{Cov}\left(\mathcal{H}(\mathcal{S}_\phi(\pi_i))_\pi,\, \mathcal{H}(\mathcal{S}_\phi(\pi_i))_{\pi'}\right) = 0 - \Pr[\mathcal{S}_\phi(\pi_i) = \pi] \cdot \Pr[\mathcal{S}_\phi(\pi_i) = \pi'].$$

For diagonal entries where $\pi = \pi'$, since the values of our random variables are always 0 or 1, we have

$$\mathrm{Cov}\left(\mathcal{H}(\mathcal{S}_\phi(\pi_i))_\pi,\, \mathcal{H}(\mathcal{S}_\phi(\pi_i))_\pi\right) = \Pr[\mathcal{S}_\phi(\pi_i) = \pi] - \Pr[\mathcal{S}_\phi(\pi_i) = \pi]^2.$$

With these in hand, we now prove the lemma statement. Fix $\mathcal{S}$, ϕ, and $\boldsymbol{\pi}$. Recall that $\mathrm{Cov}[\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi}))] = {}^1\!/\!_{n^2} \sum_{i=1}^{n} \mathrm{Cov}[\mathcal{H}(\mathcal{S}_\phi(\pi_i))]$. Hence, we first consider $\mathrm{Cov}[\mathcal{H}(\mathcal{S}_\phi(\pi))]$ for individual rankings π.

Fix an arbitrary ranking π. First, we will prove that all eigenvalues of $\mathrm{Cov}[\mathcal{H}(\mathcal{S}_\phi(\pi))]$ are positive and the minimum is at least $\textsc{min-prob}(\mathcal{S}_\phi)/m!$.

To simplify notation in the subsequent computations, for the j'th ranking π', let $q_j = \Pr[\mathcal{S}_\phi(\pi) = \pi']$. Recall that in $\mathrm{Cov}[\mathcal{H}(\mathcal{S}_\phi(\pi))]$, the (j, k)-th entry when $j = k$ (a diagonal entry) has value $q_j(1 - q_j)$ and for $j \neq k$, the entry has value $-q_j \cdot q_k$. We can then write the covariance matrix as

$$\mathrm{Cov}[\mathcal{H}(\mathcal{S}_\phi(\pi))] = \begin{pmatrix} q_1(1-q_1) & -q_1 \cdot q_2 & \cdots & -q_1 \cdot q_{m!-1} \\ -q_2 \cdot q_1 & q_2(1-q_2) & \cdots & \vdots \\ \vdots & \ddots & \ddots & \vdots \\ -q_{m!-1} \cdot q_1 & \cdots & \cdots & q_{m!-1}(1-q_{m!-1}) \end{pmatrix}.$$

Note that $q_{m!}$ is the probability of the "missing" ranking, and that $\sum_{j=1}^{m!} q_j = 1$. We first demonstrate an inverse of $\mathrm{Cov}[\mathcal{H}(\mathcal{S}_\phi(\pi))]$. Consider the matrix:

$$M^{inv} = \begin{pmatrix} \frac{1}{q_1} + \frac{1}{q_{m!}} & \frac{1}{q_{m!}} & \cdots & \frac{1}{q_{m!}} \\ \frac{1}{q_{m!}} & \frac{1}{q_2} + \frac{1}{q_{m!}} & \cdots & \vdots \\ \vdots & \ddots & \ddots & \vdots \\ \frac{1}{q_{m!}} & \cdots & \cdots & \frac{1}{q_{m!-1}} + \frac{1}{q_{m!}} \end{pmatrix}.$$

More formally, the jth diagonal entry is $1/q_j - 1/q_{m!}$ and all off-diagonal entries are simply $1/q_{m!}$.

We now show that $\mathrm{Cov}[\mathcal{H}(\mathcal{S}_\phi(\pi))] \cdot M^{inv} = I_{m!-1}$ where $I_{m!-1}$ is the identity matrix, that is, M^{inv} is in fact the inverse of $\mathrm{Cov}[\mathcal{H}(\mathcal{S}_\phi(\pi))]$. To that end, let us consider the i'th diagonal entry of the product. It is precisely

$$\begin{aligned}
(\mathrm{Cov}[\mathcal{H}(\mathcal{S}_\phi(\pi))] \cdot M^{inv})_{jj} &= \sum_{k=1,k\neq j}^{m!-1} -\frac{q_j q_k}{q_{m!}} + q_j(1-q_j)\left(\frac{1}{q_j} + \frac{1}{q_{m!}}\right) \\
&= \sum_{k=1,k\neq j}^{m!-1} -\frac{q_j q_k}{q_{m!}} + (1-q_j) + \frac{q_j}{q_{m!}}(1-q_j) \\
&= \sum_{k=1,k\neq j}^{m!-1} -\frac{q_j q_k}{q_{m!}} + 1 - q_j + \frac{q_j}{q_{m!}} - \frac{q_j \cdot q_j}{q_{m!}} \\
&= \sum_{k=1}^{m!-1} -\frac{q_j q_k}{q_{m!}} + 1 - q_j + \frac{q_j}{q_{m!}} \\
&= \frac{-q_j(1-q_{m!})}{q_{m!}} + 1 - q_j + \frac{q_j}{q_{m!}} \\
&= \frac{-q_j + q_j \cdot q_{m!}}{q_{m!}} + 1 - \frac{q_j \cdot q_{m!}}{q_{m!}} + \frac{q_j}{q_{m!}} \\
&= 1.
\end{aligned}$$

For a non-diagonal entry j, k with $j \neq k$, we have

$$\begin{aligned}
(\mathrm{Cov}[\mathcal{H}(\mathcal{S}_\phi(\pi))] \cdot M^{inv})_{jk} &= \sum_{\ell=1,\ell\neq j,k}^{m!} -\frac{q_j q_\ell}{q_{m!}} + \frac{q_j(1-q_j)}{q_{m!}} - q_j q_k \cdot \left(\frac{1}{q_k} + \frac{1}{q_{m!}}\right) \\
&= \sum_{\ell=1,\ell\neq j,k}^{m!} -\frac{q_j q_\ell}{q_{m!}} + \frac{q_j}{q_{m!}} - \frac{q_j q_j}{q_{m!}} - q_j - \frac{q_j q_k}{q_{m!}} \\
&= \sum_{\ell=1}^{m!} -\frac{q_j q_\ell}{q_{m!}} + \frac{q_j}{q_{m!}} - q_j \\
&= \frac{q_j(1-q_{m!})}{q_{m!}} + \frac{q_j}{q_{m!}} - \frac{q_j q_{m!}}{q_{m!}} \\
&= 0.
\end{aligned}$$

We now consider the eigenvalues of $\mathrm{Cov}[\mathcal{H}(\mathcal{S}_\phi(\pi))]$. Since it is a covariance matrix, it is symmetric, and therefore positive semi-definite. Since we now know it is invertible, it is in fact positive definite. This implies all of its eigenvalues exist and are positive. Further, since the eigenvalues of $M^{inv} = \mathrm{Cov}[\mathcal{H}(\mathcal{S}_\phi(\pi))]^{-1}$ are the reciprocals of the eigenvalues of $\mathrm{Cov}[\mathcal{H}(\mathcal{S}_\phi(\pi))]$, we can lower bound the eigenvalues of $\mathrm{Cov}[\mathcal{H}(\mathcal{S}_\phi(\pi))]$ by upper bounding the the eigenvalues of M^{inv}.

To upper bound the maximum eigenvalue of M^{inv}, we can upperbound the maximum absolute row sum. Note that the sum of row j is

$$\frac{1}{q_j} + \frac{m!-1}{q_{m!}} \leq \frac{m!}{\textsc{min-prob}(S_\phi)}.$$

This lower bounds the minimum eigenvalue of $\mathrm{Cov}[\mathcal{H}(\mathcal{S}_\phi(\pi))]$ by $\frac{\textsc{min-prob}(S_\phi)}{m!}$, as needed.

For our fixed profile $\boldsymbol{\pi}$, we now have that each $\mathrm{Cov}[\mathcal{H}(\mathcal{S}_\phi(\pi_i))]$ has minimum eigenvalue at least $\frac{\textsc{min-prob}(S_\phi)}{m!}$. Since the minimum eigenvalue of the sum of matrices is at least the sum of the minimum eigenvalues of each matrix, the minimum eigenvalue of $\sum_{i=1}^n \mathrm{Cov}[\mathcal{H}(\mathcal{S}_\phi(\pi_i))]$ is at least $n \cdot \frac{\textsc{min-prob}(S_\phi)}{m!}$. Finally, to get $\mathrm{Cov}[\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi}))]$, we scale this sum down by n^2 which scales the minimum eigenvalues equivalently, yielding a minimum eigenvalue of at least $\frac{\textsc{min-prob}(S_\phi)}{nm!}$. Note that the minimum eigenvalue here is positive, meaning $\mathrm{Cov}[\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi}))]$ has no zero eigenvalues, meaning it is invertible. ☐

B.2 Proof of Lemma 2

Fix $\mathcal{S}, \phi$, and $\boldsymbol{\pi} \in \Pi_n$. Since $\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi})) = {}^1\!/\!_n \sum_{i=1}^n \mathcal{H}(\mathcal{S}_\phi(\pi_i))$ where each of these summands is independent, our goal will be to apply a version of the Berry-Esseen bound, as stated below:

Lemma 27 (Restatement of Berry-Esseen as in [3]). Let $Y_1, \ldots, Y_n$ be independent, mean-zero, $\mathbb{R}^{m!-1}$-valued random variables. Let $S = Y_1 + \cdots + Y_n$, and let C^2 be the covariance matrix of S, assumed invertible. Let $\mathcal{N}(0, C^2)$ be a $m! - 1$-dimensional Gaussian with mean zero and covariance C^2. Then for any convex subset $X \subseteq \mathbb{R}^{m!-1}$,

$$|\Pr[S \in X] - \Pr[\mathcal{N}(0, C^2) \in X]| \leq O((m! - 1)^{1/4}) \cdot \left(\sum_{i=1}^n \mathbb{E}[|C^{-1} Y_i|^3] \right).$$

In order to apply the Berry-Esseen bound, we use the properties of the covariance matrix $\mathrm{Cov}[\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi}))]$ from Lemma 1.

By Lemma 1, $\mathrm{Cov}[\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi}))]$ is invertible, so we can apply the Berry-Esseen bound. For consistency with the form of the stated bound, we first translate both our distribution, $\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi}))$, and the Gaussian to which we want to show it converges, to make both mean-zero. That is, we will show the equivalent statement about $\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi})) - \mathbb{E}[\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi}))]$ approaching $\mathcal{N}(0, \mathrm{Cov}[\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi}))])$. Note that subtracting the expectations of both $\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi}))$ and $\mathcal{N}(\mathbb{E}[\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi}))], \mathrm{Cov}[\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi}))])$ translates the the distributions identically, and note that the convexity of the quantified sets X is invariant under translation. Therefore, proving the claim on the translated version of our distribution implies the claim on our original distribution.

By linearity of expectation, we can express our translated distribution as the sum of n independent random variables:

$$\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi})) - \mathbb{E}[\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi}))] = \sum_{i=1}^n {}^1\!/\!_n(\mathcal{H}(\mathcal{S}_\phi(\pi_i)) - \mathbb{E}[\mathcal{H}(\mathcal{S}_\phi(\pi_i))]).$$

We let Y_i be the random variable distributed as a single term of the above sum, ${}^1\!/\!_n(\mathcal{H}(\mathcal{S}_\phi(\pi_i)) - \mathbb{E}[\mathcal{H}(\mathcal{S}_\phi(\pi_i))])$. Note that Y_i has mean zero, and covariance ${}^1\!/\!_{n^2} \cdot \mathrm{Cov}[\mathcal{H}(\mathcal{S}_\phi(\pi_i))]$.

Now, all that remains to show is that the following bound on the convergence rate holds:

$$O((m! - 1)^{1/4}) \cdot \left(\sum_{i=1}^n \mathbb{E}\left[\left| \mathrm{Cov}[\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi}))]^{-1/2} Y_i \right|^3 \right] \right) \leq \frac{O((m! - 1)^{1/4})}{(\lambda^{min,\mathcal{S},\phi})^{3/2}} \cdot \frac{1}{\sqrt{n}}.$$

We note that the exponentiated $\mathrm{Cov}[\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi}))]^{-1/2}$ in this expression is well defined because the matrix is symmetric.

We first show that $|Y_i| \leq 2/n$ for all i, where $|Y_i|$ denotes the L_2 norm of Y_i. Recall that $\mathcal{H}(\mathcal{S}_\phi(\pi_i))$ is always either a basis vector or the all 0s vector, and $\mathbb{E}[\mathcal{H}(\mathcal{S}_\phi(\pi_i))]$ is the vector whose entry corresponding to π' is $\Pr[\mathcal{S}_\phi(\pi) = \pi']$. Then, after subtracting the second vector from the first, the negative entries in the resulting vector can sum in magnitude to at most the sum of these probabilities, which is at most 1. Similarly, the positive entries can also sum to at most 1. hence, the L_1 norm before scaling by $1/n$ is at most 2. Using the fact that L_2 norms are at most L_1 norms, we get that this continues to hold for the L_2 norm. After dividing by n, we get that $|Y_i| \leq 2/n$, as needed.

Now, per Lemma 1, $\mathrm{Cov}[\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi}))]$ has minimum eigenvalue at least $\frac{\textsc{min-prob}(S)}{m!n}$. It therefore holds that $\mathrm{Cov}[\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi}))]^{-1/2}$ has maximum eigenvalue at most

$$\left(\frac{\textsc{min-prob}(S)}{m!n}\right)^{-1/2} = \sqrt{\frac{m!n}{\textsc{min-prob}(S)}}.$$

Multiplying a vector by a matrix can scale the norm by at most the matrix's maximum eigenvalue. Thus, combined with our observation that $|Y_i| \leq 2/n$, the following bound will always hold:

$$|\,\mathrm{Cov}[\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi}))]^{-1/2} Y_i|^3 \leq \left(\sqrt{\frac{m!n}{\textsc{min-prob}(S)}} \cdot \frac{2}{n}\right)^3 = \frac{8 \cdot (m!)^{3/2}}{n^{3/2} \cdot \textsc{min-prob}(\mathcal{S}_\phi)^{3/2}}$$

Because this bound holds deterministically on the term above, it must hold also for the expectation of the term above. Thus, by summing over all i, we get that

$$\sum_{i=1}^{n} \mathbb{E}[|\,\mathrm{Cov}[\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi}))]^{-1/2} Y_i|^3] \leq \frac{8\,(m!)^{3/2}}{\sqrt{n} \cdot \textsc{min-prob}(\mathcal{S}_\phi)^{3/2}}.$$

Since $O\left((m!-1)^{1/4}\right) \cdot 8m!^{3/2} \in O\left((m!)^{7/4}\right)$, multiplying by the Berry-Esseen constant yields the lemma statement. $\square$

B.3 Proof of Lemma 3

Fix $\mathcal{S}$, ϕ, and $\boldsymbol{\pi}$. Recall that $\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi})) = \sum_{i=1}^{n} {}^1\!/\!_n \mathcal{H}(\mathcal{S}_\phi(\pi_i))$, where each term of this sum is an independently drawn basis (or all 0s) vector, scaled down by $1/n$.

We begin by considering each entry of $\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi}))$ individually, and for each such entry, showing that with probability at least $1 - 2\exp(-2\varepsilon^2 n/m!)$, this entry is within $\varepsilon/m!$ of its expected value. To this end, fix a ranking π. Then, we are considering the concentration of the following random variable, written as a sum of scaled indicators: $\mathcal{H}(S_\phi(\boldsymbol{\pi}))_\pi = \sum_{i=1}^{n} {}^1\!/\!_n \mathcal{H}(\mathcal{S}_\phi(\pi_i))_\pi$, where each term of the sum ${}^1\!/\!_n \mathcal{H}(\mathcal{S}_\phi(\pi_i))_\pi$ is bounded in $[0, {}^1\!/\!_n]$.

A direct application of Hoeffding's inequality[13] shows that $\mathcal{H}(S_\phi(\boldsymbol{\pi}))_\pi$ is within $\varepsilon/m!$ of its mean with probability at least $1 - 2\exp(-2\varepsilon^2 n/m!)$. Union bounding over all $m! - 1$ entries of $\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi}))$ shows this will hold for all entries simultaneously with probability at least $1 - 2m!\exp(-2\varepsilon^2 n/m!)$. Conditioned on this, note that the L_1 distance between the realized vector and the expected value vector is at most ε. Since L_2 distances are upper bounded by L_1 distances, this implies the desired result.

[13] For independent random variables $X_1, \ldots, X_n$ bounded in $[a, b]$, their sum $S_n = \sum_{i=1}^{n} X_i$ is such that $\Pr[|S_n - \mathbb{E}[S_n]| \geq t] \leq 2\exp\left(-\frac{2t^2}{n(b-a)^2}\right)$.

B.4 Proof of Lemma 4

Fix a noise model $\mathcal{S}$ and $\varepsilon > 0$. By continuity of $\mathcal{S}$, there exists a $\phi > 0$ such that for all $\phi' \in [0, \phi]$, $\Pr[\mathcal{S}_{\phi'}(\pi) = \pi] > 1 - \varepsilon/2$. Choose this to be our ϕ.

Fix such a ϕ' and a profile $\boldsymbol{\pi} \in \Pi_n$. Notice that after applying $\mathcal{S}$, each ranking π_i will stay the same with probability at least $1 - \varepsilon/2$. As this is independent accross voters, a straightforward application of Hoeffding's inequality tells us that at most a $1 - \varepsilon$ fraction of rankings will not change with probability at least $1 - \exp(\varepsilon^2 n/2)$, as needed. □

C Supplemental materials for Section 3

C.1 Proof of Proposition 10

For a decisive hyperplane rule, note that all profiles that fail Resolvablity must fall on one of the ℓ hyperplanes (where by definition ℓ is finite). Further, note that these hyperplanes are convex and measure-zero. Hence, we can immediately apply Theorem 7 using the hyperplanes as convex sets.

For non-decisive hyperplane rules, there must exist a profile $\boldsymbol{\pi}$ not lying on any hyperplane for which Resolvablity fails. Further, $\mathbf{h}^{\boldsymbol{\pi}}$ is at least some L^1 distance $\varepsilon > 0$ from all hyperplanes. All such profiles with histograms in this ball have the same outcome as $\boldsymbol{\pi}$ and hence do not satisfy Resolvablity. We can then directly apply Theorem 8 using this profile $\boldsymbol{\pi}$ and $r = \varepsilon$.

C.2 Proof of Proposition 11

We prove this result for each axiom separately. For each axiom, we define a sufficient condition for a counterexample $\boldsymbol{\pi}$ to be "strict", i.e., all profiles $\boldsymbol{\pi}'$ with histograms nearby to $\mathbf{h}^{\boldsymbol{\pi}}$ will also be counterexamples. We then show that the existence of a strict counterexample implies smoothed-violation via Theorem 8. We later give (or point to existing) strict counterexamples for all relevant pairs of rules and axioms. In the following arguments, we say a profile $\boldsymbol{\pi}$ is *robust* with respect to a voting rule R if there is an $\varepsilon > 0$ such that all profiles $\boldsymbol{\pi}'$ with $\|\mathbf{h}^{\boldsymbol{\pi}} - \mathbf{h}^{\boldsymbol{\pi}'}\|_1 < \varepsilon$, $R(\boldsymbol{\pi}') = R(\boldsymbol{\pi})$. Notice that for hyperplane rules, all profiles that do not fall on hyperplanes are robust.

Condorcet: We say a counterexample $\boldsymbol{\pi}$ is a strict counterexample to R satisfying Condorcet if $\boldsymbol{\pi}$ is robust with respect to R, $\boldsymbol{\pi}$ has a strict Condorcet winner a (i.e., a candidate that beats every other candidate on *strictly more* than half of the voters), and $R(\boldsymbol{\pi}) \neq a$. If such a strict counterexample $\boldsymbol{\pi}$ exists, if it is robust with value ε_1 and a wins with at least a $1/2 + \varepsilon_2$ fraction against each candidate, all profiles whose histogram falls within $\varepsilon = \min(\varepsilon_1, \varepsilon_2)$ of $\mathbf{h}^{\boldsymbol{\pi}}$ have the same Condorcet winner and same output under R, and thus are also counterexamples. Hence, Theorem 8 implies smoothed-violation.

Majority: We say a counterexample $\boldsymbol{\pi}$ is a strict counterexample to R satisfying Majority if $\boldsymbol{\pi}$ is robust and $\boldsymbol{\pi}$ has a strict Majority winner a (i.e., a candidate that is ranked first by *strictly more* than half of the voters), and $R(\boldsymbol{\pi}) \neq a$. If such a strict counterexample $\boldsymbol{\pi}$ exists, if it is robust with value ε_1 and a is ranked first by at least a $1/2 + \varepsilon_2$ fraction of the voters, all profiles whose histogram falls within $\varepsilon = \min(\varepsilon_1, \varepsilon_2)$ of $\mathbf{h}^{\boldsymbol{\pi}}$ have the same Condorcet winner and same output under R, and thus are also counterexamples. Hence, Theorem 8 implies smoothed-violation.

Consistency: We say a counterexample $\boldsymbol{\pi}$ is a strict counterexample to R satisfying Consistency if $\boldsymbol{\pi} = \boldsymbol{\pi}^1 \cup \cdots \cup \boldsymbol{\pi}^t$ such that $R(\boldsymbol{\pi}^1) = \cdots = R(\boldsymbol{\pi}^t) \neq R(\boldsymbol{\pi})$ and all of $\boldsymbol{\pi}^1, \ldots, \boldsymbol{\pi}^t$ and $\boldsymbol{\pi}$ are robust with respect to R. Suppose such a strict counterexample $\boldsymbol{\pi} = \boldsymbol{\pi}^1 \cup \cdots \cup \boldsymbol{\pi}^t$ exists. Let ε^{min} be an amount by which all the relevant profiles are robust. Suppose each profile $\boldsymbol{\pi}^j$ has n_j voters and let $n = n_1 + \cdots + n_t$ be the number of voters in $\boldsymbol{\pi}$. Let $p = \min_j n_j/n$. Notice that any $\boldsymbol{\pi}'$ on zn voters within $\varepsilon = p\varepsilon^{min}$ of $\boldsymbol{\pi}$ can be decomposed into $\boldsymbol{\pi}^{1'}, \ldots, \boldsymbol{\pi}^{t'}$ such that each $\boldsymbol{\pi}^{j'}$ has zn_j voters and $\mathbf{h}^{\boldsymbol{\pi}^{j'}}$ is at most ε^{min} away from $\mathbf{h}^{\boldsymbol{\pi}^j}$. Hence, $R(\boldsymbol{\pi}^{1'}) = \cdots = R(\boldsymbol{\pi}^{t'}) \neq R(\boldsymbol{\pi}')$, so $\boldsymbol{\pi}'$ is a counterexample to Consistency. Hence, Theorem 8 implies smoothed-violation.

IIA: We say a counterexample $\boldsymbol{\pi}^1$ is a strict counterexample to R satisfying IIA if there is another profile $\boldsymbol{\pi}^2$ such that both $\boldsymbol{\pi}^1$ and $\boldsymbol{\pi}^2$ are robust with respect to R, and there are candidates a, b such that the relative ranking of a and b are the same under π_i^1 and π_i^2 for all voters i, yet $R(\boldsymbol{\pi}^1) = a$ and $R(\boldsymbol{\pi}^2) = b$. Suppose such a strict counterexample $\boldsymbol{\pi}^1$ and $\boldsymbol{\pi}^2$ that are both robust by at least ε. Consider any profile $\boldsymbol{\pi}^{1'}$ where $\mathbf{h}^{\boldsymbol{\pi}^{1'}}$ is within ε of $\mathbf{h}^{\boldsymbol{\pi}^1}$. Notice that there must be a profile $\boldsymbol{\pi}^{2'}$ such that $\mathbf{h}^{\boldsymbol{\pi}^{2'}}$ is within ε of $\mathbf{h}^{\boldsymbol{\pi}^2}$ that matches $\boldsymbol{\pi}^{1'}$ in terms of all voter's relative rankings of a and b, and by robustness $R(\boldsymbol{\pi}^{1'}) = a$ while $R(\boldsymbol{\pi}^{2'}) = b$. Hence, Theorem 8 implies smoothed-violation.

Counterexamples: The following table points to strict counterexamples that can be used to show the claims above. Most can be found on Wikipedia pages. The missing ones are presented afterwards.

	Axioms			
Voting Rules	CONDORCET	MAJORITY	CONSISTENCY	IIA
PLURALITY	**[a]**	satisfied	satisfied	**[c]**
(non-PLURALITY) PSRs	**Example 28**	**Example 28**	satisfied	**Example 29**
MINIMAX	satisfied	satisfied	**[b]**	**[c]**
KEMENY-YOUNG	satisfied	satisfied	**[b]**	**[c]**
COPELAND	satisfied	satisfied	**[b]**	**[c]**

Table 2: Externally-referenced examples come from the following Wikipedia pages [a]: *Condorcet Criterion*, [b]: *Consistency Criterion*, [c]: *Independence of Irrelevant Alternatives*.

Example 28. Let R be the positional scoring rule represented by the weights vector $(1, \alpha, \ldots, 0)$ without loss of generality. We assume $\alpha > 0$ is separated from 0 (otherwise, R is just Plurality). Let

$$\boldsymbol{\pi} : \begin{array}{r|l} n\left(\frac{1}{2} - \frac{\alpha}{4(2-\alpha)}\right) \text{ voters} & \boldsymbol{a_1} \succ a_3 \succ \cdots \succ a_m \succ \boldsymbol{a_2} \\ n\left(\frac{1}{2} + \frac{\alpha}{4(2-\alpha)}\right) \text{ voters} & \boldsymbol{a_2} \succ \boldsymbol{a_1} \succ a_3 \succ \cdots \succ a_m \end{array}.$$

Example 29. Let R be the positional scoring rule represented by the weights vector $(1, \alpha, \ldots, 0)$ without loss of generality. We assume $\alpha > 0$ is separated from 0 (otherwise, R is just Plurality).

$$\boldsymbol{\pi}^1 : \begin{array}{r|l} n/2 \text{ voters} & \boldsymbol{a_1} \succ a_3 \succ \cdots \succ a_m \succ \boldsymbol{a_2} \\ n/4 \text{ voters} & \boldsymbol{a_2} \succ a_3 \succ \cdots \succ a_m \succ \boldsymbol{a_1} \\ n/4 \text{ voters} & \boldsymbol{a_2} \succ \boldsymbol{a_1} \succ a_3 \succ \cdots \succ a_m \end{array} \qquad \boldsymbol{\pi}^2 : \begin{array}{r|l} n/2 \text{ voters} & \boldsymbol{a_1} \succ \boldsymbol{a_2} \succ a_3 \succ \cdots \succ a_m \\ n/4 \text{ voters} & \boldsymbol{a_2} \succ a_3 \succ \cdots \succ a_m \succ \boldsymbol{a_1} \\ n/4 \text{ voters} & \boldsymbol{a_2} \succ \boldsymbol{a_1} \succ a_3 \succ \cdots \succ a_m \end{array}.$$

This concludes the proof. □

D Supplemental materials for Section 4

D.1 Proof of Theorem 13

Proof. Fix $\mathcal{S}$, ϕ, and $\delta(n)$. Next, fix a number n, a hyperplane $G \in \mathcal{G}$, a $\phi' \in [\phi, 1]$, and a $\boldsymbol{\pi}^\star \in \Pi_n$. Here, G is a hyperplane in $m! - 1$ dimensions, meaning it is defined by a linear equation whose variables are the profile proportions for rankings in $\mathcal{L}_{-1}$, each weighted by a coefficient, a_π, along with a constant, b. In other words, it is the set of proportions $\mathbf{h}$ satisfying

$$b = \sum_{\pi \in \mathcal{L}_{-1}} a_\pi \cdot h_\pi. \tag{2}$$

Formally, we need to consider only the intersection of G with our convex hull H, but for our results, it is irrelevant whether we formally make this restriction.

We will upper bound $\Pr[d(\mathcal{H}(\mathcal{S}_{\phi'}(\boldsymbol{\pi}^\star)), G) \leq \delta(n)]$ by an $o(1)$ function of n (allowing this convergence rate to depend on ϕ, m, and $\mathcal{S}$, but not G, ϕ', or $\boldsymbol{\pi}^\star$).

Recall that the hyperplane G has a coefficient for each ranking except the "missing" ranking, π_{-1}. We choose $\pi^{max} \in \mathcal{L}$ such that π^{max} is the ranking corresponding to a largest-magnitude coefficient in the definition of G, i.e., $\pi^{max} \in \operatorname{argmax}_{\pi \in \mathcal{L}_{-1}} |a_\pi|$. Now, we will define two types of events.

First, let $\mathcal{E}^{few} = \{\boldsymbol{\pi} \in \Pi_n \,:\, |\{i \in N \,:\, \pi_i \in \{\pi^{max}, \pi_{-1}\}\}| < np\}$ – that is, $\mathcal{E}^{few}$ is the set of profiles in which fewer than np agents end up voting either π^{max} or π_{-1} in $\mathcal{H}(\mathcal{S}_{\phi'}(\boldsymbol{\pi}^\star))$.

Let $\mathcal{V}^{\geq np} = \{V \subseteq N \,:\, |V| \geq np\}$ be the collection of all sets of agents of size at least np. For a set of such agents $V \in \mathcal{V}^{\geq np}$, we denote the complement of this set of agents as $\overline{V} = N \setminus V$. Now, slightly abusing notation, define $\widetilde{\Pi}^{\overline{V}} = \left\{(\pi_i)_{i \in \overline{V}} \,:\, \pi_i \in \mathcal{L} \setminus \{\pi^{max}, \pi_{-1}\}\right\}$ to be the set of all *partial* profiles in which all agents in $\overline{V}$ have a ranking other than π^{max} or π_{-1}. For $V \in \mathcal{V}^{\geq np}$ and partial profile $\widetilde{\boldsymbol{\pi}} \in \widetilde{\Pi}^{\overline{V}}$, let $\mathcal{E}^{V,\widetilde{\boldsymbol{\pi}}}$ be the event that agents in V either vote π or π', and all other agents vote as in $\widetilde{\boldsymbol{\pi}}$, that is,

$$\mathcal{E}^{V,\widetilde{\boldsymbol{\pi}}} = \left\{\boldsymbol{\pi} \in \Pi \,:\, \pi_i = \widetilde{\pi}_i \text{ for } i \in \overline{V} \text{ and } \pi_i \in \{\pi^{max}, \pi_{-1}\} \text{ for } i \in V\right\}.$$

In the remainder of the proof, we will use that the event $\mathcal{E}^{few}$, along with the set of all events of the form $\mathcal{E}^{V,\widetilde{\boldsymbol{\pi}}}$ (that is, for all subsets of agents $V \in \mathcal{V}^{\geq np}$ and partial profiles $\widetilde{\boldsymbol{\pi}} \in \widetilde{\Pi}^{\overline{V}}$), form a partition of the full space of profiles. To see how we will use this, recall that we are trying to show that $d(\mathcal{H}(\mathcal{S}_{\phi'}(\boldsymbol{\pi}^\star), G)$ is likely to be somewhat large. We will show this by showing that, conditioned on any individual event $\mathcal{E}^{V,\widetilde{\boldsymbol{\pi}}}$, the probability of this distance being large is fairly high, and the remaining event, $\mathcal{E}^{few}$ is unlikely to occur.

We will show two claims:

1. There is some $f(n) \in o(1)$ such that
$$\Pr\left[\mathcal{E}^{few}\right] \leq f(n).$$

2. There is some $g(n) \in o(1)$ such that for all $\mathcal{E}^{V,\boldsymbol{\pi}}$,
$$\Pr\left[d(\mathcal{H}(\mathcal{S}_{\phi'}(\boldsymbol{\pi}^\star)), G) \leq \delta(n) \mid \mathcal{E}^{V,\widetilde{\boldsymbol{\pi}}}\right] \leq g(n).$$

We first show that together these are sufficient to prove the bound. Indeed, by the law of total probability, as these events form a partition

$$\Pr\left[[d(\mathcal{H}(\mathcal{S}_{\phi'}(\boldsymbol{\pi}^\star)), G) \leq \delta(n)\right] = \Pr\left[d(\mathcal{H}(\mathcal{S}_{\phi'}(\boldsymbol{\pi}^\star)), G) \leq \delta(n) \mid \mathcal{E}^{few}\right] \cdot \Pr\left[\mathcal{E}^{few}\right]$$

$$
\begin{aligned}
&+ \sum_{\mathcal{E}^{V,\boldsymbol{\pi}}} \Pr\left[d(\mathcal{H}(\mathcal{S}_{\phi'}(\boldsymbol{\pi}^\star)), G) \le \delta(n) \mid \mathcal{E}^{V,\boldsymbol{\pi}}\right] \cdot \Pr\left[\mathcal{E}^{V,\boldsymbol{\pi}}\right] \\
&\le \Pr\left[d(\mathcal{H}(\mathcal{S}_{\phi'}(\boldsymbol{\pi}^\star)), G) \le \delta(n) \mid \mathcal{E}^{few}\right] \cdot f(n) + \sum_{\mathcal{E}^{V,\boldsymbol{\pi}}} g(n) \cdot \Pr\left[\mathcal{E}^{V,\boldsymbol{\pi}}\right] \\
&= \Pr\left[d(\mathcal{H}(\mathcal{S}_{\phi'}(\boldsymbol{\pi}^\star)), G) \le \delta(n) \mid \mathcal{E}^{few}\right] \cdot f(n) + g(n) \sum_{\mathcal{E}^{V,\boldsymbol{\pi}}} \Pr\left[\mathcal{E}^{V,\boldsymbol{\pi}}\right] \\
&\le 1 \cdot f(n) + g(n) \cdot 1 \\
&= f(n) + g(n) \in o(1).
\end{aligned}
$$

We now show the claims. The first claim follows from a straightforward Chernoff bound: each of the n agents places probability at least $2p$ on ending up in either π^{max} or π_{-1}, hence the expected number of agents with either π^{max} or π_{-1} is at least $2np$. Then, given that agents' rankings are sampled independently, the probability of the total number of such agents being less than half of this expectation is exponentially small in n.

We now show the second claim. Fix an arbitrary $\mathcal{E}^{V,\widetilde{\boldsymbol{\pi}}}$ and let us consider the conditional distribution of $\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi}^\star))$ conditioned on this event. Note that all proportions in the support of this distribution match on all entries except that of π^{max}. Let $\widetilde{\mathbf{h}}_{-\pi^{max}}$ be the partial proportions over all these matching entries, i.e., all rankings except π^{max}, so that $\mathbf{h}_{-\pi^{max}} = \widetilde{\mathbf{h}}_{-\pi^{max}}$ for all $\mathbf{h} \in \mathcal{H}(\mathcal{E}^{V,\widetilde{\boldsymbol{\pi}}})$.

Now, since $a_{\pi^{max}} \neq 0$, by Equation (2), there exists exactly one completion of the partial proportions $\mathbf{h}_{-\pi^{max}}$ such that the resulting proportion lies on G. Namely,

$$
p^{V,\widetilde{\boldsymbol{\pi}}} := \frac{b - \sum_{\pi \in \mathcal{L} \setminus \{\pi^{max}, \pi_{-1}\}} a_\pi \cdot \tilde{h}_\pi}{a_{\pi^{max}}}.
$$

Note that this completion need not be a valid proportion, as the value assigned to the proportion π^{max}, $p^{V,\widetilde{\boldsymbol{\pi}}}$, could need to be irrational, negative, or make all proportions add up to strictly more than one.

Now, we consider all completions of $\mathbf{h}_{-\pi^{max}}$ which assign values to the proportion of ranking π^{max} at least $\delta(n)$ larger or smaller than $p^{V,\widetilde{\boldsymbol{\pi}}}$. In particular, we will show (1) that all such proportions are at least $\delta(n)$ L_1 distance from all points on G, and (2) that we are likely to draw such a profile from $\mathcal{H}(\mathcal{S}_\phi(\boldsymbol{\pi}^\star))$ conditioned on $\mathcal{E}^{V,\widetilde{\boldsymbol{\pi}}}$.

Fix such a completion of $\widetilde{\mathbf{h}}_{-\pi^{max}}$, call it $\mathbf{h}$, so that $|h_{\pi^{max}} - p^{V,\widetilde{\boldsymbol{\pi}}}| > \delta(n)$. By this assumption, we have the following, where in the first step we replace b according to Equation (2):

$$
\begin{aligned}
\left| \sum_{\pi \in \mathcal{L}_{-1}} a_\pi h_\pi - b \right| &= \left| \sum_{\pi \in \mathcal{L}_{-1}} a_\pi h_\pi - \left(p^{V,\widetilde{\boldsymbol{\pi}}} \cdot a_{\pi^{max}} + \sum_{\pi \in \mathcal{L} \setminus \{\pi^{max}, \pi_{-1}\}} a_\pi h_\pi \right) \right| \\
&= |a_{\pi^{max}} \cdot h_{\pi^{max}} - p^{V,\widetilde{\boldsymbol{\pi}}} \cdot a_{\pi^{max}}| \\
&> |\delta(n) \cdot a_{\pi^{max}}|
\end{aligned}
$$

Since $a_{\pi^{max}}$ is maximal in magnitude over all coefficients in G, the $\delta(n)$-radius L_1-ball around $\mathbf{h}$ does not intersect G, because changing any entry of $\mathbf{h}$ by at most $\delta(n)$ cannot change $\sum_{\pi \in \mathcal{L}_{-1}} a_\pi p_\pi$ by more than $\delta(n) \cdot a_{\pi^{max}}$.

We have shown that all completions of $\mathbf{h}_{-\pi^{max}}$ assigning values of π^{max} outside $p^{V,\widetilde{\boldsymbol{\pi}}} \pm \delta(n)$ must be more than $\delta(n)$ L_1 distance from G. Next, we show that the probability of drawing such a completion is likely. We will do this by upper-bounding the probability placed on $h_{\pi^{max}}$ falling in the interval $p^{V,\widetilde{\boldsymbol{\pi}}} \pm$

$\delta(n)$ by showing its distribution is approximately normal using the Berry-Esseen bound, and showing the normal distribution does not place much mass on this small interval.

To this end, let I_i be the indicator random variable that agent $i \in V$ chooses π^{max} (if they do not choose π^{max}, they choose π_{-1}). Let $S = \sum_{i \in V} I_i$ be the random variable representing the total number of agents in V that vote for π^{max}. Note that the resulting proportion $\mathcal{H}_{\pi^{max}}$ will be in the range $p^{\overline{V}, \widetilde{\boldsymbol{\pi}}} \pm \delta(n)$ iff S is in the range $n(p^{\overline{V}, \widetilde{\boldsymbol{\pi}}} \pm \delta(n))$, an interval of size $2n\delta(n)$. We now show that the probability S is in *any* interval of size $2n\delta(n)$ is $o(1)$.

To do this, we first state the Berry-Esseen bound.

Lemma 30 (Berry-Esseen Bound [4, 13]). Let $X_1, \ldots, X_n$ be independent random variables with $\mathbb{E}[X_i] = \mu_i$, $\mathbb{E}[|X_i^2 - \mu_i|] = \sigma_i^2 > 0$, and $\mathbb{E}[|X_i^3 - \mu_i|] = \rho_i < \infty$. Let $T = X_1 + \cdots + X_n$. Let F be the cdf of $\frac{T - \sum_{i=1}^n \mu_i}{\sqrt{\sum_{i=1}^n \sigma_i^2}}$. Then, there exists an absolute constant C_1 such that

$$\sup_{x \in \mathbb{R}} |F(x) - \Phi(x)| \leq C_1 \cdot \left(\sum_{i=1}^n \sigma_i^2 \right)^{-1/2} \cdot \max_{1 \leq i \leq n} \frac{\rho_i}{\sigma_i^2}.$$

We would like to directly apply Lemma 30 to our S, where each $X_i = I_i$. We first consider the quantity σ_i^2 for each i. Note that $\textsc{min-prob}(\mathcal{S}'_\phi) \leq \Pr[I_i = 1] = 1 - \Pr[I_i = 0] \leq 1 - \textsc{min-prob}(\mathcal{S}'_\phi)$. Further, by monotonicity (Assumption 3), $\textsc{min-prob}(\mathcal{S}'_\phi) \geq \textsc{min-prob}(\mathcal{S}'_\phi)$, so $\textsc{min-prob}(\mathcal{S}_\phi) \leq \Pr[I_i = 1] \leq 1 - \textsc{min-prob}(\mathcal{S}_\phi)$. Since each I_i is a Bernoulli, each $\sigma_i^2 \geq \textsc{min-prob}(\mathcal{S}_\phi) \cdot (1 - \textsc{min-prob}(\mathcal{S}_\phi)) \geq {}^{\textsc{min-prob}(\mathcal{S}_\phi)}/_2$. This implies that $\sum_{i=1}^n \sigma_i^2 \geq n \cdot {}^{\textsc{min-prob}(\mathcal{S}_\phi)}/_2$. Further, note that since each I_i is bounded in $[0, 1]$, each $\rho_i \leq 1$. This implies the error bound of Lemma 30 is at most

$$C_1 \cdot \left(n \cdot \frac{\textsc{min-prob}(\mathcal{S}_\phi)}{2} \right)^{-1/2} \cdot \frac{2}{\textsc{min-prob}(\mathcal{S}_\phi)} = O\left(\frac{1}{\textsc{min-prob}(\mathcal{S}_\phi)^{3/2} \cdot \sqrt{n}} \right).$$

Further, the probability S is in some range of size $2n\delta(n)$ will be equivalent to the probability $\frac{S - \sum_{i=1}^n \mu_i}{\sqrt{\sum_{i=1}^n \sigma_i^2}}$ is in a specific range of size $\frac{2n\delta(n)}{\sqrt{\sum_{i=1}^n \sigma_i^2}} \leq \frac{2\delta(n)\sqrt{2n}}{\sqrt{\textsc{min-prob}(S_\phi)}}$. Let F be the CDF of this random variable. The probability this random variable is in a range of this size is at most

$$\begin{aligned}
&\sup_x \left(F\left(x + \frac{2\delta(n)\sqrt{2n}}{\sqrt{\textsc{min-prob}(S_\phi)}} \right) - F(x) \right) \\
\leq &\sup_x \left(\Phi\left(x + \frac{2\delta(n)\sqrt{2n}}{\sqrt{\textsc{min-prob}(S_\phi)}} \right) - \Phi(x) \right) + 2O\left(\frac{1}{\textsc{min-prob}(\mathcal{S}_\phi)^{3/2} \cdot \sqrt{n}} \right) \\
= &\sup_x \left(\int_x^{x + \frac{2\delta(n)\sqrt{2n}}{\sqrt{\textsc{min-prob}(S_\phi)}}} \phi(x) \right) + 2O\left(\frac{1}{\textsc{min-prob}(\mathcal{S}_\phi)^{3/2} \cdot \sqrt{n}} \right) \\
\leq &\sup_x \left(\int_x^{x + \frac{2\delta(n)\sqrt{2n}}{\sqrt{\textsc{min-prob}(S_\phi)}}} \frac{e}{\sqrt{2\pi}} \right) + 2O\left(\frac{1}{\textsc{min-prob}(\mathcal{S}_\phi)^{3/2} \cdot \sqrt{n}} \right) \\
= &\sup_x \left(\frac{e}{\sqrt{2\pi}} \cdot \frac{2\delta(n)\sqrt{2n}}{\sqrt{\textsc{min-prob}(S_\phi)}} \right) + 2O\left(\frac{1}{\textsc{min-prob}(\mathcal{S}_\phi)^{3/2} \cdot \sqrt{n}} \right)
\end{aligned}$$

$$
\begin{aligned}
&= O\left(\frac{1}{\sqrt{\textsc{min-prob}(S_\phi)}} \cdot \delta(n)\sqrt{n}\right) + O\left(\frac{1}{\textsc{min-prob}(\mathcal{S}_\phi)^{3/2} \cdot \sqrt{n}}\right) \\
&\in o(1).
\end{aligned}
$$

This is our desired $g(n)$ which completes the proof of the second claim. $\square$

E Supplemental materials from Section 6

E.1 Proof of Proposition 23

Consider the following profile $\boldsymbol{\pi}$ with 300 voters and 3 candidates:

$$\begin{array}{c|l} \text{36 voters} & a \succ b \succ c \\ \text{80 voters} & a \succ c \succ b \\ \text{115 voters} & b \succ a \succ c \\ \text{69 voters} & c \succ b \succ a. \end{array}$$

Fix an arbitrary rule $R \in PSR$ and Mallow's noise parameter $\phi \in [0,1)$. Notice that R can be represented by its scoring vector on 3 candidates, which, without loss of generality, has been scaled and translated to $(1, s, 0)$ for some $s \in [0,1]$. We will show that, with high probability as z grows large, $R(\mathcal{S}_\phi^{Mallows}(z\boldsymbol{\pi})) \neq b$ while b is a Condorcet winner.

To do this, we will find an $\varepsilon > 0$ such that for all profiles $\boldsymbol{\pi'}$ with $\|\mathcal{H}(\boldsymbol{\pi'}) - \mathbb{E}[\mathcal{H}(\mathcal{S}_\phi^{Mallows}(\boldsymbol{\pi})]\|_1 < \varepsilon$, $R(\boldsymbol{\pi'}) \neq b$ yet b is the Condorcet winner. The proposition will then follow from Lemma 3.

To show this holds, we consider $\boldsymbol{\mu} := \mathbb{E}[\mathcal{H}(\mathcal{S}_\phi^{Mallows}(\boldsymbol{\pi})]$. For shorthand, we will write xyz instead of $x \succ y \succ z$. Notice that, by Definition 21,

$$\mu_{xyz} = \frac{h^{\boldsymbol{\pi}}_{xyz}\phi^0 + (h^{\boldsymbol{\pi}}_{yxz} + h^{\boldsymbol{\pi}}_{xzy})\phi^1 + (h^{\boldsymbol{\pi}}_{yzx} + h^{\boldsymbol{\pi}}_{zxy})\phi^2 + h^{\boldsymbol{\pi}}_{zyx}\phi^3}{1 + 2\phi + 2\phi^2 + \phi^3}.$$

We next consider the margin that b pairwise beats other candidates a and c, as well as the margin by which a beats b under R. More concretely, these margins correspond to the following expressions

$$(\mu_{bac} + \mu_{bca} + \mu_{cba}) - (\mu_{abc} + \mu_{acb} + \mu_{cab})$$
$$(\mu_{bca} + \mu_{bac} + \mu_{abc}) - (\mu_{cba} + \mu_{cab} + \mu_{acb})$$
$$(\mu_{abc} + \mu_{acb} + s\mu_{bac} + s\mu_{cab}) - (\mu_{bac} + \mu_{bca} + s\mu_{abc} + s\mu_{cba}).$$

Expanding the definitions of each μ_{xyz} and simplifying (noting that all occurrences of s cancel), we find that these expressions are respectively equal to

$$1/75(1-\phi)(17 - 23\phi + 17\phi^2)$$
$$1/150(1-\phi)(1 + 116\phi + \phi^2)$$
$$1/300(1-\phi)(1+\phi)(1 + 11\phi).$$

Since $\phi \in [0,1)$, it is clear that all terms except for $(17 - 23\phi + 17\phi^2)$ are positive. It can further be checked that $(17 - 23\phi + 17\phi^2)$ is in fact positive for all values of ϕ because it has no real roots and is a convex parabola. Hence, there exists $\varepsilon > 0$ that is smaller than all these expressions. Notice that any profile $\boldsymbol{\pi'}$ with $\|\mathbf{h}^{\boldsymbol{\pi}'} - \boldsymbol{\mu}\|_1 < \varepsilon$, each of these expressions with $\mathbf{h}^{\boldsymbol{\pi}'}$ can change by at most ε, and hence remains positive. This implies the condition necessary to apply Lemma 3, as needed. □